From biodiversity response to ecological inference: context- and scale-dependent relationships between land use and biodiversity change in tropical forest landscapes

by

Juan Andrés Martínez-Lanfranco

A thesis submitted in partial fulfillment of the requirements for the degree of

Doctor of Philosophy
in
Ecology

Department of Biological Sciences
University of Alberta

© Juan Andrés Martínez-Lanfranco, 2026
This work is licensed under CC Attribution-NoDerivatives 4.0 International

ABSTRACT

Land-use change does more than remove habitat for species. It reorganises vegetation structure, climatic exposure, spatial configuration and the composition of the communities that remain. Biodiversity change is nonetheless commonly summarised through average responses, and through model coefficients whose ecological meaning is assumed rather than demonstrated. This dissertation asks when environmental context changes the biological consequences of land use, and when the structure of observational landscapes changes what those consequences can be inferred to represent.

Across 2,986 georeferenced, effort-reported assemblages spanning seven dataset-defined animal groups across five broad taxonomic groups (mammals, birds, reptiles, amphibians and butterflies) and broad climatic, topographic, and productivity gradients in the Atlantic Forest of South America, land-use effects on species richness varied strongly with ecological context. Thermal variability intensified negative land-use effects, most consistently among forest specialists. That contrast is consistent with the loss of climatic buffering when forest is removed, although canopy microclimate was not measured directly. Heterogeneity carried no universal biodiversity value either. Persistent geophysical variation supported higher baseline richness, whereas remotely sensed vegetation greenness variability compounded cropland effects, showing that different heterogeneity axes can carry contrasting ecological meanings. Guild responses opposed one another along the heterogeneity and productivity axes, so aggregation obscured moderation that was present within guilds.

The same land-use transformation that alters climatic and vegetational context jointly changes how much habitat remains and how it is spatially arranged. Fragmentation therefore provides a

particularly demanding case for asking whether correlated landscape changes can be separated strongly enough to support the ecological interpretation assigned to a model coefficient. The problem appears in two contrasting fragmentation cases, a global multi-taxa forest dataset and an Atlantic Forest tree beta-diversity study. In the global case, habitat amount and patch number remained nonlinearly coupled in the realised predictor space, pushing the patch-number coefficient towards zero while conventional collinearity screening gave little warning. In the Atlantic Forest case, a supported negative patch-number coefficient documented a real subdivision association without identifying the mechanism invoked. Contrasting coefficient outcomes led to the same inferential lesson. What a model estimates is not automatically what an ecological claim requires.

Competing fragmentation frameworks imply different response surfaces along the habitat-amount gradient. The habitat amount hypothesis predicts invariance once amount is controlled. Scarcity and threshold dynamics predict stronger costs where habitat is limited. Coexistence can generate sign reversal, whereas extinction filtering predicts stronger detectable effects where sensitive taxa remain. Across Atlantic Forest vertebrate analyses spanning multiple forest definitions, configuration metrics, spatial scales and predictor formulations, forest amount remained the dominant landscape component. Configuration associations were smaller but became more negative towards higher forest amount, most consistently among forest specialists. The overall geometry was most consistent with extinction filtering, although some supported endpoint patterns departed from its prospective prediction, and cross-sectional data and reduced recoverability at low cover prevent unique identification of the mechanism.

Biodiversity responses to land-use change are therefore conditional twice over. Ecological context governs which responses are expressed, and inferential context governs which of those responses can be recovered and interpreted. Assemblage composition determines which signals remain in the

response and whether aggregation preserves them. Spatial scale determines which processes enter the analysed grain and extent, and therefore which contrasts an observational analysis can recover. Aggregate richness can preserve shared land-use damage while missing the contexts and assemblage components in which that damage intensifies. Forest retention and restoration remain the clearest conservation priorities supported by the Atlantic Forest data, and nothing here justifies trading forest amount for configuration. Even so, configuration was consistently associated with an additional biodiversity cost after accounting for forest amount, particularly for forest specialists and in better-forested landscapes.

Interpreting land-use effects therefore requires knowing which species remain, the spatial scale at which responses are observed, and whether the available landscape contrasts identify the ecological process of interest before a fitted coefficient can be interpreted as evidence for the process it is meant to represent.

PREFACE

This thesis is an original work by Juan Andrés Martínez-Lanfranco (ORCID 0000-0002-1692-1168), licensed under CC Attribution-NoDerivatives 4.0 International (CC BY-ND 4.0).

Generative artificial intelligence (AI) tools (Claude by Anthropic and ChatGPT by OpenAI) were used during the preparation of this thesis to assist with coding, editorial review of my own drafted prose for clarity and concision, and structural assembly of the final document. No part of the scientific content of this thesis was generated by AI.

I remained responsible for the generation of ideas, research questions, hypotheses and predictions, data compilation and standardization, study design, statistical analyses and modelling, interpretation of results, inference, and the scientific claims presented here. I reviewed and verified all AI-assisted output, and I take full responsibility for the entire contents of this thesis.

The research presented in this thesis involved no new fieldwork, handling, experimental manipulation, or sampling of live animals. The empirical analyses used existing biodiversity and environmental data obtained from freely available open-access data repositories and other publicly available sources, which are cited in the relevant chapters.

No chapter of this thesis has been formally published in the form presented here.

Chapters 1, 2 and 5 have not yet been prepared or submitted as manuscripts. I led the conceptual development, conducted the analyses, and wrote these chapters. My supervisor, Erin M. Bayne, contributed to conceptual development, review and editing, and supervision throughout. I expect these chapters to form the basis of future co-authored manuscripts.

Chapter 3 is a synthesis that integrates, substantially rewrites, and extends material developed in three sole-authored accepted and in-review manuscripts, posted as preprints:

- Martínez-Lanfranco, J. A. (2026). Habitat amount control is necessary but not sufficient to resolve the fragmentation debate. *Conservation Biology*, accepted. *EcoEvoRxiv*: https://doi.org/10.32942/X2KT05

- Martínez-Lanfranco, J. A. (2026). Habitat-loss-driven predictor coupling limits inference about the independent effects of configuration in additive habitat-amount models: implications for the fragmentation debate. *bioRxiv*: https://doi.org/10.64898/2026.04.12.718042

- Martínez-Lanfranco, J. A. (2026). Configuration coefficients can mislead fragmentation-based conservation inference from beta diversity. *EcoEvoRxiv*: https://ecoevorxiv.org/repository/view/14049/

Chapter 4 corresponds to the following in-review manuscript, also posted as a preprint. I led the conceptual development and wrote the manuscript. Erin M. Bayne contributed to conceptual development, and review and editing. The chapter is substantially the same as that manuscript, differing only in formatting and in the referencing conventions required by the thesis.

- Martínez-Lanfranco, J. A., & Bayne, E. M. 2026. Reframing the habitat fragmentation debate around the inferential targets predicted by ecological theory. *EcoEvoRxiv*: https://ecoevorxiv.org/repository/view/13012/

To my parents and best friends, Alejandro and Silvana,
who have been my guiding lights all along the way.

"Science is more than a body of knowledge. It is a way of thinking; a way of skeptically interrogating the universe with a fine understanding of human fallibility."
— Carl Sagan

ACKNOWLEDGEMENTS

This work was supported in part by the CALDO-National Agency for Research and Innovation (*Agencia Nacional de Investigación e Innovación*, ANII, Uruguay) partnership, the BIOS^2 Computational Biodiversity Science and Services Program [NSERC's Collaborative Research and Training Experience (CREATE) Program], and the Bayne Lab (Department of Biological Sciences, University of Alberta). Colleagues in the Bayne Lab and BIOS^2 fellows provided valuable discussions and feedback throughout my program of study at the University of Alberta.

I am thankful to my supervisor, Erin M. Bayne, and to my committee members, Rolf Vinebrooke and Scott Nielsen, for their guidance throughout this work. I thank James Cahill and Fangliang He, who served on my candidacy examination, and Guillaume Blanchet and Fangliang He, my external examiners at the thesis defence; the discussions in both sharpened the arguments developed here. I owe particular thanks to Erin for his trust and for the freedom to follow my own questions and ideas, and to Shelley Scott in the graduate office and Georgina Lazarini (ANII), whose help with questions and program formalities I relied on throughout.

Finally, I am indebted to my family, without whose continued support I would not have made it this far: to my parents, who always supported me in pursuing my passions; to my brothers, Alejandro and Agustín; and to my nephew Euse and my nieces Dani, Cami, Pauli and Gabi, who are a reminder of what's truly important in life. And to my kitties, Nalah and Callie, who brought much joy when I needed it most.

TABLE OF CONTENTS

Introduction 1

Land-use change as climatic, structural, and spatial reorganisation 2

Shared Atlantic Forest biodiversity data and study design 5

Where the observations come from 5

The spatial design 7

What happens to a source assemblage 7

Two biodiversity universes 8

What the responses are 8

Predictor geometry and inference from observational gradients 9

Aims and chapter structure 11

References 14

Chapter 1. Amplitude, not extremes: thermal variability amplifies land-use effects on species diversity in the Atlantic Forest 29

Introduction 29

Hypotheses and predictions 33

Methods 33

Results 43

Discussion 46

Conservation implications 52

Conclusions 52

References 53

Chapter 2. Not all heterogeneity is opportunity: its biodiversity value depends on origin and metric in the transformed Atlantic Forest 76

Introduction 76

Hypotheses and predictions 79
Methods 80
Results 92
Discussion 96
Conservation implications 104
Conclusions 104
References 106
Chapter 3. Habitat amount control is necessary but not sufficient: sign-symmetric diagnostics of fragmentation-per-se inference 132
Introduction 132
Methods 136
Results 143
Discussion 158
Conclusions 165
References 166
Chapter 4. Reframing the habitat fragmentation debate around the inferential targets predicted by ecological theory 180
The debate and its structure 180
Ecological mechanisms generating conditional fragmentation effects 183
From ecological mechanism to statistical identifiability 188
Translating mechanism into testable predictions 195
Four informative outcomes 201
Implications for fragmentation research 203
Conclusions 204
References 206

Chapter 5. Where forest remains, configuration matters most: amount-conditioned associations in Atlantic Forest vertebrates 218
Introduction 218
Hypotheses and predictions 221
Methods 223
Results 235
Discussion 243
Conservation implications 253
Conclusions 254
References 256
Chapter 6. Putting it all together: coupled ecological and inferential conditionality in biodiversity responses to land-use change 275
Ecological conditionality in Atlantic Forest assemblages 275
Thermal variability and forest dependence 276
Heterogeneity origin, energy, and response aggregation 277
Inferential conditionality in fragmentation research 280
From additive effects to conditional response surfaces 284
Joint application in the Atlantic Forest 286
What ecological and inferential conditionality imply 288
Shared inferential boundaries 288
Future directions 292
Conservation implications 293
Conclusions 295
References 295
Complete list of literature cited 306
Appendix 0.S1. Species assemblages and spatial aggregation 331

Appendix 0.S2. Thesis-wide glossary of recurring terms 344
Appendix 1.S1. Chapter-specific response documentation 355
Appendix 1.S2. Predictor definitions, data provenance, and predictor geometry 357
Appendix 1.S3. Predictor roles and support-classification rules 369
Appendix 1.S4. Model specification, model-building diagnostics, and adequacy 374
Appendix 1.S5. Land-use × climate interaction results and bootstrap validation 386
Appendix 2.S1. Chapter-specific response documentation 397
Appendix 2.S2. Environmental predictors, provenance, and geometry 398
Appendix 2.S3. Predictor roles and support-classification rules 415
Appendix 2.S4. Model specification, model-building diagnostics, and adequacy 421
Appendix 2.S5. Land-use × heterogeneity interaction results and bootstrap validation 433
Appendix 3.S1. Diagnostic terminology used in Chapter 3 443
Appendix 5.S1. Chapter-specific response documentation 447
Appendix 5.S2. Forest definitions, fragmentation metrics, and predictor geometry 450
Appendix 5.S3. Mechanism classification, eligibility, and evidential rules 467
Appendix 5.S4. Model specification, model building, and adequacy 476
Appendix 5.S5. Fragmentation × amount results and robustness validation 487

LIST OF TABLES

Table 1.1. Predictions tested in this study, with hypothesised mechanisms and alternative hypotheses. 60

Table 1.2. Focal land-use × climate interactions across forest-dependence strata. 63

Table 2.1. Predictions tested in this study, with hypothesised mechanisms and alternative hypotheses. 113

Table 2.2. Focal land-use × heterogeneity and land-use × energy interactions across forest-affiliation strata. 119

Table 3.1. Cross-case diagnostic synthesis. 177

Table 4.1. Limiting-case predictions of four fragmentation frameworks 212

Table 5.1. Predicted fragmentation–amount interaction signatures 262

Table 5.2. Formulation-separated summary of additive and conditional fragmentation results across richness strata. 265

Table 6.1. Shared inferential conditions across the Atlantic Forest empirical chapters and the conclusions retained after each condition is applied. 302

Table 0.S1.1. Curated assemblages by animal group and sampling method 341

Table 0.S1.2. Observed species richness across aggregation scales by animal group and sampling method 342

Table 1.S1.1. Univariate total-versus-restricted-universe benchmark by predictor and scale 355

Table 1.S2.1. Predictor set, operational meaning, and analytical role. 358

Table 1.S2.2. Predictor specifications: source product, temporal anchor, native resolution. 359

Table 1.S2.3. Correlation stability across the 18 grain × extent combinations (nine retained predictors), sorted by mean strength. 361

Table 1.S2.4. Variance inflation factors by model configuration (representative scale). 362

Table 1.S2.5. Strongest predictor–predictor nonlinearity at the representative scale 364

Table 1.S3.1. Predictor families, ecological mechanisms, expected roles, and key references. 369

Table 1.S4.1. Five-stratum effort-formulation audit..... 376

Table 1.S4.2. Selected polynomial structure by stratum (counts over the 144 interaction models per stratum). 377

Table 1.S4.3. Selected random-slope structure by stratum (counts over the 144 interaction models per stratum)..... 379

Table 1.S4.4. Residual-diagnostic adequacy by response stratum 381

Table 1.S5.1. Five-stratum primary interaction support 387

Table 1.S5.2. Forest-specialist primary and controlled interaction support detail 388

Table 1.S5.3. Parametric-bootstrap primary-tier interval envelopes for the load-bearing forest-specialist interactions. 391

Table 1.S5.4. Bootstrap validation: target models whose 95% interval excludes zero, by interaction (forest-specialist stratum). 392

Table 1.S5.5. Provenance of the six load-bearing forest-specialist rows of the main-text Table 1.2: per interaction, the primary support count (n/18), the nested controls-robust count (n/n_primary), and the mean… 394

Table 2.S2.1. Predictor operationalisation, provenance, and design status. 398

Table 2.S2.2. Variance inflation by predictor configuration. 402

Table 2.S2.3. Structure-coefficient classification of the five focal predictors..... 405

Table 2.S2.4. Strongest predictor–predictor nonlinearity at the representative combination..... 407

Table 2.S3.1. Criteria used to classify and to set aside interaction surfaces..... 419

Table 2.S4.1. Selected polynomial structure by stratum (counts over the 180 interaction models per stratum). 423

Table 2.S4.2. Selected random-slope structure by response stratum, counted over the 180 selected cells per stratum drawn from the 720 additive fits..... 423

Table 2.S4.3. Residual-diagnostic adequacy by response stratum (scaled-residual tests at $\alpha = 0.05$). 425

Table 2.S4.4. Residual spatial autocorrelation before and after the spatial autocovariate, by response stratum and predictor family. 426

Table 2.S4.5. Aggregation-cancellation fractions by process 429

Table 2.S5.1. Five-stratum primary interaction support 434

Table 2.S5.2. Forest-specialist controlled interaction support detail 436

Table 2.S5.3. Proxy-swap survival for energy-conditioned interactions 437

Table 2.S5.4. Parametric-bootstrap interval envelopes for the load-bearing interactions. 438

Table 2.S5.5. Provenance of the load-bearing rows of Table 2.2 439

Table 5.S1.1. Classified taxon pools and forest-affiliation composition. 448

Table 5.S2.1. Forest raster streams: product, resolution, temporal anchor and provenance, with sources given in the final column 450

Table 5.S2.2. Fragmentation metric definitions and orientation (all oriented so larger = more fragmented). 452

Table 5.S2.3. Retained-metric geometry, in two panels. 455

Table 5.S2.4. Residualization validation and amount-metric geometry 457

Table 5.S2.5. Control geometry. 459

Table 5.S3.1. Operational classification: how an estimated interaction is assigned to a mechanism class 468

Table 5.S3.2. Support and eligibility fields and the primary-support rule. 471

Table 5.S3.3. Design dimensions that condition the mapping from observed interaction signatures to ecological interpretation. 473

Table 5.S4.1. Polynomial-form selection by response and formulation. 478

Table 5.S4.2. Random-slope candidate selection across the four-candidate set, by response stratum and formulation. 479

Table 5.S4.3. Simulation-based residual diagnostics, forest-specialist richness. 481

Table 5.S4.4. Targeted environmental-control accounting. 483

Table 5.S5.1. Design, attempt, eligibility, and failure accounting, by stratum and formulation. 487

Table 5.S5.2. Primary interaction support by stratum × metric, evaluated within each formulation (not pooled). 488

Table 5.S5.3. Additive forest-amount and fragmentation coefficients by response stratum, fragmentation metric and formulation. 495

Table 5.S5.4. Robustness across product, resolution, threshold, scale, and formulation, matched within cell..... 497

Table 5.S5.5. Guild-specificity matched contrasts. 500

Table 5.S5.6. Environmental-control sensitivity, targeted coverage: the 1,248 primary-supported interaction cells. 502

Table 5.S5.7. Bootstrap validation and main-text provenance. 504

Table 5.S5.8. Direction of bootstrap interval support for the forest-specialist marginal effect at the lower cover percentiles, among primary-supported cells (214 raw, 197 residualized). 505

Table 5.S5.9. Cells whose interval was positive at the 5th cover percentile and negative at the 95th, by stream, metric and formulation, over primary-supported cells..... 506

Table 5.S5.10. Threshold, response-grain and predictor-extent profiles of the opposite-signed endpoint pattern, over primary-supported specialist cells. 507

Table 5.S5.11. Paired within-cell ratio of the residualized to the raw coefficient magnitude, across the 648 forest-specialist design cells matched between formulations. 508

LIST OF FIGURES

Figure 0.1. The Atlantic Forest study region in South America 22

Figure 0.2. Cropland cover across the Atlantic Forest..... 23

Figure 0.3. Geographic distribution of the curated species assemblages across the Atlantic Forest 24

Figure 0.4. Hexagonal response units at the three response grains 25

Figure 0.5. Sample-size attrition by animal group across aggregation stages 26

Figure 0.6. Forest-affiliation composition of the habitat-classifiable taxon pool..... 27

Figure 0.7. Conceptual structure of the dissertation 28

Figure 1.1. Geographic distribution of the curated species assemblages (n = 2,986) across the Atlantic Forest (AF). 68

Figure 1.2. Hypothesised mechanism, competing climatic predictions, response strata and analytical design..... 69

Figure 1.3. Land-use main-effect contrasts across forest-dependence strata..... 71

Figure 1.4. All fitted land-use × climate interaction coefficients across forest-dependence guilds. 72

Figure 1.5. Scale coherence of forest-specialist thermal-variability amplification. 74

Figure 1.6. Climatic moderation of the cropland effect on the response scale 75

Figure 2.1. Study area and assemblages. 124

Figure 2.2. Why heterogeneity is ambiguous, prospective moderation pathways, response strata and analytical design..... 125

Figure 2.3. Cropland × greenness-magnitude compounding and its guild specificity. 127

Figure 2.4. Energy conditioning of the specialist cropland × greenness-magnitude surface. 129

Figure 2.5. Process-specific cancellation under aggregation. 131

Figure 3.1. Diagnostic architecture for fragmentation-per-se inference. 170

Figure 3.2. Realised habitat amount–patch number geometry in the two diagnostic case studies. 172

Figure 3.3. Coefficient architecture of habitat amount and patch number across 18 analytical variants (α and γ diversity) in the GS25–F26 case. 174

Figure 3.4. Mechanism triangulation for fragmentation-per-se inference from beta diversity in the K26 Atlantic Forest dataset. 176

Figure 4.1. Conceptual response-surface framework for four competing fragmentation theories. 214

Figure 4.2. Statistical translation of the canonical alternatives in Fig. 4.1. 216

Figure 5.1. Study area and assemblages. 267

Figure 5.2. The estimand, prospective theoretical signatures, response strata and design space 268

Figure 5.3. Additive fragmentation and forest-amount coefficients by metric and response stratum 270

Figure 5.4. Specialist interaction coefficient across predictor extent. 272

Figure 5.5. Conditional fragmentation effect by metric and formulation. 273

Figure 5.6. Interaction coefficient by response stratum and fragmentation metric. 274

Figure 0.S1.1. Geographic and temporal distribution of curated species assemblages across the Atlantic Forest 338

Figure 0.S1.2. Sample-size attrition by sampling method across aggregation stages. 339

Figure 0.S1.3. Spatial distribution of curated assemblages after aggregation to 40-km units 340

Figure 1.S2.1. Land-use surrogates across the Atlantic Forest 366

Figure 1.S2.2. Climate moderators across the Atlantic Forest 367

Figure 1.S2.3. Environmental controls across the Atlantic Forest 368

Figure 1.S4.1. Model-building and inference sequence for the climatic-moderation analysis ... 385

Figure 1.S5.1. Resolved interaction atlas across all five response strata. 395

Figure 1.S5.2. Land-use main-effect contrasts resolved by scale. 396

Figure 2.S2.1. Land-use surrogates across the Atlantic Forest 409
Figure 2.S2.2. Vegetational heterogeneity across the Atlantic Forest 410
Figure 2.S2.3. Geophysical heterogeneity across the Atlantic Forest 411
Figure 2.S2.4. Realised water–energy availability across the Atlantic Forest 412
Figure 2.S2.5. Predictor geometry. 413
Figure 2.S4.1. Model-building and inference sequence for the heterogeneity-moderation analysis 432
Figure 2.S5.1. Resolved interaction-support atlas. 441
Figure 5.S2.1. Forest-cover streams across the Atlantic Forest 461
Figure 5.S2.2. Forest definitions at a representative landscape 463
Figure 5.S2.3. Environmental controls across the Atlantic Forest 464
Figure 5.S2.4. Amount–configuration geometry and residualization validation 465
Figure 5.S4.1. Model-building and inference sequence for the amount-conditioned fragmentation analysis 486
Figure 5.S5.1. Marginal fragmentation effect across the forest-cover gradient at a fixed display cell, by response stratum and metric 509
Figure 5.S5.2. Primary-support frequency across predictor extent, by metric and formulation. 511
Figure 5.S5.3. Additive fragmentation coefficient in every fitted forest-specialist design cell. 513
Figure 5.S5.4. Interaction coefficient in every fitted forest-specialist design cell. 515
Figure 5.S5.5. Ordered guild gradient in the additive fragmentation coefficient. 516

INTRODUCTION

Human land-use change reshapes contemporary assemblages by reducing native habitat while altering the spatial arrangement of what remains, vegetation structure and quality, matrix conditions, the local climate experienced by organisms, and the set of species able to persist (Meyer and Turner 1992; Foley et al. 2005; Fischer and Lindenmayer 2007; Newbold et al. 2015). These changes reshape the abiotic and spatial template of community assembly across scales, linking local richness responses to broader questions about habitat specificity, dispersal, compositional turnover, and conservation intervention (Chase et al. 2020). A central challenge is to determine when, where, and through which ecological contexts those effects are amplified, buffered, obscured, or misidentified.

Land-use effects rarely have a single expression across landscapes. Community ecology has long emphasised that local assemblages cannot be interpreted independently of broader spatial and historical context (Ricklefs 1987, 2008; Pärtel et al. 2016). Landscape ecology and macroecology add that ecological processes and conservation threats operate across spatial domains that rarely align with a single sampling grain (Wiens 1989; Levin 1992; Teng et al. 2020; Banks-Leite et al. 2022). Landscape moderation frames the ecological questions of this dissertation because context can alter the strength, direction and consequences of biodiversity responses (Tscharntke et al. 2012).

The scale at which richness is aggregated need not match the scale at which land use, climate, heterogeneity, productivity, or forest configuration acts (Wiens 1989; Levin 1992; Jackson and Fahrig 2015). Local habitat structure, canopy-mediated microclimate, edge exposure, dispersal limitation, and regional environmental gradients can operate across different extents, so a land-use effect may appear weak, strong, or absent depending on the response grain and predictor extent being evaluated. Local richness and compositional turnover can also respond differently to the same anthropogenic driver across spatial and organisational scales (Chase et al. 2020).

We treat land-use change as a compound reorganisation of habitat amount, spatial configuration, vegetation structure, matrix conditions, climatic exposure, and assemblage composition rather than as a single environmental gradient. The central question is which

component of that transformation carries a biodiversity signal, under which ecological contexts, at which spatial scale, and for which response component.

Land-use change as climatic, structural, and spatial reorganisation

Tropical forests are among the most species-rich systems on Earth and have been strongly affected by land-use transformation (Laurance et al. 2012; Edwards et al. 2019). Forest conversion typically replaces structurally complex native vegetation with agriculture, pasture, plantation, urban cover, secondary vegetation, or degraded matrix environments. This transformation reduces the amount of forest, subdivides what remains into smaller and more isolated units, increases edge exposure, changes matrix conditions, and erodes forest quality (Fischer and Lindenmayer 2007; Haddad et al. 2015; Alroy 2017). It also acts as a directional filter because species differ in forest dependence, matrix tolerance, and capacity to use modified environments (Devictor et al. 2008, 2010; Clavel et al. 2011; Newbold et al. 2018; Donald et al. 2019).

Habitat conversion can intensify or redirect biodiversity loss through its climatic consequences. Forest canopies buffer temperature and humidity, reducing short-term thermal extremes beneath the canopy. Clearing, degradation, and fragmentation weaken this buffering capacity and expose organisms to warmer or more variable microclimates (Ewers and Banks-Leite 2013; De Frenne et al. 2019; Williams and Newbold 2020). This pathway is especially relevant in tropical forests, where many species are forest-associated and may have narrow thermal safety margins or limited physiological flexibility under climatic extremes (Khaliq et al. 2014; Scheffers et al. 2016; Newbold et al. 2020). Interactions between land use and climate drive biodiversity change, but the local climatic consequences of forest degradation require attention alongside broad-scale climate change (Brook et al. 2008; Mantyka-Pringle et al. 2012; Oliver and Morecroft 2014).

Environmental heterogeneity can also alter land-use effects. Niche-based explanations have long linked environmental heterogeneity to higher species richness by expanding environmental space and opportunities for resource partitioning (Hutchinson 1957; MacArthur 1972; Tews et al. 2004; Stein et al. 2014). Yet heterogeneity is not a single ecological quantity (Stein and Kreft 2015). Because a landscape of fixed area divided among more environmental

states leaves less of each, greater heterogeneity can reduce the effective area available per state and increase extinction risk for species with narrow niches or limited dispersal (Kadmon and Allouche 2007; Allouche et al. 2012; Laanisto et al. 2013; Ben-Hur and Kadmon 2020). Its origin matters as well. Persistent geophysical heterogeneity can provide environmental opportunity and refugial conditions (Hjort et al. 2015), whereas land-use-responsive vegetational heterogeneity can reflect land-cover replacement, edge proliferation, degraded forest structure, and secondary-regrowth mosaics (Seiferling et al. 2014; Heidrich et al. 2020).

Forest fragmentation is a related but more specific expression of landscape reorganisation. Habitat fragmentation per se refers to effects of spatial configuration at a given amount of habitat, yet its role in biodiversity loss remains one of the most contested issues in landscape ecology (Fahrig 2017; Fletcher et al. 2018; Fahrig et al. 2019; Miller-Rushing et al. 2019; Riva et al. 2024). The controversy persists partly because habitat amount and configuration are difficult to separate in real landscapes. Habitat loss co-generates spatial subdivision, and many configuration metrics covary nonlinearly with habitat amount (Ewers and Didham 2006; Didham et al. 2012; Neel et al. 2004; Wang et al. 2014). Total richness can also mask declines of forest-dependent species when matrix-tolerant species persist or colonise (Newbold et al. 2018). Fragmentation effects are also expected to vary along the habitat-amount gradient. Competing frameworks place their strongest or most detectable effects in different domains of that gradient, ranging from low-cover scarcity and threshold conditions to higher-cover landscapes where sensitive taxa remain or coexistence processes can operate (Andrén 1994; Villard and Metzger 2014; Betts et al. 2019; Rybicki et al. 2020; Saura 2021; Fahrig 2021; Zhang et al. 2024).

These difficulties mean that fragmentation per se is not identified by the sign of a configuration coefficient alone. Controlling for habitat amount is necessary (Fahrig et al. 2026) but may remain insufficient when the realised predictor space does not support the intended contrast (Martínez-Lanfranco 2026a). Identification also requires a response estimand that matches the mechanism and a configuration metric that represents the proposed process. Structural configuration is not equivalent to functional connectivity, because the same mapped arrangement can represent different habitat accessibility depending on species' movement

capacities, matrix permeability, and the spatial scale at which organisms integrate multiple patches (Taylor et al. 1993; Bélisle 2005; Villard and Metzger 2014).

The ecological questions developed above concern where land-use effects change across environmental contexts, assemblage components and spatial scales. Fragmentation raises a further question of interpretation, namely when an estimated association represents the ecological contrast it is used to describe. These are two forms of conditionality. Ecological conditionality concerns where and for whom land-use effects change, whereas inferential conditionality concerns when a fitted quantity represents the contrast of interest. The General Discussion returns to where they meet. We approach that question through the response analysed, the realised predictor space, the landscape metric, the model structure and the domain over which an estimate is interpreted. Assemblage composition and spatial scale therefore remain explicit design dimensions throughout the dissertation rather than fixed background choices.

The Atlantic Forest of South America is the empirical anchor of this dissertation. The biome is not a single uncontested spatial unit. Eleven delimitation proposals and four sectorisation schemes define different analytical domains, so statements about its extent and remaining vegetation require explicit geographic and vegetation definitions (Marques et al. 2021). Across the integrative delimitation reviewed by Marques et al. (2021), the biome spans roughly 31 degrees of latitude, extends from sea level to 2,891 m and contains eight Köppen–Geiger climate types (Figure 0.1).

Present landscapes reflect successive social–ecological regimes rather than one recent episode of forest conversion, as extractive, agricultural, pastoral, mining and industrial land uses repeatedly reorganised forest amount, structure and matrix conditions (Solórzano et al. 2021). Most remaining forest occurs in small, edge-exposed remnants, and only a small fraction is formally protected (Ribeiro et al. 2009; Vancine et al. 2024; Sloan et al. 2014). These remnants are embedded in anthropogenic matrices that vary sharply in structure and permeability, making matrix contrast and edge exposure part of the ecological context in which forest amount and configuration operate (Tabarelli et al. 2010; Vancine et al. 2024). Severe transformation alongside pronounced topographic, climatic and historical variation supports comparisons of the

mechanisms that condition, amplify, and obscure biodiversity responses across contexts and scales (Figure 0.2).

Shared Atlantic Forest biodiversity data and study design

The three Atlantic Forest chapters analyse one body of biodiversity data under one spatial design. We set out both here, so that those chapters document only what changes between them. Across the empirical chapters, land-cover composition is represented by cropland prevalence and non-natural cover in Chapters 1 and 2 and by forest cover in Chapter 5, while land-cover configuration is analysed explicitly in Chapter 5. Appendix 0.S1 reports the sources, classifications and summaries in full, and Appendix 0.S2 defines the terms that recur across chapters.

Species richness provides the common response axis across the three Atlantic Forest empirical chapters and the principal empirical throughline of the dissertation. We model richness within each response unit, so the diversity component being estimated is defined by the response grain rather than fixed a priori: at fine grains it approximates local assemblage richness, and at coarser grains it approaches the richness of the pool consolidated within the unit (Whittaker 1972, Anderson et al. 2011). Richness alone cannot identify turnover, abundance compensation, trait-mediated filtering, or the species-level processes underlying community change. It can, however, be harmonised across taxa, sampling designs, response grains and landscape extents in large observational datasets, which makes it the empirical throughline of the dissertation.

Where the observations come from

We compiled 2,986 species assemblages from published open-access data papers, covering seven dataset-defined animal groups: mid-large mammals, small mammals, bats, passerines, snakes, anurans and fruit-feeding butterflies (Figure 0.3). These are the groupings the source data papers themselves use, and they represent five broad taxonomic groups. Every assemblage carries georeferenced coordinates, a standardised sampling protocol and reported survey effort. We restricted the compilation to assemblages sampled from 1980 onward, because earlier records are sparse and unevenly distributed among groups; the retained data span 1981 to 2018. No single group covers that interval: group-level sampling windows run from 26 to 36 years, begin between 1981 and 1991, and end between 2015 and 2018. The pooled range is

therefore the union of seven shorter and unequal windows rather than a period over which every group was sampled. We excluded opportunistic and presence-only records. Sampling methods differ among these groups, and sometimes within them, because protocols, effort units and survey scope all vary, so we retain sampling method as a second organising level alongside animal group, giving sixteen animal-group by sampling-method combinations. These are assemblage richness values observed under standardised but heterogeneous survey protocols, not exhaustive inventories of the species present in each landscape. Thus, richness is an observed count rather than a rarefied or coverage-standardised estimate. Sampling effort enters the models as a covariate and sampling method as a grouping factor, so inference is conditional on reported effort and broad method-level differences. These adjustments do not establish complete detection or model variation in species detection probability that may itself covary with environmental conditions.

Observed richness is therefore not comparable in absolute terms among sampling methods. Including sampling method as a grouping factor accounts for average method-level differences before the focal environmental associations are estimated, but it does not place methods on a common detection scale or test whether those associations are invariant among them. Method-dependent detection that covaries with environmental gradients remains a limitation, and inference concerns relative associations in observed richness after effort and method adjustment rather than absolute richness differences among protocols.

Source assemblages also differ in how long they were sampled. Some represent a single survey period and others aggregate several years, and where a source dataset reported a multi-year period without year-specific resolution we assigned its endpoint year. The number of years contributing to an assemblage is not retained in the modelling data and does not enter the models, so sampling effort is represented by the number of source assemblages consolidated into each record and by their reported effort, not by the temporal span over which those surveys accumulated. Two consequences follow. Assemblages of equal reported effort but unequal temporal span enter as equivalent observations, and if temporal span were itself associated with land use or climate — for instance if long-running sites were concentrated in particular landscapes — that association would not be adjusted for. We cannot evaluate this from the

frozen data and therefore treat it as a limitation on the precision of the effort adjustment rather than as a resolved issue.

The spatial design

Ecological processes and analytical resolution need not operate at matching scales, so we separate them by design. The response grain is the scale at which we count species; the predictor extent is the scale at which we describe the landscape surrounding them. We aggregate richness into hexagonal response units equivalent in area to circles of 10, 25 and 40 km radius (Figure 0.4), representing local, intermediate and broader landscape contexts of approximately 314, 1,963 and 5,027 km². The finest grain sets the functional grain of the design, the smallest extent at which predictors can meaningfully be extracted given the spatial precision of the source assemblages, and the coarsest retains enough records for the guild-resolved strata to remain estimable. Environmental predictors are measured within concentric extents of 10, 25, 40, 55, 70, 85 and 100 km radius, in 15 km increments, and predictor extent is never smaller than response grain, because biodiversity responses are best explained by predictors measured at matching or broader scales (Jackson and Fahrig 2015; Miguet et al. 2016).

We require the extent to be at least as large as the grain, which yields eighteen grain-by-extent combinations: seven at the 10 km grain, six at 25 km and five at 40 km. We evaluate every effect across all eighteen rather than at one chosen resolution. These combinations are nested, overlapping windows on the same assemblages, so agreement among them indicates coherence across scales rather than independent replication. The hexagons serve as an aggregation template that determines which assemblages combine into one record, but they do not define where predictors are measured. Predictor extents are centred on the centroid of the assemblages aggregated into each record rather than on the hexagon's geometric centre, so the predictor window is anchored to where sampling actually occurred.

What happens to a source assemblage

A source assemblage is not a modelling observation. For the year-aggregated responses used in the reported analyses, several source assemblages can fall within the same spatial unit, so we consolidate assemblages of the same animal group sampled by the same method within each response unit into a single modelling record. Modelling records therefore represent animal group

by sampling method by spatial unit combinations rather than individual field sites. The 2,986 source assemblages yield 1,689 modelling records at the 10 km grain, 1,306 at 25 km and 1,069 at 40 km. Records and locations are also distinct counts: at the 10 km grain those 1,689 records occupy 1,651 sampled locations, because a location surveyed for more than one group, or by more than one method, contributes more than one record. Consolidation ratios vary by group, from roughly 1.5 source assemblages per record in mid-large mammals to 2.2 in songbirds, which reflects how densely we sampled each group relative to the grain rather than any differential exclusion (Figure 0.5).

Two biodiversity universes

Several analyses require knowing whether a species depends on forest. We assign forest affiliation from species-level habitat records in the IUCN Red List, which is possible for five of the seven groups, so two universes follow. The full universe covers all seven dataset-defined animal groups and all 2,986 assemblages. The habitat-classifiable universe comprises 2,653 of those same assemblages from five vertebrate groups, namely mid-large mammals, small mammals, bats, passerines and anurans, together with fourteen of the sixteen animal-group by sampling-method combinations. Snakes and fruit-feeding butterflies drop out for different reasons: snakes reach us as published richness totals rather than species lists, so no snake can be matched to a habitat record, whereas butterflies are identified to species, but Red List assessment of the group remains incomplete. The habitat-classifiable universe is a subset of the full universe rather than a separate compilation. Under the same aggregation, those 2,653 assemblages yield 1,360 modelling records at the 10 km grain, 1,020 at 25 km and 833 at 40 km, occupying 1,323 sampled locations at the 10 km grain.

What the responses are

Within the habitat-classifiable universe we assign 1,199 taxa to three forest-affiliation classes: 302 forest specialists, 702 forest-facultative taxa and 195 non-forest taxa. The three classes are mutually exclusive, so they partition restricted total richness exactly: restricted total equals specialist plus facultative plus non-forest richness, assemblage by assemblage (Figure 0.6). Chapters 1 and 2 analyse five response strata, adding full-universe total richness to the restricted total and the three guilds. Chapter 5 analyses four, because it requires forest affiliation throughout and therefore carries no full-universe aggregate. We compare aggregate and guild-

resolved responses because aggregate richness can conceal opposing changes. Losses of forest-associated species can be offset by gains of tolerant or generalist species, leaving totals stable while composition shifts (McKinney and Lockwood 1999; Matthews et al. 2014; Newbold et al. 2018). The three guilds form a habitat-affiliation partition rather than a measured trait axis. Through them, each chapter asks whether land-use responses concentrate in forest-dependent species, follow an affiliation ordering, or differ in direction among guilds.

Sampling is not spatially even. Figueiredo et al. (2021) identify a centre-south concentration in the available occurrence knowledge for Atlantic Forest vertebrates, a pattern inherited by the five habitat-classifiable vertebrate groups in our compilation. Restricted-total and guild-resolved responses therefore describe the realised sampling domain rather than a spatially balanced census of the biome.

Predictor geometry and inference from observational gradients

Observational studies in landscape ecology, macroecology and related fields commonly relate an ecological response to several environmental predictors at once. Individual regression coefficients are then interpreted as relationships with particular environmental gradients, but that interpretation depends on a property of the sampled data as well as on the fitted model. Consider the simple additive model $Y = \beta_0 + \beta_1 F + \beta_2 A$. The coefficient β_1 represents the fitted change in Y with F while A is held fixed. For that coefficient to correspond closely to an observed ecological contrast, the data must contain enough observations in which F varies while A remains similar. A model can return an estimable coefficient even when that comparison is sparse, restricted to part of the sampled domain, or supported mainly by the assumed functional form.

Collinearity is the familiar version of this problem, but it describes only part of it. Pairwise correlations and variance inflation factors quantify aspects of linear dependence among predictors and identify situations in which coefficients become difficult to separate statistically. They do not describe the full arrangement of observations in multivariable predictor space. Two predictors can be only moderately correlated overall while sharing a dominant gradient, spanning very different ranges, occupying only part of their possible joint space, or remaining coupled nonlinearly. Conventional collinearity diagnostics can therefore appear acceptable even when few observations

provide the holding-fixed comparison that a partial coefficient needs for a straightforward ecological interpretation (Smith et al. 2009; Ray-Mukherjee et al. 2014).

I use predictor geometry for the realised joint structure of the sampled predictors. It includes the length of each gradient, the overlap among their ranges, the combinations that occur, how densely different regions are represented, and which combinations are rare or absent. This geometry is fixed by the environmental system and by how that system was sampled, before any regression coefficient is interpreted. It matters especially in observational ecology because predictors are often generated or constrained by the same processes. Habitat loss changes both habitat amount and spatial configuration, so those quantities need not vary independently across real landscapes (Koper et al. 2007; Ruffell et al. 2016). Climate, terrain, productivity and land use can likewise be geographically aligned. Entering all of them into one model does not create environmental contrasts that were absent from the sampled data.

Interaction models make gradient coverage still more important. An interaction asks whether the relationship between one predictor and the response changes along another. That comparison is informed most directly where the focal predictor spans a useful range at multiple positions along the moderating gradient. If high values of one predictor occur only where another is also high, some portions of the fitted response surface rest on substantial observed contrast while others depend much more heavily on interpolation, extrapolation or functional-form assumptions. Gradient length, overlap and sampling density determine which conditional relationships can be estimated directly from the observations. Changing spatial grain or predictor extent can alter that geometry even when the named variables remain unchanged.

These general problems recur across the empirical analyses in this dissertation. Across the Atlantic Forest chapters, pairwise correlations and variance inflation factors are combined with nonlinear coupling diagnostics, structure coefficients, mapped predictor distributions and comparisons across spatial scales. Chapter 5 additionally uses amount-residualised configuration metrics as a complementary representation of the amount–configuration relationship. Chapters 3 and 4 develop the inferential consequences more explicitly, by asking whether the realised predictor space contains the ecological contrasts assigned to fitted coefficients. These diagnostics establish what comparisons the sampled gradients actually support, and how directly a statistical estimate can sustain the ecological interpretation assigned to it. They complement model

comparison, effect estimates and uncertainty, and they do not convert an observational relationship into a causal effect (Lundberg et al. 2021).

Aims and chapter structure

The dissertation asks how climatic exposure, heterogeneity origin and productivity context modify land-use effects on Atlantic Forest species richness, and whether aggregate responses conceal graded or opposing effects across forest-specialist, forest-facultative and non-forest guilds (Figure 0.7). Diagnostic case studies and a response-surface perspective then examine the conditions required to interpret evidence about fragmentation per se, distinguishing habitat amount, forest configuration, response estimands and metric–mechanism correspondence. An Atlantic Forest application evaluates amount-conditioned configuration associations across habitat definitions, spatial scales, restricted total richness and the three forest-affiliation guilds in five habitat-classifiable vertebrate groups. The general discussion draws together the ecological and interpretive lessons of this progression and considers what they imply for conserving biodiversity in transformed tropical forest landscapes.

Chapter 1 examines climatic exposure as a moderator of land-use effects. Across full-universe total richness, restricted total richness, and the three forest-affiliation guilds, it asks whether agriculture and broader non-natural cover interact with daily and annual temperature range, maximum temperature of the warmest month, and dry-season precipitation, across response grains and predictor extents. Daily and annual temperature range represent thermal variability, because forest canopies buffer both peak temperatures and the amplitude of thermal fluctuations relative to open conditions (Ewers and Banks-Leite 2013; De Frenne et al. 2019). Maximum temperature of the warmest month provides a peak-heat axis, motivated by evidence that extreme exposed temperatures can approach upper thermal limits and that human-altered habitats are often warmer than natural habitats (Sunday et al. 2014; Williams and Newbold 2020). This contrasts peak exposure with variability across daily and annual cycles, so the temperature metrics enter as distinct rather than interchangeable axes. Comparing the aggregate and guild-resolved responses tests the prediction that climate-conditioned land-use effects strengthen with forest dependence. Water–energy availability and terrain heterogeneity enter additively to test whether climatic moderation persists after accounting for these broader

environmental gradients, and no land-use interaction with either variable is fitted. The chapter leaves the structural pathway underlying any detected moderation unresolved.

The same response architecture supports Chapter 2, which advances those two environmental dimensions into the focal question and asks whether productivity and distinct forms of environmental heterogeneity themselves condition land-use effects. It distinguishes persistent geophysical heterogeneity from land-use-responsive vegetation texture and further separates variability in greenness magnitude from configurational greenness diversity. This design draws on heterogeneity-diversity theory and the area-heterogeneity tradeoff, treating heterogeneity origin as the key distinction in a transformed forest biome (Tews et al. 2004; Kadmon and Allouche 2007; Allouche et al. 2012; Stein et al. 2014; Seiferling et al. 2014; Heidrich et al. 2020). It also evaluates productivity and water–energy balance as candidate resource-related moderators, without treating either as evidence of abundance-mediated buffering (Hawkins et al. 2003; Currie et al. 2004; Storch et al. 2018). Across the same full- and restricted-universe response architecture, the chapter asks whether heterogeneity acts as a single, uniformly beneficial axis or whether its ecological meaning depends on origin, metric, and the assemblage component used as a response. Structure-coefficient and marginal-association diagnostics are used to distinguish ecological direction from distortions in partial slopes under predictor coupling, so that the sign of a heterogeneity-diversity relationship is read from coefficient geometry and not from a partial slope alone. Because those vegetation-texture metrics do not separate spatial configuration from habitat amount, the chapter treats them as heterogeneity signals rather than estimates of fragmentation per se.

The diagnostic analysis in Chapter 3 evaluates whether configuration coefficients from amount-controlled models correspond to the fragmentation-per-se contrasts they are used to represent. It compares a global richness analysis with an Atlantic Forest tree beta-diversity analysis to examine how null and non-zero configuration coefficients should be interpreted (Gonçalves-Souza et al. 2025; Fahrig et al. 2026; Kramer et al. 2026). The four conditions are predictor-space adequacy, response-estimand alignment, metric–mechanism correspondence, and domain stability. The chapter uses predictor-coupling, coefficient-architecture, response, metric, scale, and subset diagnostics (Lundberg et al. 2021; Smith et al. 2009; Koper et al. 2007; Ruffell et al. 2016; Ray-Mukherjee et al. 2014). Statistical detectability is treated as an intervening

diagnostic tier rather than a fifth condition. The chapter asks what evidence is required before either a null or a non-zero configuration coefficient can be interpreted as evidence about fragmentation per se. Because local richness can mask compositional change, Chapter 3 also examines beta-diversity responses to test whether configuration coefficients persist once the response estimand is aligned with the mechanism being claimed (Chase et al. 2020).

A response-surface framework for fragmentation inference follows in Chapter 4. When configuration effects vary with habitat amount, an additive coefficient cannot describe that variation, so the inferential target becomes the configuration effect as a function of the habitat-amount gradient. The chapter treats the fragmentation debate as a question about how the configuration effect varies with habitat amount rather than as a single additive fragmentation coefficient. Habitat amount invariance, threshold dynamics, colonisation–competition coexistence, and extinction filtering differ in whether the configuration effect is absent, intensifies under scarcity, reverses sign along the habitat-amount gradient, or fades after sensitive species have been filtered from disturbed assemblages (Andrén 1994; Fahrig 2003, 2013; Betts et al. 2019; Rybicki et al. 2020; Saura 2021; Zhang et al. 2024). The chapter argues that fragmentation × habitat amount is the minimum linear structure needed to represent this conditionality, while beta-diversity decomposition, gradient position, and metric–mechanism correspondence provide additional discriminating evidence (Baselga 2010). The interaction is the model structure the competing theories require rather than a guarantee that the corresponding configuration effect is identified.

The Atlantic Forest returns in Chapter 5 through an amount-conditioned analysis of forest configuration across five habitat-classifiable vertebrate groups, with data distinct from the external datasets used in the diagnostic case studies. The chapter asks whether configuration associations remain after modelling their dependence on forest amount and whether those associations change along the habitat-amount gradient. The analysis evaluates multiple forest products, canopy thresholds, normalised and enumeration-based fragmentation metrics, raw and residualised formulations, response grains, and predictor extents, across restricted total richness and the three mutually exclusive forest-specialist, forest-facultative, and non-forest guilds. Habitat definition is treated as an analytical variable rather than a fixed input, because mapped vegetation classes approximate species-specific habitat only imperfectly (Almeida-Gomes et al.

2016). Restricted total richness provides the broad assemblage response, while the guild partition tests whether configuration effects vary with sensitivity to forest subdivision, edge exposure, matrix contrast, and loss of interior habitat conditions (Fahrig 2003, 2013, 2017; Ewers and Didham 2006; Villard and Metzger 2014; Fletcher et al. 2018, 2023; Fahrig et al. 2019). The forest cover by fragmentation interaction is evaluated across products, thresholds, metrics, formulations, scales, and response strata, so no single coefficient carries the interpretation. The fitted surface is interpreted only after evaluating predictor geometry, metric correspondence, formulation, response strata, gradient coverage and uncertainty, rather than as self-validating evidence of one ecological mechanism.

The final Atlantic Forest application places ecological moderation and fragmentation inference within the same amount-conditioned, guild-resolved analysis. Broad landscape indicators can guide conservation, but their interpretation depends on how closely they represent the response and mechanism under study. Poor alignment calls for more specific response variables and landscape metrics.

References

Allouche O, Kalyuzhny M, Moreno-Rueda G, Pizarro M, Kadmon R (2012) Area-heterogeneity tradeoff and the diversity of ecological communities. Proceedings of the National Academy of Sciences 109:17495-17500.

Almeida-Gomes M, Prevedello JA, Crouzeilles R (2016) The use of native vegetation as a proxy for habitat may overestimate habitat availability in fragmented landscapes. Landscape Ecology 31:711-719.

Alroy J (2017) Effects of habitat disturbance on tropical forest biodiversity. Proceedings of the National Academy of Sciences 114:6056-6061.

Andrén H (1994) Effects of habitat fragmentation on birds and mammals in landscapes with different proportions of suitable habitat: a review. Oikos 71:355-366.

Banks-Leite C, Betts MG, Ewers RM, Orme CDL, Pigot AL (2022) The macroecology of landscape ecology. Trends in Ecology and Evolution 37:480-487.

Baselga A (2010) Partitioning the turnover and nestedness components of beta diversity. Global Ecology and Biogeography 19:134-143.

Bélisle M (2005) Measuring landscape connectivity: the challenge of behavioral landscape ecology. Ecology 86:1988-1995.

Ben-Hur E, Kadmon R (2020) An experimental test of the area-heterogeneity tradeoff. Proceedings of the National Academy of Sciences 117:4815-4822.

Betts MG, Wolf C, Pfeifer M, Banks-Leite C, Arroyo-Rodríguez V, Ribeiro DB, et al. (2019) Extinction filters mediate the global effects of habitat fragmentation on animals. Science 366:1236-1239.

Brook BW, Sodhi NS, Bradshaw CJA (2008) Synergies among extinction drivers under global change. Trends in Ecology and Evolution 23:453-460.

Chase JM, Jeliazkov A, Ladouceur E, Viana DS (2020) Biodiversity conservation through the lens of metacommunity ecology. Annals of the New York Academy of Sciences 1469:86-104.

Clavel J, Julliard R, Devictor V (2011) Worldwide decline of specialist species: toward a global functional homogenization? Frontiers in Ecology and the Environment 9:222-228.

Currie DJ, Mittelbach GG, Cornell HV, et al. (2004) Predictions and tests of climate-based hypotheses of broad-scale variation in taxonomic richness. Ecology Letters 7:1121-1134.

De Frenne P, et al. (2019) Global buffering of temperatures under forest canopies. Nature Ecology and Evolution 3:744-749.

Devictor V, Julliard R, Jiguet F (2008) Distribution of specialist and generalist species along spatial gradients of habitat disturbance and fragmentation. Oikos 117:507-514.

Devictor V, Clavel J, Julliard R, Lavergne S, Mouillot D, Thuiller W, Venail P, Villéger S, Mouquet N (2010) Defining and measuring ecological specialization. Journal of Applied Ecology 47:15-25.

Didham RK, Kapos V, Ewers RM (2012) Rethinking the conceptual foundations of habitat fragmentation research. Oikos 121:161-170.

Donald PF, Arendarczyk B, Spooner F, Buchanan GM (2019) Loss of forest intactness elevates global extinction risk in birds. Animal Conservation 22:341-347.

Edwards DP, Socolar JB, Mills SC, Burivalova Z, Koh LP, Wilcove DS (2019) Conservation of Tropical Forests in the Anthropocene. Current Biology 29:R1008-R1020.

Ewers RM, Didham RK (2006) Confounding factors in the detection of species responses to habitat fragmentation. Biological Reviews 81:117-142.

Ewers RM, Banks-Leite C (2013) Fragmentation impairs the microclimate buffering effect of tropical forests. PLoS One 8:e58093.

Fahrig L (2003) Effects of habitat fragmentation on biodiversity. Annual Review of Ecology, Evolution, and Systematics 34:487-515.

Fahrig L (2013) Rethinking patch size and isolation effects: the habitat amount hypothesis. Journal of Biogeography 40:1649-1663.

Fahrig L (2017) Ecological responses to habitat fragmentation per se. Annual Review of Ecology, Evolution, and Systematics 48:1-23.

Fahrig L (2021) What the habitat amount hypothesis does and does not predict: a reply to Saura. Journal of Biogeography 48:1530-1535.

Fahrig L, Arroyo-Rodríguez V, Bennett JR, et al. (2019) Is habitat fragmentation bad for biodiversity? Biological Conservation 230:179-186.

Fahrig L, Galán-Acedo C, Edwards BPM, et al. (2026) Why controlling for habitat amount is critical for resolving the fragmentation debate. Conservation Biology e70245.

Figueiredo MSL, Weber MM, Brasileiro CA, et al. (2021) Tetrapod diversity in the Atlantic Forest: maps and gaps. In: Marques MCM, Grelle CEV (eds) The Atlantic Forest: history, biodiversity, threats and opportunities of the mega-diverse forest. Springer, Cham, pp 185-203.

Fischer J, Lindenmayer DB (2007) Landscape modification and habitat fragmentation: a synthesis. Global Ecology and Biogeography 16:265-280.

Fletcher RJ Jr, Didham RK, Banks-Leite C, et al. (2018) Is habitat fragmentation good for biodiversity? Biological Conservation 226:9-15.

Fletcher RJ Jr, Betts MG, Damschen EI, Hefley TJ, Hightower J, Smith TA, Fortin MJ, Haddad NM (2023) Addressing the problem of scale that emerges with habitat fragmentation. Global Ecology and Biogeography 32:828-841.

Foley JA, DeFries R, Asner GP, et al. (2005) Global consequences of land use. Science 309:570-574.

Gonçalves-Souza T, Chase JM, Haddad NM, et al. (2025) Species turnover does not rescue biodiversity in fragmented landscapes. Nature 640:702-706.

Haddad NM, Brudvig LA, Clobert J, et al. (2015) Habitat fragmentation and its lasting impact on Earth's ecosystems. Science Advances 1:e1500052.

Hawkins BA, Field R, Cornell HV, et al. (2003) Energy, water, and broad-scale geographic patterns of species richness. Ecology 84:3105-3117.

Heidrich L, Bae S, Friess N, et al. (2020) Heterogeneity-diversity relationships differ between and within trophic levels in temperate forests. Nature Ecology and Evolution 4:1204-1212.

Hjort J, Gordon JE, Gray M, Hunter ML Jr (2015) Why geodiversity matters in valuing nature's stage. Conservation Biology 29:630-639.

Hutchinson GE (1957) Concluding remarks. Cold Spring Harbor Symposia on Quantitative Biology 22:415-427.

Jackson HB, Fahrig L (2015) Are ecologists conducting research at the optimal scale? Global Ecology and Biogeography 24:52-63.

Kadmon R, Allouche O (2007) Integrating the effects of area, isolation, and habitat heterogeneity on species diversity: a unification of island biogeography and niche theory. The American Naturalist 170:443-454.

Khaliq I, et al. (2014) Interspecific variation in responses to thermal environments across vertebrate ectotherms and endotherms. Oecologia 174:1277-1287.

Koper N, Schmiegelow FKA, Merrill EH (2007) Residuals cannot distinguish between ecological effects of habitat amount and fragmentation: implications for the debate. Landscape Ecology 22:811-820.

Kramer JMF, Zwiener VP, Camana M, de Lima RAF, Müller SC (2026) Effects of deforestation and fragmentation on tree beta diversity and endemism across landscapes in the Atlantic Forest biodiversity hotspot. Global Change Biology 32:e70855.

Laanisto L, Tamme R, Hiiesalu I, et al. (2013) Microfragmentation concept explains non-positive environmental heterogeneity-diversity relationships. Oecologia 171:217-226.

Laurance WF, Useche DC, Rendeiro J, et al. (2012) Averting biodiversity collapse in tropical forest protected areas. Nature 489:290-294.

Levin SA (1992) The problem of pattern and scale in ecology. Ecology 73:1943-1967.

Lundberg I, Johnson R, Stewart BM (2021) What is your estimand? Defining the target quantity connects statistical evidence to theory. American Sociological Review 86:532-565.

MacArthur RH (1972) Geographical Ecology. Harper and Row, New York.

Mantyka-Pringle CS, Martin TG, Rhodes JR (2012) Interactions between climate and habitat loss effects on biodiversity: a systematic review and meta-analysis. Global Change Biology 18:1239-1252.

Marques MCM, Trindade W, Bohn A, Grelle CEV (2021) The Atlantic Forest: an introduction to the megadiverse forest of South America. In: Marques MCM, Grelle CEV (eds) The Atlantic Forest: history, biodiversity, threats and opportunities of the mega-diverse forest. Springer, Cham, pp 3-23.

Martínez-Lanfranco JA (2026a) Habitat amount control is necessary but not sufficient to resolve the fragmentation debate. Conservation Biology, accepted. EcoEvoRxiv. DOI 10.32942/X2KT05.

Matthews TJ, Cottee-Jones HE, Whittaker RJ (2014) Habitat fragmentation and the species-area relationship: a focus on total species richness obscures the impact of habitat loss on habitat specialists. Diversity and Distributions 20:1136-1146.

McKinney ML, Lockwood JL (1999) Biotic homogenization: a few winners replacing many losers in the next mass extinction. Trends in Ecology and Evolution 14:450-453.

Meyer WB, Turner BL (1992) Human population growth and global land-use/cover change. Annual Review of Ecology and Systematics 23:39-61.

Miguet P, Jackson HB, Jackson ND, Martin AE, Fahrig L (2016) What determines the spatial extent of landscape effects on species? Landscape Ecology 31: 1177-1194.

Miller-Rushing AJ, Primack RB, Sekercioglu CH, et al. (2019) How does habitat fragmentation affect biodiversity? A controversial question at the core of conservation biology. Biological Conservation 232:271-273.

Neel MC, McGarigal K, Cushman SA (2004) Behavior of class-level landscape metrics across gradients of class aggregation and area. Landscape Ecology 19:435-455.

Newbold T, Hudson LN, Hill SLL, et al. (2015) Global effects of land use on local terrestrial biodiversity. Nature 520:45-50.

Newbold T, Hudson LN, Contu S, et al. (2018) Widespread winners and narrow-ranged losers: land use homogenizes biodiversity in local assemblages worldwide. PLoS Biology 16:e2006841.

Newbold T, Oppenheimer P, Etard A, Williams JJ (2020) Tropical and Mediterranean biodiversity is disproportionately sensitive to land-use and climate change. Nature Ecology and Evolution 4:1630-1638.

Oliver TH, Morecroft MD (2014) Interactions between climate change and land use change on biodiversity: attribution problems, risks, and opportunities. WIREs Climate Change 5:317-335.

Pärtel M, Bennett JA, Zobel M (2016) Macroecology of biodiversity: disentangling local and regional effects. New Phytologist 211:404-410.

Ray-Mukherjee J, Nimon K, Mukherjee S, Morris D, Slotow R, Hamer M (2014) Using commonality analysis in multiple regressions: a tool to decompose regression effects in the face of multicollinearity. Methods in Ecology and Evolution 5:320-328.

Ribeiro MC, Metzger JP, Martensen AC, Ponzoni FJ, Hirota MM (2009) The Brazilian Atlantic Forest: how much is left, and how is the remaining forest distributed? Implications for conservation. Biological Conservation 142:1141-1153.

Ricklefs RE (1987) Community diversity: relative roles of local and regional processes. Science 235:167-171.

Ricklefs RE (2008) Disintegration of the ecological community. The American Naturalist 172:741-750.

Riva F, Koper N, Fahrig L (2024) Overcoming confusion and stigma in habitat fragmentation research. Biological Reviews 99:1411-1424.

Ruffell J, Banks-Leite C, Didham RK (2016) Accounting for the causal basis of collinearity when measuring the effects of habitat loss versus habitat fragmentation. Oikos 125:117-125.

Rybicki J, Abrego N, Ovaskainen O (2020) Habitat fragmentation and species diversity in competitive communities. Ecology Letters 23:506-517.

Saura S (2021) The habitat amount hypothesis implies negative effects of habitat fragmentation on species richness. Journal of Biogeography 48:11-22.

Scheffers BR, et al. (2016) The broad footprint of climate change from genes to biomes to people. Science 354:aaf7671.

Seiferling I, Proulx R, Wirth C (2014) Disentangling the environmental-heterogeneity-species-diversity relationship along a gradient of human footprint. Ecology 95:2084-2095.

Sloan S, Jenkins CN, Joppa LN, et al. (2014) Remaining natural land cover in the global biodiversity hotspots. Biological Conservation 177:12-24.

Smith AC, Koper N, Francis CM, Fahrig L (2009) Confronting collinearity: comparing methods for disentangling the effects of habitat loss and fragmentation. Landscape Ecology 24:1271-1285.

Solórzano A, Brasil LSCA, Oliveira RR (2021) The Atlantic Forest ecological history: from pre-colonial times to the Anthropocene. In: Marques MCM, Grelle CEV (eds) The Atlantic Forest: history, biodiversity, threats and opportunities of the mega-diverse forest. Springer, Cham, pp 25-44.

Stein A, Gerstner K, Kreft H (2014) Environmental heterogeneity as a universal driver of species richness across taxa, biomes and spatial scales. Ecology Letters 17:866-880.

Stein A, Kreft H (2015) Terminology and quantification of environmental heterogeneity in species-richness research. Biological Reviews 90:815-836.

Storch D, Bohdalková E, Okie J (2018) The more-individuals hypothesis revisited: the role of community abundance in species richness regulation and the productivity-diversity relationship. Ecology Letters 21:920-937.

Sunday JM, et al. (2014) Thermal-safety margins and the necessity of thermoregulatory behavior across latitude and elevation. Proceedings of the National Academy of Sciences USA 111:5610-5615.

Tabarelli M, Aguiar AV, Ribeiro MC, Metzger JP, Peres CA (2010) Prospects for biodiversity conservation in the Atlantic Forest: lessons from aging human-modified landscapes. Biological Conservation 143:2328-2340.

Taylor PD, Fahrig L, Henein K, Merriam G (1993) Connectivity is a vital element of landscape structure. Oikos 68:571-573.

Teng SN, Svenning J-C, Santana J, Reino L, Abades S, Xu C (2020) Linking landscape ecology and macroecology by scaling biodiversity in space and time. Current Landscape Ecology Reports 5:25-34.

Tews J, Brose U, Grimm V, et al. (2004) Animal species diversity driven by habitat heterogeneity/diversity: the importance of keystone structures. Journal of Biogeography 31:79-92.

Tscharntke T, Tylianakis JM, Rand TA, et al. (2012) Landscape moderation of biodiversity patterns and processes - eight hypotheses. Biological Reviews 87:661-685.

Vancine MH, Muylaert RL, Niebuhr BB, et al. (2024) The Atlantic Forest of South America: spatiotemporal dynamics of the vegetation and implications for conservation. Biological Conservation 291:110499.

Villard M-A, Metzger JP (2014) Beyond the fragmentation debate: a conceptual model to predict when habitat configuration really matters. Journal of Applied Ecology 51:309-318.

Wang X, Blanchet FG, Koper N (2014) Measuring habitat fragmentation: an evaluation of landscape pattern metrics. Methods in Ecology and Evolution 5:634-646.

Wiens JA (1989) Spatial scaling in ecology. Functional Ecology 3:385-397.

Whittaker RH (1972) Evolution and measurement of species diversity. Taxon 21:213–251.

Williams JJ, Newbold T (2020) Local climatic changes affect biodiversity responses to land use: a review. Diversity and Distributions 26:76-92.

Zhang H, Chase JM, Liao J (2024) Habitat amount modulates biodiversity responses to fragmentation. Nature Ecology and Evolution 8:1437-1447.

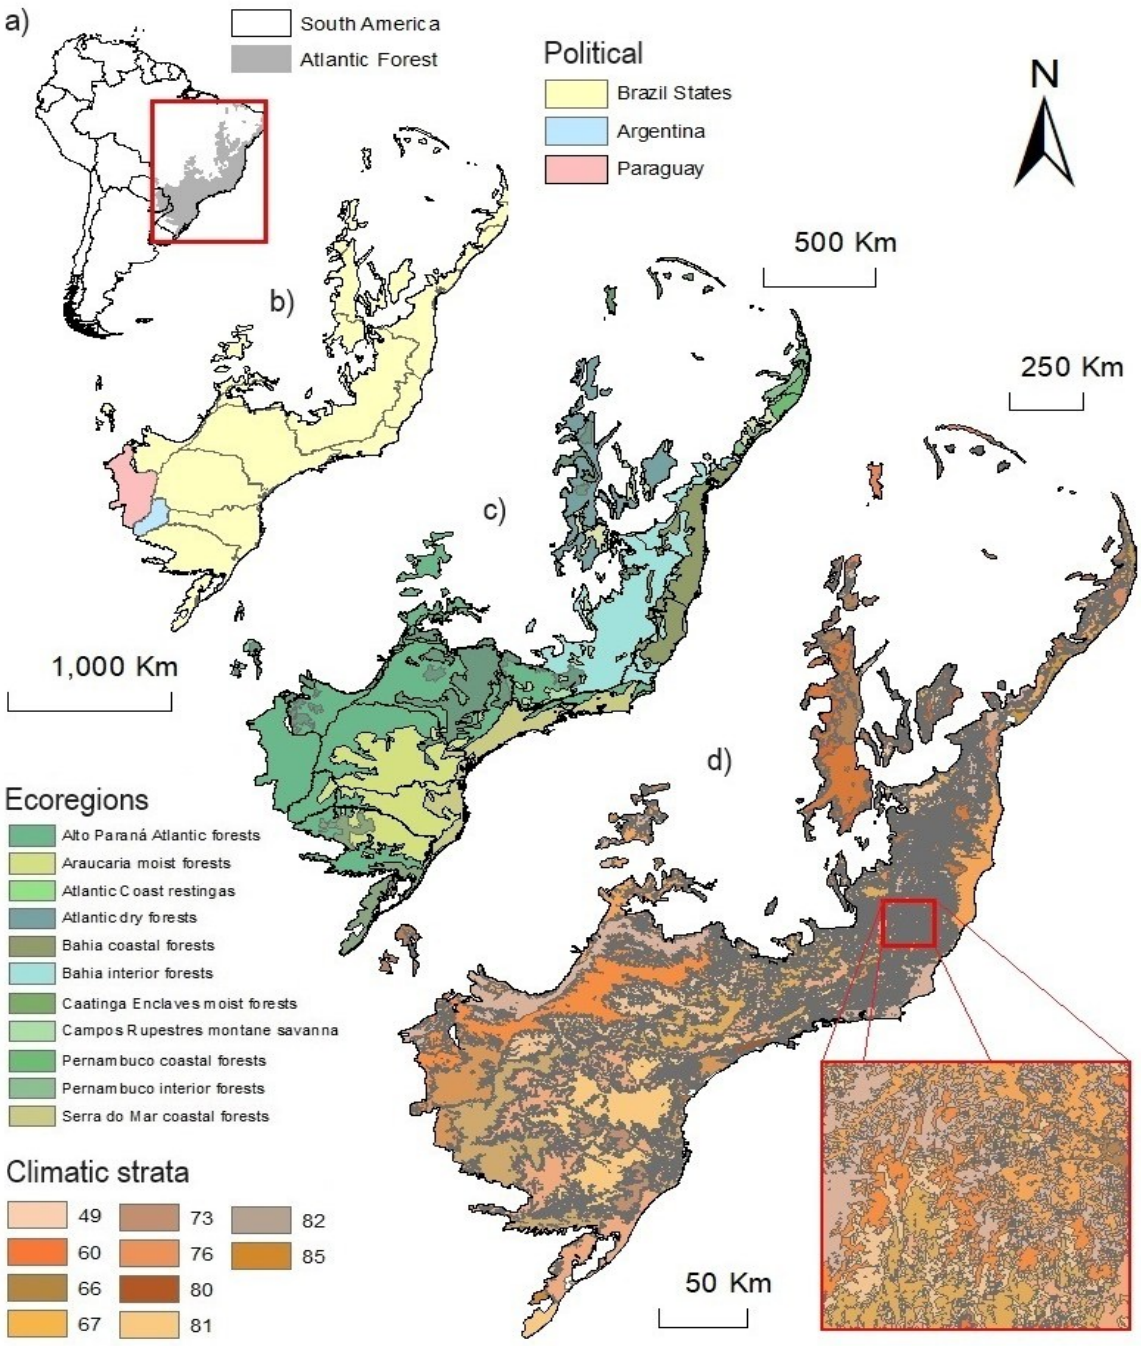


Figure 0.1. The Atlantic Forest study region in South America. (a) Location within South America, with the biome shaded. (b) Political divisions spanning Brazilian states, Argentina and Paraguay. (c) Terrestrial ecoregions within the biome. (d) Climatic strata, with an inset showing the fine-grained interdigitation of strata at landscape scale. The biome spans a wide latitudinal range, several ecoregions and marked climatic variation, which is the environmental context across which the dissertation evaluates land-use effects.

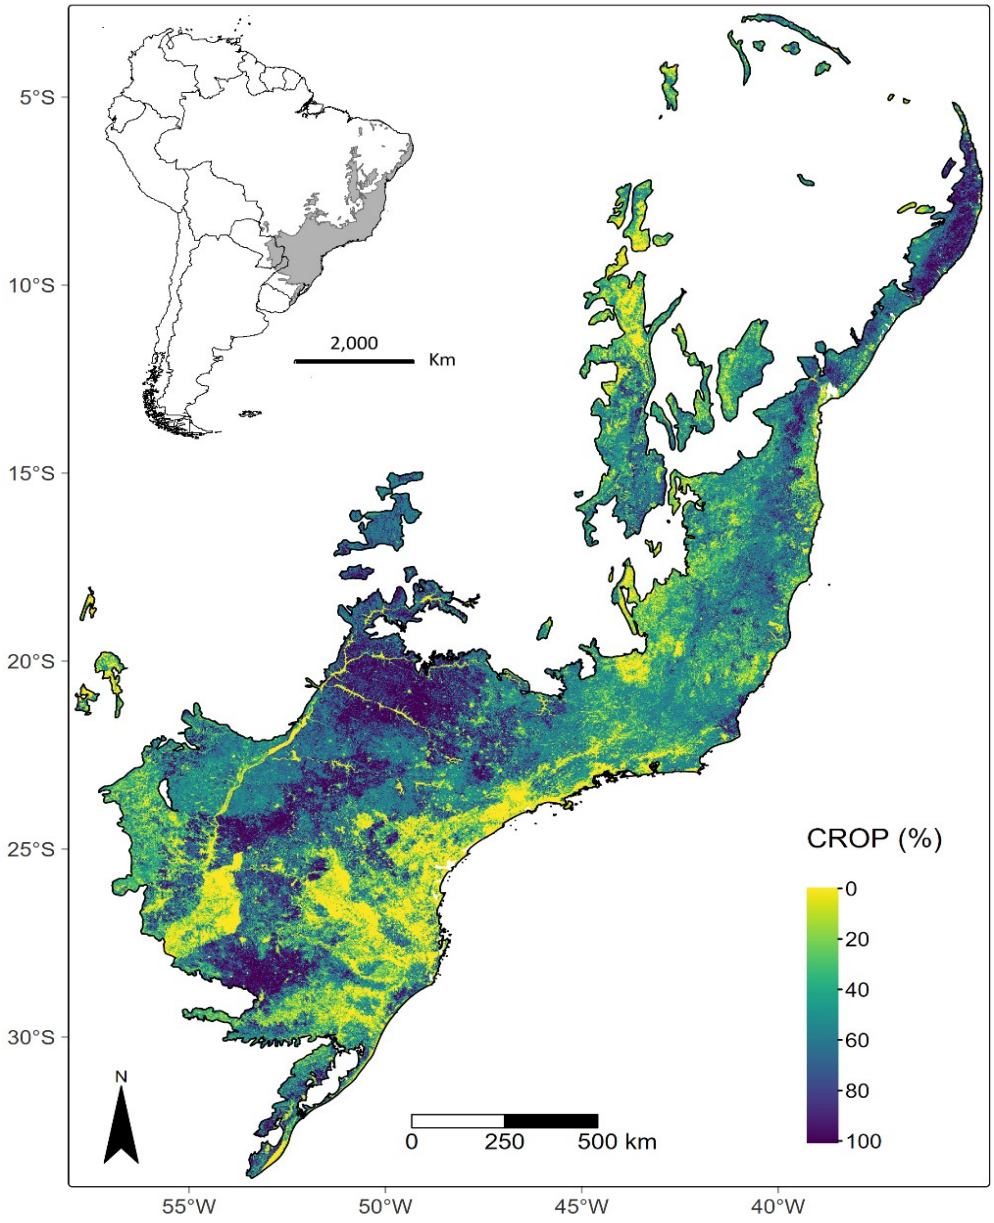


Figure 0.2. Cropland cover across the Atlantic Forest. Percentage of cropland per pixel across the biome, shown here to illustrate the extent and spatial structure of agricultural transformation rather than as one of the predictors used in the analyses. Cropland is unevenly distributed, forming extensive contiguous tracts in the interior south and west while coastal and montane areas retain lower cover, so land-use transformation varies both in amount and in spatial configuration across the biome. Chapter-specific appendices document the provenance, resolution and collinearity of all environmental predictors.

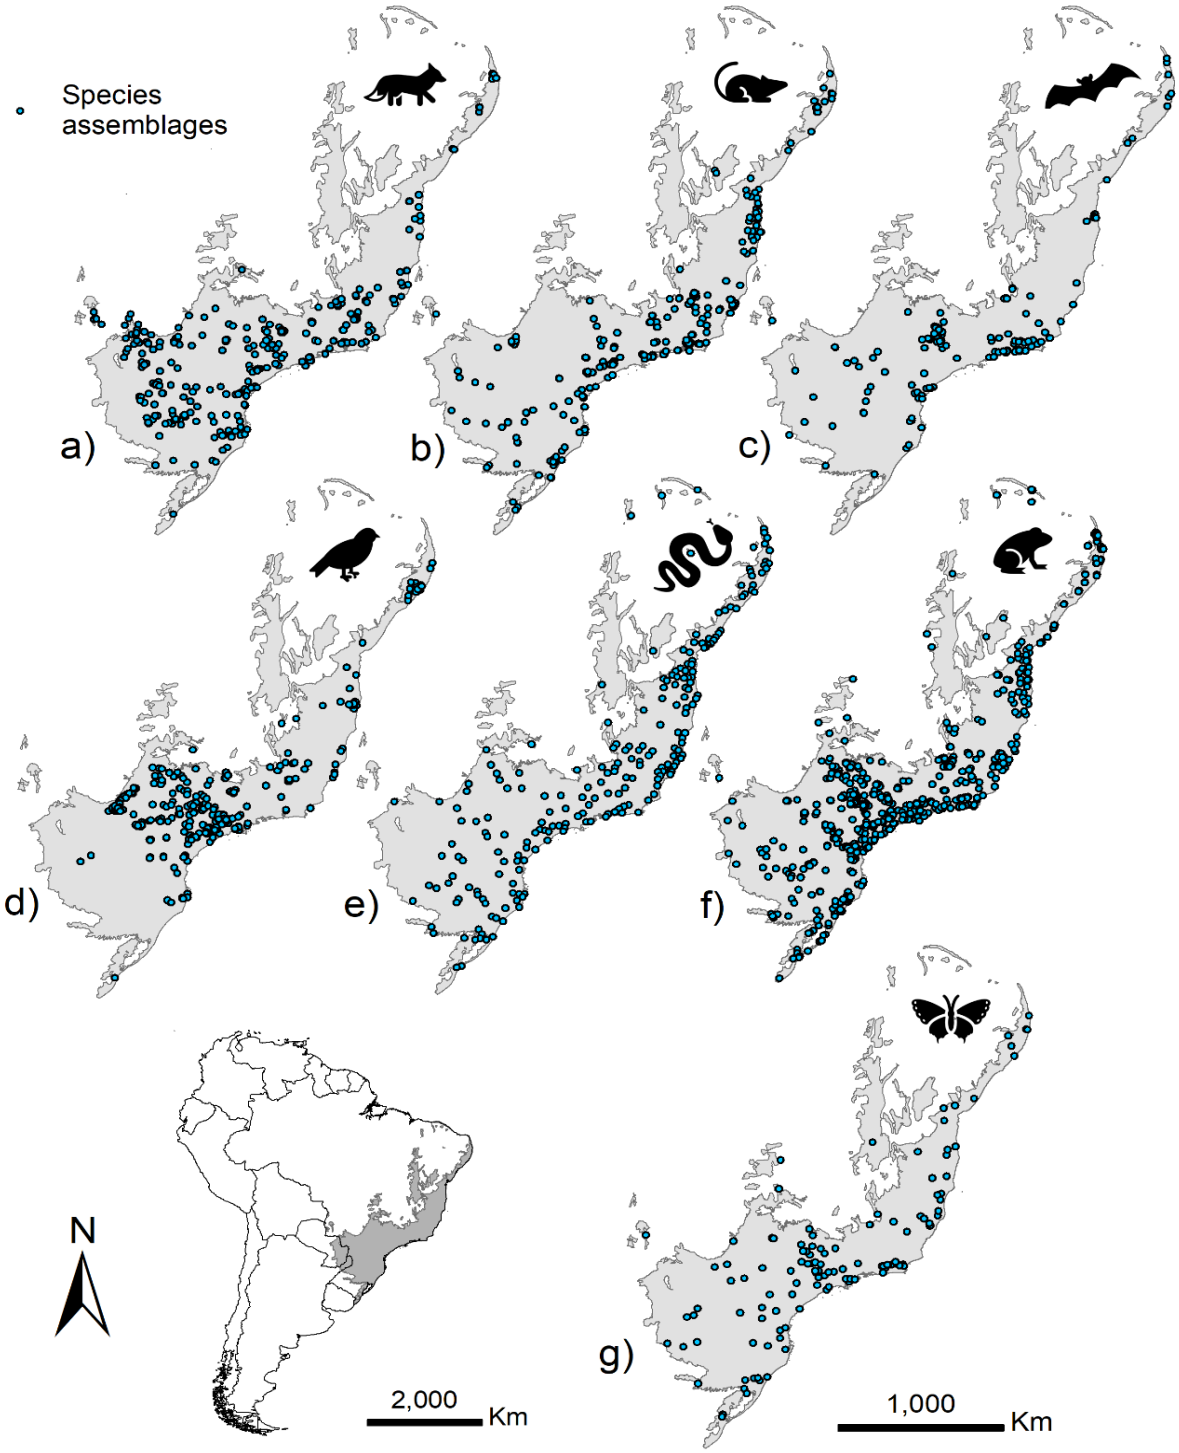


Figure 0.3. Geographic distribution of the curated species assemblages across the Atlantic Forest. Points show the locations of raw, curated assemblages (n = 2,986) across the seven dataset-defined animal groups or taxa between 1981 and 2018: mid-large mammals (a), small mammals (b), bats (c), songbirds (d), snakes (e), anurans (f), and butterflies (g).

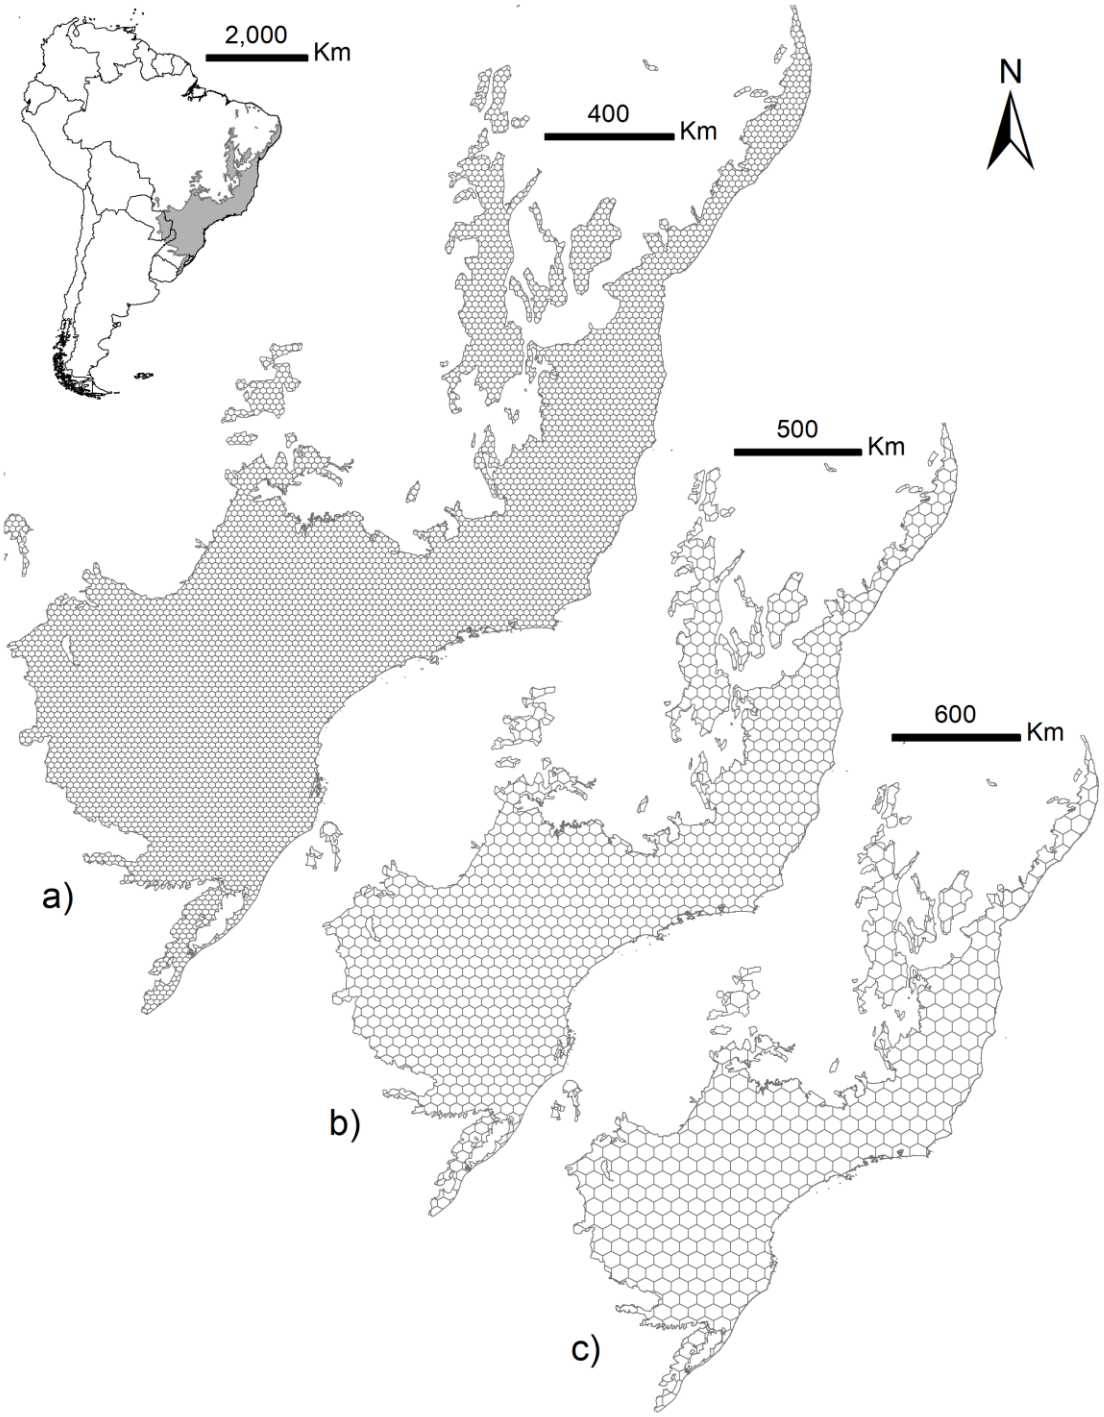


Figure 0.4. Hexagonal response units at the three response grains. Hexagonal grids equivalent in area to circles of 10 km (a), 25 km (b) and 40 km (c) radius across the Atlantic Forest. The hexagons serve as an aggregation template determining which assemblages combine into one record; they do not define where environmental predictors are measured. Figure 0.S1.3 shows which units are occupied after aggregation, by animal group and sampling method.

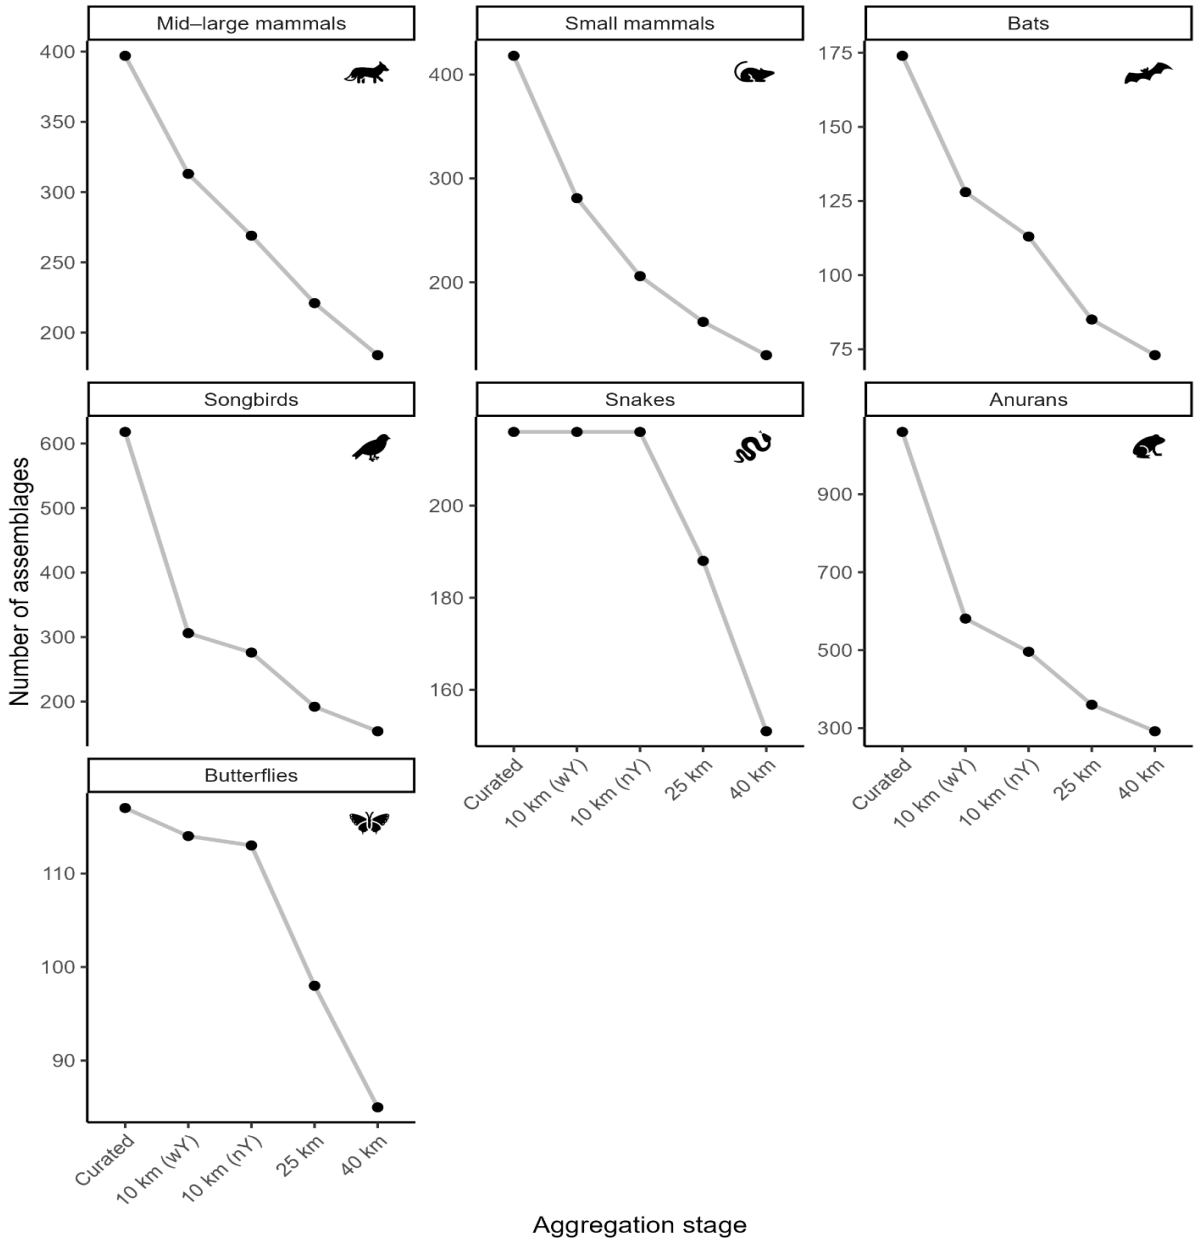


Figure 0.5. Sample-size attrition by animal group across aggregation stages. Number of assemblages retained at each stage, from the raw, curated data (n = 2,986) through the three response grains. Stages labelled 10 km (wY) and 10 km (nY) distinguish the year-retained and year-aggregated variants at the 10 km grain; the analyses reported in this dissertation use the year-aggregated form throughout. Attrition is steepest for groups sampled densely relative to the grain, because more source assemblages consolidate into each modelling record. Snake counts are unchanged from the curated data through the 10 km grain because snake assemblages were already more than 10 km apart.

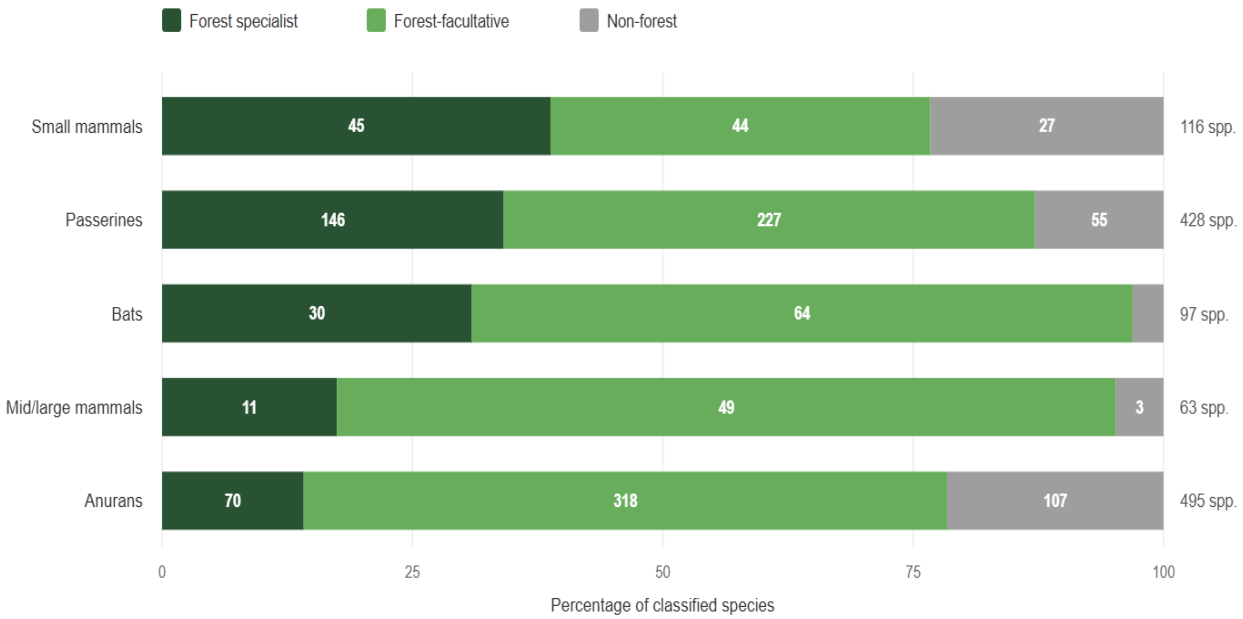


Figure 0.6. Forest-affiliation composition of the habitat-classifiable taxon pool. Horizontal 100% stacked bars show, for each of the five vertebrate groups with species-level IUCN habitat coding — small mammals, passerines, bats, mid-large mammals and anurans — the proportion of classified taxa assigned as forest specialist, forest-facultative, or non-forest, with counts labelled within each bar and the number of classified taxa per group at right. Groups are ordered by decreasing forest-specialist prevalence. Mid-large mammals contribute 63 classified entries, 56 resolved species and seven genus-level records, against a curated gamma of 70 resolved species. These proportions describe the classified pool from which the guild responses are constructed and are not model-weighted assemblage proportions.

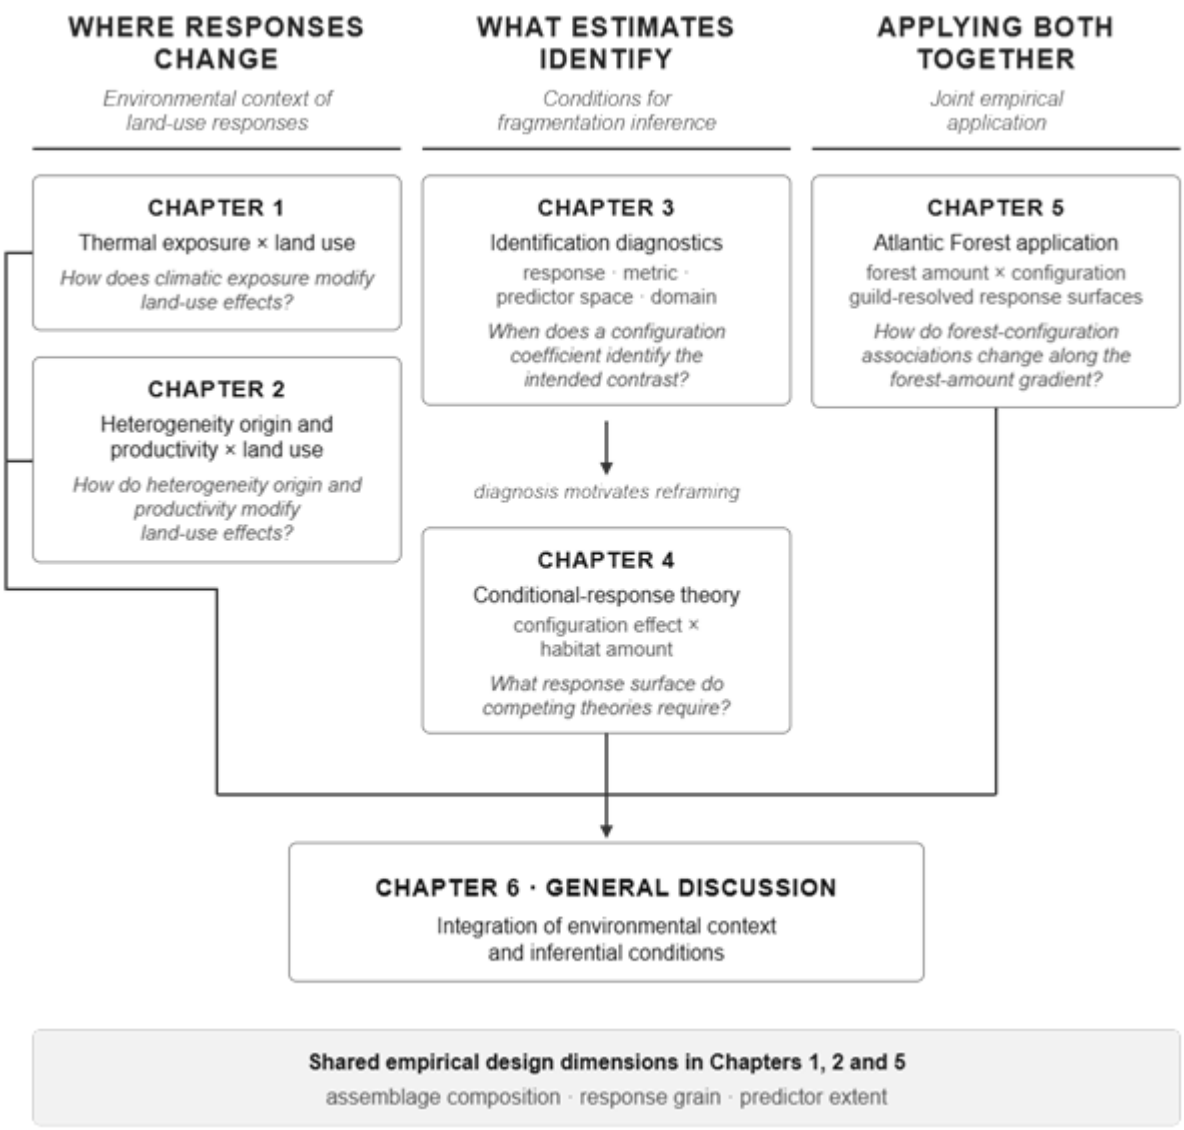


Figure 0.7. Conceptual structure of the dissertation. Chapters 1 and 2 examine how environmental context modifies land-use responses in Atlantic Forest assemblages and are parallel analyses rather than a sequence. Chapter 3 evaluates the conditions under which an amount-controlled configuration coefficient identifies the ecological contrast assigned to it, and Chapter 4 develops the conditional response-surface framework that competing ecological theories imply. Chapter 5 returns to Atlantic Forest vertebrate assemblages and asks how forest-configuration associations change along the forest-amount gradient. Assemblage composition, response grain and predictor extent are shared design dimensions of the empirical analyses in Chapters 1, 2 and 5. Chapter 6 integrates the resulting strands in the general discussion.

CHAPTER 1

Amplitude, not extremes: thermal variability amplifies land-use effects on species diversity in the Atlantic Forest

Introduction

Land-use intensification is a major driver of biodiversity loss. It reduces habitat area, simplifies resources, and filters assemblages toward species that persist in human-modified landscapes (Newbold et al. 2015). In forests, however, habitat and resource loss are not the only pathways linking land-use transformation to biodiversity change. By reducing canopy cover and structural complexity, transformation also alters the local thermal and hydric conditions organisms experience, so the biodiversity cost of land use may depend on the climatic context in which it occurs (Williams and Newbold 2020; Williams et al. 2022). Tropical forests make this coupling especially consequential because they combine high biodiversity with extensive deforestation, which has converted continuous forest into cropland, pasture, and infrastructure across millions of hectares (Laurance et al. 2012; Edwards et al. 2019).

The climatic consequence of this transformation begins at the canopy. Forests buffer local thermal and hydric conditions through canopy cover and structural complexity, moderating both daytime heating and nighttime cooling and stabilising under-canopy humidity (De Frenne et al. 2019; Ewers and Banks-Leite 2013; Jucker et al. 2020). Pantropical near-ground sensor networks confirm that diurnal temperature range inside tropical forest understories is on average 1.7°C lower than open-air conditions, with spatial variability regulated by vegetation structure and topography rather than captured by existing macroclimate grids (Ismaeel et al. 2024).

Global forest-change analyses further show that tropical deforestation increases local land-surface temperature through changes in evapotranspiration and albedo, whereas forest gain can produce local cooling, indicating that land-cover change directly modifies the thermal template in which biodiversity responses occur (Prevedello et al. 2019). Forest conversion also creates edges that alter the thermal and hydric conditions experienced by forest-core vertebrates, with abundances of 85% of vertebrate species globally affected by forest edges, and forest-core

species being 3.7 times more likely to be IUCN-listed as threatened, confirming that biophysical restructuring from land-use transformation disproportionately amplifies exposure of forest-affiliated assemblages (Pfeifer et al. 2017). Conversion to agriculture can produce strong local daytime warming at organism-relevant scales, reaching up to 13.6°C above primary forest in pantropical measurements, whereas forest degradation produces comparatively smaller, non-significant thermal changes, indicating that conversion type matters for the magnitude of local climatic disruption (Senior et al. 2017). Removing canopy increases organismal exposure to atmospheric variability and dry-season stress, particularly in tropical systems where many species exhibit narrow thermal niches and limited physiological tolerance to climatic extremes (Sunday et al. 2014; Khaliq et al. 2014; Senior et al. 2017).

Such mechanisms predict that the biodiversity costs of land-use transformation depend on the climatic context in which transformation occurs. Where atmospheric thermal variability is already high, forest canopies buffer a larger absolute amplitude of temperature fluctuation, so the same degree of canopy removal produces a greater increment of organismal thermal exposure; the biodiversity cost of land-use transformation should therefore scale with the atmospheric variability from which canopy buffering is withdrawn (De Frenne et al. 2019; Kemppinen et al. 2024). Global assemblage analyses are consistent with this filtering pathway. Vertebrate communities in human-dominated land uses contain higher proportions of individuals affiliated with warmer, drier, and more climatically extreme niches than communities in primary vegetation, with stronger compositional shifts in tropical than in temperate assemblages, and population abundances in human-altered habitats are most reduced, relative to natural habitats, where thermal conditions approach species' tolerance limits (Williams et al. 2020; Williams and Newbold 2021).

In the Atlantic Forest, large-scale biodiversity analyses have often emphasised climatic, energetic, or historical gradients (Carnaval and Moritz 2008; Marques and Grelle 2021), whereas landscape studies have more often focused on habitat conversion and fragmentation (Banks-Leite et al. 2014; Ribeiro et al. 2009). Few analyses ask whether contemporary climatic context modifies land-use effects across both response grains and predictor extents. Retaining approximately 12 to 16 % of its original forest cover, depending on the forest definition applied (Ribeiro et al. 2009) and strong gradients in climate, elevation, and land-use intensity, the

Atlantic Forest provides a high-resolution gradient across which the climatic conditionality of land-use effects can be tested.

Land-use-driven local climatic change has been identified as a key mechanism through which habitat conversion may alter biodiversity responses to land use (Williams and Newbold 2020). Existing studies have largely quantified either the physical climatic consequences of forest loss or the climatic filtering of assemblages across land-use types. Global analyses show that tropical biodiversity is disproportionately sensitive to both drivers, partly because tropical species often occur near their upper thermal tolerance limits and modified land uses tend to be warmer, drier, and less thermally buffered than natural habitats (Newbold et al. 2020). Yet few studies provide scale-explicit assemblage-level tests of whether land-use effects on richness are themselves amplified along climatic axes associated with reduced thermal or hydric buffering. Precipitation affiliation can also couple directly to land-use response, sorting species between forest and agriculture along rainfall gradients (Frishkoff et al. 2016). Where habitat-loss and climate effects on biodiversity have been tested together, designs have often used binary effect-direction outcomes, global macroclimatic predictors, or focused on population-level responses rather than assemblage richness across spatial scales (Mantyka-Pringle et al. 2012; Williams et al. 2022; Oliver and Morecroft 2014). Long-term population data show that synergistic climate–land-use effects can be misattributed or missed when only one driver is assessed (Northrup et al. 2019). Land use, climate, and topography operate at different but overlapping spatial scales, and whether climate-conditioned land-use effects are spatially coherent across a scale matrix or restricted to one aggregation grain is an open inferential question.

Here we test whether contemporary climatic context modifies the effects of land-use transformation on species richness across the Atlantic Forest. We use two complementary land-use surrogates — cropland prevalence and non-natural cover — to distinguish agriculture-specific conversion from broader anthropogenic transformation. We evaluate four climate moderators representing distinct exposure pathways, namely daily temperature range (DTR), annual temperature range (ATR), maximum temperature of the warmest month (MAXT), and precipitation of the driest quarter (PDRY). DTR captures the timescale most directly relevant to canopy-mediated buffering — daily thermal swings during which forest cover moderates exposure. ATR captures broader thermal seasonality, which acts over longer ecological timescales but indexes a similar variability dimension. Forest canopies buffer both peak

temperatures and the amplitude of thermal fluctuations, so canopy loss can expose organisms to higher maximum temperatures as well as greater daily and seasonal variability (De Frenne et al. 2019; Ewers and Banks-Leite 2013; Jucker et al. 2020). Because earlier global synthesis identified maximum temperature as the strongest climatic determinant of habitat-loss effects, while excluding temperature seasonality owing to collinearity with maximum temperature (Mantyka-Pringle et al. 2012), we retained MAXT as a competing peak-heat moderator and DTR and ATR as thermal-variability moderators, testing whether land-use damage was compounded primarily by peak heat or by thermal variability. Forest loss can also reduce precipitation at landscape-to-regional scales through disruption of moisture recycling and evapotranspiration, potentially increasing hydric stress where dry-season water availability already constrains forest-dependent organisms (Smith et al. 2023). PDRY indexes hydric stress during the driest quarter. We used these climate variables as atmospheric contexts in which land-use-driven loss of canopy buffering should be most consequential, not as direct measurements of under-canopy thermal variability or hydric exposure.

Two alternative pathways could generate apparent climate-conditioned land-use effects without requiring a direct canopy-buffering mechanism. First, realised water–energy flux may structure both richness and land-use placement. Productive landscapes can sustain larger populations and richer species pools, while also covarying with climatic exposure and agricultural suitability (Currie 1991; Hawkins et al. 2003; Evans et al. 2005; Storch et al. 2018; Hurlbert and Stegen 2014). We therefore used actual evapotranspiration (AET), the realised water–energy flux, as a resource-availability control to test whether land-use × climate interactions persisted beyond productivity-linked demographic capacity. Second, terrain can structure both climatic exposure and land-use placement. Rugged landscapes can provide environmental opportunity, microclimatic refugia, and species-sorting gradients, while flatter areas are more readily converted to agriculture (Stein et al. 2014; Hjort et al. 2015; Tukiainen et al. 2017; Bouchet et al. 2015; Räsänen et al. 2016). We therefore used topographic ruggedness (TOPOD) as an abiotic-template control to test whether LU × climate interactions were separable from terrain-structured land-use placement.

We additionally partition richness by forest affiliation, because the buffering hypothesis predicts that climate-conditioned land-use effects should be strongest in the assemblage components most dependent on forest structure. We distinguish forest-specialist, forest-

facultative, and non-forest richness, and analyse these guilds alongside aggregate richness. This partition avoids the arithmetic dependence of comparing forest specialists only with the total richness measure in which they are nested. It also provides a diagnostic comparison. Species with no listed forest habitat should not show the amplification predicted for forest-dependent species if the pattern reflects forest-affiliation filtering rather than a generic land-use or biogeographic gradient. The prediction is grounded in evidence that human-altered land uses filter assemblages toward species with broader or more extreme climatic affiliations, with forest-dependent taxa disproportionately reduced where land-use transformation and climatic exposure coincide (Williams et al. 2020; Sweeney and Jarzyna 2022).

Hypotheses and predictions

Building from this framework, we tested five predictions across a scale-explicit grain × extent design (Table 1.1). We predicted negative main effects of both land-use surrogates on richness (H1), negative land-use × DTR and land-use × ATR interactions if thermal variability amplified land-use effects (H2), negative land-use × PDRY interactions if dry-season stress amplified land-use effects (H3), negative land-use × MAXT interactions if peak heat, rather than thermal variability, compounded land-use damage (H4), and negative land-use × climate interactions ordered by forest dependence, from forest-specialist through forest-facultative to non-forest richness (H5). We paired these predictions with alternatives involving aggregate-richness masking by disturbance-tolerant or matrix-associated species, productivity gradients, terrain-structured land-use placement, biogeographic context, and taxonomic-subset effects, and we assessed whether support was scale-consistent or scale-restricted.

Methods

Study region and species data

The Atlantic Forest is one of the world's most species-rich and threatened tropical forests, spanning approximately 1.5 million km² along the Atlantic coast of South America, primarily in Brazil with extensions into Argentina and Paraguay (Ribeiro et al. 2009; Marques and Grelle 2021). We assembled 2,986 georeferenced species assemblages (Figure 1.1) derived from standardised field protocols with reported survey effort between 1981 and 2018, covering seven dataset-defined animal groups representing five broad taxonomic groups, namely mid–

large terrestrial mammals (Lima et al. 2017; Souza et al. 2019), small mammals (Bovendorp et al. 2017; Figueiredo et al. 2017), bats (Muylaert et al. 2017), passerine birds (Hasui et al. 2018), snakes (Moura et al. 2017a, b), anurans (Vancine et al. 2018), and fruit-feeding butterflies (Santos et al. 2018). We used the dataset-defined animal group as the TAXA random-effect grouping factor, with seven levels in the full-universe response strata and five in the restricted-universe strata.

We classified each species by forest affiliation from species-level habitat records in the IUCN Red List, coded under the IUCN Habitats Classification Scheme (Version 3.1; IUCN 2012). Following Donald et al. (2019), we classified species whose listed habitats fall entirely within the Forest class as forest specialists, species listing Forest together with one or more non-forest classes as forest-facultative, and species with no Forest habitat listed as non-forest. These categories describe exclusivity within the IUCN-listed habitat types and do not imply experimentally demonstrated habitat obligacy.

We used this classification to construct five response strata across two taxonomic universes. The full universe spans all seven dataset-defined groups (total_full). The restricted universe spans the five habitat-classifiable dataset-defined groups — mid–large mammals, small mammals, bats, passerines, and anurans — and excludes butterflies and snakes for different reasons. Snakes enter the datasets as analysis-derived richness estimates rather than species-level records, so individual species cannot be matched to IUCN habitat records, and fruit-feeding butterflies, though resolved to species, are not comprehensively assessed on the IUCN Red List. Within the restricted universe we modelled total richness (total_restr) and three mutually exclusive forest-dependence guilds — forest-specialist, forest-facultative, and non-forest richness — which partition it exactly, so that total_restr = specialist + facultative + non_forest. We classified 1,199 species into 302 forest specialists, 702 forest-facultative species, and 195 non-forest species. Restricting the response to the five habitat-classifiable groups did not create or reverse the baseline land-use and climate associations, with univariate coefficients agreeing in sign with the full universe at 99.1% of predictor–scale combinations (107 of 108; Appendix 1.S1.1). We fitted interaction models for all five response strata.

Multi-level design and spatial aggregation

We implemented a multilevel, multiscale design to examine how ecological processes vary across hierarchical spatial contexts (McGarigal et al. 2016). We aggregated species assemblages by METHOD into spatial response units at three nested grains (10, 25, and 40 km) and characterised landscape-level predictors at seven extents (10, 25, 40, 55, 70, 85, and 100 km). The hexagonal units serve as an aggregation template; predictor extents are centred on the centroid of the assemblages aggregated into each record rather than on the geometric centre of the hexagon. The factorial of grains and extents (with extent ≥ grain) generated 18 scale combinations, allowing inference about both grain-restricted and extent-coherent moderation patterns. The canopy-buffering mechanism, response-stratum design, land-use and climate predictor families, and scale matrix are summarised in Figure 1.2.

Land-use surrogates

We indexed land-use transformation with two complementary composition surrogates that differed in temporal anchors and resolutions. Cropland prevalence (CROP) and non-natural cover (NNCOV) index how much of a landscape has been converted, not how intensively it is managed. CROP was derived from the 1-km global consensus land-cover product centred on the temporal core of biodiversity sampling (1992–2006; Tuanmu and Jetz 2014). NNCOV was derived from the regional 30-m Atlantic Forest land-cover product for 2020 as the complement of natural vegetation (Souza et al. 2020). We used their convergence to assess whether the inferred pattern persisted across two distinct operationalisations of land-use transformation that differed in ecological definition, source product, resolution, classification framework, and temporal anchoring.

Climate moderators

We represented contemporary climate using four variables capturing thermal variability, maximum heat context, and dry-season hydric stress (Hijmans et al. 2005; Title and Bemmels 2018). Daily temperature range (DTR; Bio02) indexed diurnal thermal variability and was treated as the more mechanistically proximate proxy for short-term canopy-buffering exposure. Because forest canopies buffer this diurnal range relative to open-air conditions (Ismaeel et al. 2024), atmospheric DTR indexes the variability increment to which organisms are exposed when canopy cover is removed. Annual temperature range (ATR; Bio07) indexed broader

thermal seasonality and was retained as a longer-timescale variability moderator, not as a redundant comparator. Maximum temperature of the warmest month (MAXT; Bio05) indexed peak-heat context and was retained as a competing thermal moderator against DTR and ATR. Precipitation of the driest quarter (PDRY; Bio17) indexed dry-season hydric limitation; we multiplied PDRY by −1 so that higher values represent drier conditions, aligning its sign with the other stress variables. Annual precipitation and precipitation seasonality were excluded from the focal moderator set because annual precipitation was strongly coupled to the water–energy axis (AET) and precipitation seasonality represented an overlapping seasonal-hydric axis (Appendix 1.S2).

Additive controls

We selected AET and TOPOD a priori as hypothesis-based additive controls. AET indexes realised water–energy flux, a primary broad-scale positive correlate of species richness and a proxy for resource-mediated demographic capacity (Hawkins et al. 2003; Currie et al. 2004). It was derived as the product of potential evapotranspiration (Title and Bemmels 2018) and the Priestley–Taylor coefficient ALPHA (Trabucco and Zomer 2010) prior to z-transformation; it therefore tests whether apparent LU × climate moderation reflects productivity gradients rather than climate-conditioned land-use damage. TOPOD (Amatulli et al. 2018) indexes persistent terrain structure, which can provide microclimatic refugia and environmental opportunity expected to support richness (Stein et al. 2014; Hjort et al. 2015) while also constraining land-use placement because rugged areas are less readily converted; it therefore tests whether apparent LU × climate moderation reflects terrain-structured conversion rather than climatic exposure alone. These controls were included additively only — we did not fit second-tier productivity moderation or three-way LU × CLIM × AET interactions in this study. AET and TOPOD entered all controlled models as linear fixed-effect terms without polynomial expansion or random-slope selection, consistent with their empirically consistent additive associations across taxa and their role as conditioning rather than focal predictors. ALPHA was retained as a sensitivity proxy for AET, and we never fitted AET and ALPHA in the same model because ALPHA enters the construction of AET and the two predictors were correspondingly strongly correlated ($r \approx +0.81$).

Statistical modelling

Predictor geometry and baseline specification

We evaluated each land-use × climate interaction against a priori, mechanism-based predictions and an explicit support logic (Appendix 1.S3). Before fitting interaction models, we characterised predictor coupling across the full grain × extent matrix using pairwise Pearson correlations and variance inflation factors. Standard linear collinearity diagnostics, typically $r < 0.7$ and $VIF < 5$, are insufficient to detect all gradient-sharing problems in multi-predictor systems where two predictors can be moderately correlated pairwise yet share a dominant spatial gradient that systematically distorts additive partial coefficients (Dormann et al. 2013). We additionally used predictor-coupling, scale-stability, nonlinearity and sign-diagnostic summaries to identify climatic moderators whose interaction products with the land-use surrogates might reflect broad geographic co-alignment rather than separable biological moderation. These diagnostics informed interpretation but did not determine model support; support was based on model-fit improvement and product-term structure-coefficient concordance rather than on coupling geometry alone. Full coupling profiles and pairwise correlations at representative scale combinations are reported in Appendix 1.S2.

We modelled observed species richness using generalised linear mixed-effects models with a Poisson distribution, log link and an observation-level random effect (OLRE) for overdispersion (Harrison 2014; O'Hara and Kotze 2010). All focal continuous predictors were z-transformed before fitting, so coefficients represent changes per standard-deviation change in the corresponding predictor (Schielzeth 2010). The core maximal candidate interaction model can be represented as:

$$\text{resp} \sim \text{EFFsb} + \text{LU} + \text{CLIM} + \text{LU:CLIM} + \text{space}$$

$$+ (1 \mid \text{METHOD}) + (1 + \text{LU} + \text{CLIM} \mid \text{TAXA}) + (1 \mid \text{obs})$$

where EFFsb is site-based sampling effort, LU is either cropland prevalence or non-natural cover, CLIM is one of the four climatic moderators, space is the residual spatial autocovariate described below, METHOD accounts for baseline differences among sampling protocols, TAXA represents the dataset-defined animal groups, and the observation-level intercept carries overdispersion. Polynomial terms were added where supported by the preceding functional-form

selection. The displayed TAXA term represents the fullest candidate random-slope structure; the fitted structure was selected separately for each response stratum, land-use × climate family and scale cell as described below. Appendix 1.S4 provides the complete specification and diagnostics. We carried out all analyses in R (R Core Team 2024) and fitted every model with the glmmTMB package (Brooks et al. 2017).

Before hypothesis testing, we selected the nuisance structure that would be held fixed during each focal comparison. For each response stratum, land-use × climate family and grain × extent cell, we compared four TAXA-level random-effect candidates: random intercepts only, a random slope for land use, a random slope for the climatic moderator, and correlated random slopes for both. We selected the simplest structure within $\Delta\text{AICc} \leq 2$ of the lowest-AICc candidate (Burnham and Anderson 2002). Polynomial terms for land use and climate were selected in a preceding step and then inherited unchanged by the corresponding full and reduced interaction models. We selected polynomial and random-slope structures independently for each response stratum, land-use × climate family, and scale cell. The selected polynomial and random-effect structure was retained across the corresponding primary, controlled and reduced models, except where convergence required a simpler structure, in which case the cell was flagged. This sequencing prevents the focal interaction comparison from being confounded with simultaneous changes in functional form or random-effect structure. Because these specifications could differ among cells, we interpreted coefficient magnitudes across the interaction matrix only where model specification was held fixed, and otherwise summarised the cell-specific selected models through coefficient direction, support and control retention.

We assessed sensitivity to sampling-effort formulation across all five response strata and to temporal aggregation for total richness at the 10-km grain before interaction testing. These audits supported carrying the site-based effort covariate and year-aggregated responses into the focal analysis; their construction and agreement summaries are reported in Appendix 1.S4. We used the Poisson-OLRE family throughout and assessed its adequacy using simulation-based scaled-quantile residual diagnostics rather than treating an alternative count family as part of the focal model-selection sequence.

Land-use and interaction model sequence

We first fitted moderator-free land-use models to estimate the main land-use relationships. These models omitted climatic moderators and land-use × climate interactions. Adjusted models included sampling effort, the land-use surrogate, AET, TOPOD, the residual spatial autocovariate, TAXA and METHOD random intercepts, and the OLRE; unadjusted counterparts omitted AET and TOPOD to gauge control-driven attenuation. We selected land-use functional form on the adjusted model, because functional form defines the response surface that carries the main-effect inference. For each response stratum, land-use surrogate, response grain, and predictor extent, we compared linear and quadratic land-use forms and retained quadratic terms where they improved fit consistently across the scale matrix. Cropland was retained as linear in every response stratum, whereas non-natural cover was quadratic in forest-specialist and aggregate richness, borderline but retained linear in forest-facultative richness, and linear in non-forest richness. Because the two surrogates therefore differed in functional form, we expressed their effects as predicted contrasts in log-richness across the shared 10th–90th percentile interval of each surrogate rather than comparing raw coefficients. Because selected random-effect structures differed among strata, we estimated these cross-stratum contrasts under a common random-intercept structure, keeping the comparison estimand identical across guilds and aggregate anchors; stratum-specific random-slope fits are reported as sensitivity analyses in Appendix 1.S4.

The focal climatic analysis then followed two nested tiers. The Tier 1 model estimated each LU × CLIM interaction using the selected polynomial and random-effect structure together with sampling effort and the residual spatial autocovariate. Its reduced comparator retained the same land-use and climate main effects, polynomial terms, sampling structure and spatial adjustment but removed only the LU × CLIM product term. Tier 2 asked whether an interaction retained support after conditioning on broader water–energy and terrain structure. It added AET and TOPOD as additive fixed effects to both the full and reduced models; the reduced Tier 2 model again differed only by omission of the LU × CLIM interaction. We did not fit LU × AET, LU × ALPHA or LU × CLIM × AET interactions in this study.

Spatial adjustment was constructed separately for each relevant model family rather than by fitting a coordinate smooth jointly with the environmental predictors. From the corresponding full non-spatial model, we extracted Pearson residuals and fitted a thin-plate spline of those

residuals against geographic coordinates using the mgcv package (Crase et al. 2012; Wood 2004). The fitted spline values were z-transformed and entered as a fixed-effect spatial autocovariate in the spatial models.

We used this residual approach rather than fitting a smooth of geographic coordinates jointly with the environmental predictors because the focal environmental gradients are themselves strongly spatially structured across the Atlantic Forest. A jointly estimated coordinate smooth can compete with those predictors for shared broad-scale variation, making their coefficients conditional on a flexible spatial surface that may absorb part of the environmental structure the analysis is intended to estimate. By first fitting the hypothesised environmental model and deriving the spatial term from its residuals, we instead targeted spatial dependence remaining after the measured predictors had been given the opportunity to explain the response. The autocovariate therefore provides defensive adjustment for residual spatial structure; it does not identify the processes generating that structure or guarantee removal of bias from unmeasured environmental variables.

For controlled Tier 2 models, we recomputed the spatial autocovariate from Pearson residuals of the corresponding non-spatial controlled model rather than carrying it forward from the primary model, ensuring that AET- and TOPOD-structured spatial variation was not absorbed into the autocovariate before the controlled fit and that within-tier model comparisons rested on a consistent spatial footing (Fletcher and Fortin 2018). Within each tier, the full and reduced models used the same spatial autocovariate, so their comparison differed in the focal interaction rather than in spatial adjustment. Sensitivity comparisons of model coefficients with and without the spatial autocovariate are reported in Appendix 1.S4 to verify that the LU and LU × CLIM effects are not artefacts of spatial adjustment.

Evidence, robustness and model adequacy

AICc served two distinct roles in this sequence. During random-slope selection, candidates within $\Delta AICc \leq 2$ of the lowest-AICc model were treated as statistically indistinguishable and we retained the simplest member of that set. During hypothesis testing, by contrast, an interaction had to improve fit by more than two AICc units relative to its matched reduced model to earn primary support. Keeping these two uses separate prevents model-component selection from being conflated with evidence for the ecological interaction itself.

We complemented this parsimony criterion with coefficient-geometry diagnostics. For each predictor term we calculated the structure coefficient, SC = cor($\hat{\eta}$, x*), where $\hat{\eta}$ is the fitted linear predictor and x* is the z-transformed predictor or product term (Ray-Mukherjee et al. 2014). SC is a geometric reliability diagnostic, not an effect size. The standardised partial regression coefficient (BW) is the primary effect-size measure, while SC indicates whether the fitted surface associates with that term in the same direction as the partial slope. Sign concordance between SC and the coefficient indicates that the term is doing geometrically coherent work in the linear predictor; sign discordance indicates a suppressor geometry in which a collinear predictor dominates $\hat{\eta}$ and the coefficient cannot be interpreted ecologically.

Primary support for a LU × CLIM interaction therefore required two conditions: the full model had to improve on its matched reduced model by ΔAICc > 2, and the product-term SC had to be sign-concordant with the interaction coefficient. Likelihood-ratio-test $p < 0.05$ was retained as a secondary consistency check where nested comparisons were available, rather than as the primary decision rule. An interaction that cleared only the likelihood-ratio threshold without ΔAICc > 2, or that produced an SC-discordant product term, was not treated as primary-supported. This pre-specified combination of model-fit and coefficient-geometry criteria was applied across the complete scale factorial rather than selecting interactions according to individual p-values.

We then tested robustness only for primary-supported interactions. An interaction was classified as controls-robust when the Tier 2 full model retained a parsimony gain of ΔAICc > 2 over its controlled reduced comparator after AET and TOPOD entered. We also tracked sign stability and product-term SC concordance under controls as diagnostics, which distinguished interactions that retained model-selection support from interactions whose direction remained stable but whose parsimony support was absorbed by the control set. Because AET and TOPOD entered jointly, this comparison tests survival under conditioning and does not identify which control is responsible for any attenuation.

Because the hierarchical random-effect structure absorbs the dominant baseline heterogeneity across dataset-defined groups and survey methods, fixed-effect variance shares are structurally small, so we summarise interaction support through coefficient direction, structure-coefficient concordance and bootstrap intervals rather than through variance partitioning

(Nakagawa and Schielzeth 2013). Fixed-effect inference was supported by parametric bootstrap with full distributional summaries, reporting standardised partial coefficients (BW, bootstrap median and 95% interval) and structure coefficients (SC), as defined and verified against the bootstrap script in Appendix 1.S5 (Fieberg et al. 2020). Predictive adequacy — predictive p-values, coverage and relative RMSE, with relative RMSE not interpreted when estimates were close to zero — is reported in Appendix 1.S5. Residual diagnostics followed a simulation-based framework using scaled quantile residuals (DHARMa; Hartig 2022; Appendix 1.S4). We interpreted recurrence across the 18 grain × extent combinations as scale coherence rather than independent replication, because the cells contain nested response grains and overlapping predictor extents. Appendix 1.S4 summarises this sequence schematically (Figure 1.S4.1).

We used predictor diagnostics to characterise the planned moderator set. Because DTR and ATR were strongly correlated, we fitted them in separate models to keep their partial coefficients interpretable. Their concordant results are therefore interpreted as evidence concerning a thermal-variability axis rather than as independently identified daily- and annual-timescale mechanisms. We report pairwise correlations and variance inflation factors at the representative scale combination in Appendix 1.S2.

Finally, the analysis assumes that spatial contrasts in long-term land-use and climate conditions capture the dominant landscape context experienced by assemblages, despite temporal mismatch among biodiversity surveys (1981–2018), land-cover products (1992–2006 and 2020), and WorldClim normals (1970–2000). A second temporal limitation is internal to the biodiversity data: the effort covariate counts consolidated source assemblages and their reported survey effort, but not the number of years over which those surveys accumulated, which the modelling data do not retain.

*Abbreviations used in figures and tables***.** AET, actual evapotranspiration; ALPHA, Priestley–Taylor coefficient; ATR, annual temperature range; BW, standardised partial regression coefficient; CROP, cropland prevalence; DTR, daily temperature range; MAXT, maximum temperature of the warmest month; NNCOV, non-natural cover; PDRY, precipitation of the driest quarter; SC, structure coefficient; TOPOD, topographic ruggedness. Full definitions and provenance are given in Appendix 1.S2.

Results

Land-use conversion reduced species richness, and the strength of that reduction scaled with forest dependence (Figure 1.3). We express each land-use effect as the predicted change in log-richness across the shared 10th–90th-percentile interval of the land-use surrogate, held identical across strata so the contrasts remain comparable across guilds and functional forms. Forest specialists declined most steeply under both surrogates (cropland −0.323, non-natural cover −0.366), forest-facultative species declined less (−0.263 and −0.179), and non-forest richness declined least under cropland (−0.150) and showed no decline under non-natural cover (+0.065). The two aggregate anchors fell between the forest specialists and the weaker guilds (total_full −0.318 and −0.324, total_restr −0.290 and −0.289), so pooled richness diluted the specialist signal rather than reproducing it. This ordering was monotone across the three mutually exclusive guilds under both surrogates, matching the forest-dependence gradient predicted under H1.

Before applying the support criterion, we examined all fitted land-use × climate interaction coefficients across the 18 grain × extent combinations in all eight families (Figure 1.4). In the four thermal-variability families, fitted coefficients were negative in all three mutually exclusive guilds under both land-use surrogates, so the negative thermal-variability direction was present before support filtering. Because we selected polynomial and random-slope structures independently for each response stratum and scale cell, we read these distributions for direction and geometry and do not compare their magnitudes across strata. The hydric families differed between surrogates. CROP × PDRY showed a negative specialist-centred distribution before filtering and met the primary support criterion at all 18 scale combinations, but none retained controls-robust support after joint AET + TOPOD adjustment. NNCOV × PDRY met the primary criterion at no scale combination. MAXT was directionally surrogate-dependent, with forest-specialist coefficients predominantly positive under non-natural cover and negative under cropland. Together, the pre-filter audit and controlled analysis distinguished a recurrent thermal-variability pattern from a surrogate-specific, control-absorbed dry-season pattern and divergent peak-heat geometry. We report the full coefficient matrix in Appendix 1.S5.

Primary support concentrated strongly in forest specialists (Table 1.2). Across the eight land-use × climate families and 18 scale combinations per response stratum, 87 of 144 cells met

the support criterion in forest-specialist richness, compared with 20 and 23 in the two aggregate anchors, 14 in forest-facultative richness, and one in non-forest richness. Within the forest-specialist stratum, support concentrated in the thermal-variability families, which accounted for 62 of 72 thermal cells. We recovered primary support for cropland × daily temperature range at 16 of 18 combinations, cropland × annual temperature range at 17, non-natural-cover × daily temperature range at 13, and non-natural-cover × annual temperature range at 16. We recovered this strong concentration of thermal-variability support in forest specialists, but we did not establish the predicted monotonic ordering of interaction magnitude across the forest-dependence guilds. In matched-specification sensitivity analyses, which refitted the four thermal families across the three mutually exclusive guilds under a common random-intercept structure, the ordering changed materially between two pre-specified common functional forms, so we treat that component of H5 as specification-sensitive.

After we added AET and TOPOD, we found that every primary-supported forest-specialist thermal-variability cell retained the parsimony gain, with controls-robust support of 16/16, 17/17, 13/13, and 16/16, respectively. Across all 72 fitted forest-specialist DTR and ATR cells, we observed no sign reversals after adding AET and TOPOD, but we retained the primary support criterion for inference (Appendix 1.S4.6). Across matched fitted-layer coefficients, we generally recovered stronger forest-specialist thermal interactions after adding AET and TOPOD. The mean controlled BW coefficients were consistently negative, at −0.094, −0.111, −0.183, and −0.176. Thermal-variability moderation therefore remained after accounting additively for water–energy flux and terrain structure. We illustrate these controlled interactions on the response scale at the 25-km response grain and 25-km predictor extent in Figure 1.6.

The forest-specialist thermal-variability support was scale-coherent (Figure 1.5). The supported combinations reported above recurred across the 18-cell grain × extent matrix, appearing across multiple response grains and predictor extents rather than being confined to a single isolated scale cell. Because the 18 cells are nested realisations of the same assemblages aggregated at different grains and extents, they are statistically dependent, so we read this recurrence as scale coherence rather than as independent replication. This recurrence across the dependent scale matrix reduces the likelihood that forest-specialist thermal moderation reflects an isolated scale-specific artefact. By contrast, peak-heat support was surrogate-restricted and partly positive, whereas dry-season support was primary-scale-coherent but absorbed by the joint

AET + TOPOD control set. These contrasts distinguish the scale-coherent thermal-variability signal from the peak-heat and dry-season alternatives.

Two diagnostics clarified the weaker-guild results. For forest-facultative richness, we found 16 cropland cells that cleared the parsimony threshold, but all 16 were structure-coefficient-discordant and therefore failed the support criterion. We treated this as suppressor geometry, because the interaction coefficient and its structure coefficient opposed each other in sign. We therefore restricted interpretable forest-facultative interactions to the non-natural-cover families, with support at seven daily-temperature-range, six annual-temperature-range, and one peak-heat combinations. For non-forest richness, we recovered primary support at a single cell, and its fitted interaction coefficients under non-natural cover fell within the range spanned by forest-facultative richness under the respective selected models (Figure 1.4). The near-absence of non-forest support was consistent with lower precision rather than a clean absence of interaction geometry, because non-forest standard errors were approximately 1.8 times larger than forest-facultative standard errors under non-natural cover, so comparable coefficients less often cleared the parsimony threshold. These patterns identify two distinct limitations. Suppressor geometry limited forest-facultative cropland inference, whereas low precision limited non-forest support.

Maximum temperature did not reproduce the thermal-variability pattern. Under cropland, MAXT support appeared only in total_full richness, where CROP × MAXT reached primary and controls-robust support at 4 of 18 combinations with a negative controlled coefficient (−0.037). We recovered no forest-specialist CROP × MAXT support. Under non-natural cover, supported MAXT rows were positive wherever they appeared, in every stratum in which they reached support. NNCOV × MAXT reached primary support at 7 of 18 forest-specialist combinations and retained controls-robust support at 6 of those 7, but its controlled coefficient was positive (+0.079), opposite to the predicted compound-damage direction. We therefore interpret MAXT as a thermal-axis discrimination result. In this observational geometry, coherent land-use amplification emerged along daily and annual thermal-variability axes rather than along peak heat.

Dry-season stress behaved differently again. In forest-specialist richness, CROP × PDRY reached primary support at all 18 scale combinations, but none retained controls-robust support after we added AET and TOPOD. The joint control set therefore absorbed the parsimony gain

without reversing the direction of the fitted interaction. Aggregate CROP × PDRY rows retained partial controls-robust support, with 6 of 7 total_full cells and 4 of 7 total_restr cells retained after controls, so we do not treat PDRY as globally absorbed across all response strata. For the forest-specialist inference, however, dry-season amplification was not retained as a controls-robust moderator. Because AET and TOPOD entered jointly, this design cannot identify which control absorbed the specialist CROP × PDRY support.

The aggregate anchors behaved as pooled composites of the guild responses. Their total support counts, 20 of 144 cells in total_full and 23 of 144 in total_restr, fell between the concentrated forest-specialist support and the weaker guilds. This pattern matches the main-effect result, where pooled richness diluted rather than reproduced the forest-specialist signal. Restricting the response universe did not create the thermal-moderation pattern. Full- and restricted-universe aggregate thermal coefficients agreed in sign across all 72 thermal cells, and the few sign differences across the full moderator matrix occurred in maximum-temperature or dry-season cells where coefficients were near zero. We therefore interpret the aggregate anchors as dilution checks, not as alternative locations of the main signal.

Discussion

We found that land-use effects on Atlantic Forest animal richness were both forest-dependent and climatically conditional. This result aligns with broader evidence that anthropogenic climate change and land-use change jointly reshape biodiversity across taxa, but it locates that interaction within a scale-explicit richness framework and shows that the strongest signal is concentrated in forest-dependent assemblage components (Montràs-Janer et al. 2024). Baseline land-use filtering followed the expected forest-dependence gradient. Forest-specialist richness declined most strongly with cropland and non-natural cover, forest-facultative richness declined less, and non-forest richness responded least. The aggregate anchors carried fewer supported interactions than forest specialists, showing that pooled richness diluted rather than represented the forest-specialist signal.

Thermal-variability amplification showed a sharper pattern. Support for it concentrated in forest specialists rather than forming a smooth monotone continuum across all guilds. On the response scale, the marginal cropland slope for forest specialists steepened across the thermal-

variability gradient, and predicted forest-specialist richness declined steeply with cropland where thermal variability was high but remained nearly flat where it was low (Figure 1.6). We observed negative interaction directions across guilds before support filtering, and we recovered recurrent forest-specialist support across the dependent scale matrix that remained after we added the joint AET + TOPOD control set. The weaker guilds clarified the inference rather than simply weakening it. Forest-facultative richness retained interpretable support only under non-natural cover because cropland interactions showed suppressor geometry, whereas non-forest richness showed low precision rather than a well-supported thermal response.

The concentration of thermal-variability amplification in forest specialists is consistent with a canopy-buffering mechanism. Intact forest canopies can moderate near-ground thermal extremes, dampen diurnal and seasonal variation, and stabilise under-canopy humidity relative to more open land uses (De Frenne et al. 2019; Ewers and Banks-Leite 2013). Land-use conversion can remove or thin this canopy layer, exposing assemblages to the broader thermal variability captured by DTR and ATR (Jucker et al. 2020). Because we classified forest specialists as species whose listed habitat associations fall entirely within forest, this guild should be the component most closely tied to forest structure and the buffered conditions it provides. Our results match this expectation without requiring a complete guild continuum. We recovered negative, scale-coherent, controls-robust forest-specialist thermal-variability interactions. Forest-facultative richness retained interpretable thermal support only under non-natural cover, consistent with weaker dependence on closed-canopy buffering. Non-forest richness carried no supported thermal-variability cells, although lower precision limits interpretation of this absence. We therefore read the amplification pattern as the strongest evidence for climate-conditioned land-use filtering concentrating in the forest-dependent component of the assemblage, rather than acting uniformly across guilds.

DTR was the more mechanistically proximate proxy for canopy-buffering exposure because forest canopies operate over diurnal heating and cooling cycles. However, our prediction concerned thermal variability as an axis rather than a strict DTR-over-ATR hierarchy, and DTR and ATR were strongly correlated across the scale matrix (mean $r = 0.824$; Appendix 1.S2). We therefore interpret their concordant negative interactions as evidence for a thermal-variability amplification axis, not as separable daily- and annual-timescale mechanisms. This interpretation is consistent with multi-taxon evidence that macroclimate and microclimatic modulation can

interactively shape alpha diversity, even though our design tests land-use amplification rather than direct microclimate mediation (Geres et al. 2025).

Partitioning richness into forest-affiliation guilds was necessary to recover the climate-conditioned filtering that aggregate richness obscured. The two aggregate anchors carried fewer supported interactions than forest specialists, and their support counts fell between the specialist and weaker-guild responses. Because total richness pools forest specialists with forest-facultative and non-forest species, it averages the strong specialist response against weaker responses in the other components, so aggregate richness understates the effect on the forest-dependent guild. This dilution appeared in both the main-effect gradient and the interaction support, so it reflects a general property of pooled richness in this system rather than an artefact of a single analysis. The restricted and full aggregate anchors agreed closely, which shows that the dilution did not arise from the two dataset-defined groups lacking habitat-affiliation data. Guild partitioning therefore did more than add resolution. It changed the inference, because the conclusion that land-use effects are strongly climate-conditioned holds most clearly for forest specialists but pooling weaker-response guilds into aggregate richness progressively masks it. Analyses that rely on aggregate richness alone would under-detect this climate conditioning and could misread it as weak or scale-restricted.

This interpretation is consistent with broader evidence that disturbance and climate exposure can erode specialist components of assemblages while aggregate richness obscures those losses. Specialist-focused analyses are therefore useful because they retain ecological identity that total richness discards, even though richness remains incidence-based and does not resolve abundance, composition, or demographic rates (Devictor et al. 2008, 2010; Clavel et al. 2011; Fletcher et al. 2025). The pattern also matches evidence that forest loss can disproportionately erode biodiversity in landscapes where substantial forest cover remains, and that interacting anthropogenic pressure and climate warming can drive specialist losses while less sensitive species expand (Betts et al. 2017; Scharnhorst et al. 2026).

The peak-heat and dry-season results acted as discriminating cases that separated the thermal-variability signal from broader environmental structure. Maximum temperature did not reproduce the combination of negative direction and specialist-concentrated support shown by the variability moderators. Its supported interactions were surrogate-dependent and, under non-

natural cover, positive. We do not read the positive non-natural-cover × maximum-temperature interaction as evidence that warmth buffers land-use damage. Instead, we interpret it as a predictor-geometry result in this observational setting, where maximum temperature aligns with broad geographic and land-use context rather than isolating a peak-heat amplification mechanism. This interpretation is consistent with evidence that human-dominated land uses can favour species affiliated with warmer or more climatically extreme conditions, especially in tropical assemblages (Williams et al. 2020). The maximum-temperature result therefore provides a result-level thermal-axis discrimination. Diurnal and seasonal thermal variability, rather than the peak-heat axis, carried the coherent negative amplification of land-use effects.

Dry-season water stress discriminated the signal differently. Cropland × dry-season precipitation scantness reached primary support at every scale in forest specialists but retained no support once we conditioned on productivity and terrain. We therefore treat dry-season amplification as absorbed by the joint control set for the forest-specialist inference, rather than as a retained secondary moderator. Because actual evapotranspiration and topographic complexity entered together, this design cannot identify which control absorbed the specialist dry-season support, and we do not rank them. The dry-season result also strengthens the thermal-variability interpretation by contrast. The thermal-variability interactions remained after the same conditioning, so they were separable from the joint water–energy and terrain structure that absorbed the dry-season signal.

This result does not imply that maximum temperature is ecologically irrelevant in tropical systems. Maximum temperatures can condition habitat affiliations toward greater forest dependence where heat stress is high, as shown for North American mammals (Tourani et al. 2023), and an analogous mechanism may operate in tropical forests where many species occur near thermal limits. However, our design could not isolate peak-heat moderation from the spatial co-gradient among MAXT, land-use intensity, latitude, terrain, and biogeographic context. Whether peak heat acts as an additional moderator in tropical forest assemblages therefore remains an open question that this observational geometry cannot resolve.

This contrast shows why land-use × climate product terms require geometric diagnostics before mechanistic interpretation. Product-term support can reflect spatial co-alignment among land use, climate, and biogeographic context rather than separable ecological moderation. We

therefore interpret climatic moderation from the joint pattern of coefficient direction, structure-coefficient concordance, response-stratum concentration, control retention, and recurrence across the scale matrix, rather than from product-term support alone.

The control sequence clarified rather than erased the thermal-variability result. Adding AET and TOPOD retained the parsimony gain in all primary-supported specialist thermal-variability cells, whereas the same joint control set absorbed specialist CROP × PDRY support. This asymmetry shows that thermal-variability moderation was separable from the joint water–energy and terrain structure that absorbed the dry-season signal, even though the design cannot rank AET against TOPOD as the absorbing control.

Several design constraints bound our interpretation. First, we inferred climate-conditioned land-use effects from observational assemblage data and spatial environmental gradients rather than experimental manipulations, so we treat the results as evidence of conditional association rather than causation. Second, we used atmospheric climate layers to index thermal exposure rather than paired under-canopy measurements, a limitation we develop below. Third, we used the multiscale design to assess recurrence across analytical resolutions, not to identify the ecological scale at which canopy buffering operates. Because response grains were nested and predictor extents overlapped around many of the same assemblages, we could not separate ecological scale from predictor averaging and spatial overlap in the extent profile. Fourth, cropland and non-natural cover are broad land-use surrogates that combine habitat amount, matrix composition, edge exposure, and configuration, so we cannot isolate which land-use component drives the interaction. That decomposition matters because edge microclimates, landscape-level fragmentation signatures, and climate connectivity need not respond identically to forest loss. Edge microclimatic responses depend on matrix contrast and edge structure, fragmentation can produce landscape-scale thermal signatures that differ from forest loss alone, and forest loss can reduce climate connectivity in tropical forest landscapes (Arroyo-Rodríguez et al. 2017; Mendes and Prevedello 2020; Senior et al. 2019; Parks et al. 2020). Fifth, the AET and TOPOD controls entered jointly, so this design can show that shared water–energy and terrain structure absorbed the forest-specialist cropland × dry-season signal but cannot rank productivity against terrain as the absorbing control. Finally, the weaker-guild diagnostics constrain inference, because structure-coefficient discordance limited forest-facultative cropland interpretation and lower precision limited non-forest support. These constraints do not weaken

the central result, but they define its scope. We identify a scale-coherent, controls-robust thermal-variability signal concentrated in forest specialists, not a universal land-use × climate interaction across all guilds, climate axes, or land-use surrogates.

DTR, ATR, MAXT, and PDRY are atmospheric WorldClim-derived metrics, not direct measurements of realised under-canopy microclimate, which depends on vegetation structure, topography, and soil conditions (Kemppinen et al. 2024). The patterns reported here are consistent with loss of canopy-mediated buffering, but direct microclimatic data and explicit mediation analysis would be required to identify that pathway as a causal mechanism rather than a candidate explanation. The closest available evidence for this pathway in a Neotropical system is that Amazonian forest edges adjacent to pastures are measurably hotter and drier than edges adjacent to rubber plantations, and these microclimatic differences coincide with shifts in vegetation composition and structure, demonstrating co-variation between matrix context, local climate, and vegetation response without establishing formal mediation (Meza-Elizalde and Armenteras-Pascual 2021). Below-canopy tropical temperatures are already entering novel thermal regimes even within some intact contiguous forests (Trew et al. 2024), indicating that canopy buffering reduces but does not eliminate climatic risk — a limitation our atmospheric-proxy design cannot resolve. Observational climate–land-use studies remain vulnerable to non-independence among regional climate gradients, land-use placement, and local microclimate; fully resolving this requires sampling designs with direct microclimate measurements crossed against land-use and climate gradients (Redlich et al. 2022). The confounding of land-use-driven microclimatic effects with regional climate change is not unique to our design. Multiscale studies using plot-level thermologgers have explicitly noted that microclimatic conditions in anthropogenic habitats may be only weakly correlated with larger-scale regional climatic variation (Ganuza et al. 2022), confirming that atmospheric-proxy approaches index exposure context but not the local climatic reality organisms experience.

Finally, because we modelled richness rather than abundance, composition, or demographic rates, the scale-coherent richness patterns do not identify the population or compositional mechanisms that produced them.

Conservation implications

Our findings have implications for conservation strategies that rely on forest-cover thresholds (Banks-Leite et al. 2014; Arroyo-Rodríguez et al. 2020). The fitted biodiversity cost of land-use transformation was not uniform across climatic contexts. It was strongest in thermally variable landscapes and concentrated in forest specialists, while dry-season stress showed primary support that the joint control set absorbed. Conservation planning that assumes a single forest-cover threshold across climatic contexts may therefore underestimate biodiversity costs in thermally variable landscapes, especially for forest-dependent assemblage components. In practice, equal amounts of retained forest may not provide equal biodiversity protection across climatic contexts. Because the amplification pattern is consistent with canopy buffering, conserving forest cover, canopy structure, and local microclimatic refugia should be especially important where thermal variability is high. Forest structure can maintain microhabitat diversity and thermal refuges that buffer organisms against extremes (Gunderson 2023; Stark et al. 2023), and some forest areas already maintain relatively non-novel below-canopy temperature regimes that may function as current thermal refugia (Trew et al. 2024). As climate change reshapes thermal and hydric exposure across landscapes, the climatic contexts that condition land-use effects should be treated as dynamic conservation layers rather than fixed background conditions (Newbold 2018; Northrup et al. 2019; Williams et al. 2022).

Conclusions

We show that land-use effects on Atlantic Forest animal richness were climatically conditional and concentrated in the forest-dependent component of assemblages. Forest-specialist richness declined most strongly with land-use transformation, and high diurnal and annual thermal variability amplified those declines. Aggregate richness diluted this signal, forest-facultative richness retained weaker and surrogate-specific support, and non-forest richness showed no supported thermal-variability cells. Peak heat and dry-season stress did not reproduce the same controls-robust specialist-concentrated pattern. Our results therefore indicate that tropical forest biodiversity loss from land-use transformation depends on climatic context as well as forest-dependence structure within assemblages. Conservation assessments based only on aggregate richness or uniform land-use thresholds may underestimate losses in thermally variable landscapes and in the forest-specialist species most tied to closed-canopy conditions.

References

Amatulli G, Domisch S, Tuanmu M-N, et al. (2018) A suite of global, cross-scale topographic variables for environmental and biodiversity modelling. Scientific Data 5:180040.

Arroyo-Rodríguez V, Fahrig L, Tabarelli M, et al. (2020) Designing optimal human-modified landscapes for forest biodiversity conservation. Ecology Letters 23:1404-1420.

Arroyo-Rodríguez V, Saldaña-Vázquez RA, Fahrig L, Santos BA (2017) Does forest fragmentation cause an increase in forest temperature? Ecological Research 32:81-88.

Banks-Leite C, Pardini R, Tambosi LR, et al. (2014) Using ecological thresholds to evaluate the costs and benefits of set-asides in a biodiversity hotspot. Science 345:1041-1045.

Betts MG, Wolf C, Ripple WJ, et al. (2017) Global forest loss disproportionately erodes biodiversity in intact landscapes. Nature 547:441-444.

Bouchet PJ, Meeuwig JJ, Salgado Kent CP, et al. (2015) Topographic determinants of mobile vertebrate predator hotspots: current knowledge and future directions. Biological Reviews 90:699-728.

Bovendorp RS, Villar N, de Abreu-Junior EF, et al (2017) Atlantic small-mammal: a dataset of communities of rodents and marsupials of the Atlantic forests of South America. Ecology 98:2226.

Brooks ME, Kristensen K, van Benthem KJ, et al. (2017) glmmTMB balances speed and flexibility among packages for zero-inflated generalized linear mixed modeling. The R Journal 9:378-400.

Burnham KP, Anderson DR (2002) Model Selection and Multimodel Inference: A Practical Information-Theoretic Approach, 2nd edn. Springer-Verlag, New York.

Carnaval AC, Moritz C (2008) Historical climate modelling predicts patterns of current biodiversity in the Brazilian Atlantic forest. Journal of Biogeography 35:1187-1201.

Clavel J, Julliard R, Devictor V (2011) Worldwide decline of specialist species: toward a global functional homogenization? Frontiers in Ecology and the Environment 9:222-228.

Crase B, Liedloff AC, Wintle BA (2012) A new method for dealing with residual spatial autocorrelation in species distribution models. Ecography 35:879-888.

Currie DJ (1991) Energy and large-scale patterns of animal- and plant-species richness. The American Naturalist 137:27-49.

Currie DJ, Mittelbach GG, Cornell HV, et al. (2004) Predictions and tests of climate-based hypotheses of broad-scale variation in taxonomic richness. Ecology Letters 7:1121-1134.

De Frenne P, Zellweger F, Rodríguez-Sánchez F, et al. (2019) Global buffering of temperatures under forest canopies. Nature Ecology and Evolution 3:744-749.

Devictor V, Clavel J, Julliard R, et al. (2010) Defining and measuring ecological specialization. Journal of Applied Ecology 47:15-25.

Devictor V, Julliard R, Jiguet F (2008) Distribution of specialist and generalist species along spatial gradients of habitat disturbance and fragmentation. Oikos 117:507-514.

Donald PF, Arendarczyk B, Spooner F, Buchanan GM (2019) Loss of forest intactness elevates global extinction risk in birds. Animal Conservation 22:341-347.

Dormann CF, Elith J, Bacher S, et al. (2013) Collinearity: a review of methods to deal with it and a simulation study evaluating their performance. Ecography 36:27-46.

Edwards DP, Socolar JB, Mills SC, et al. (2019) Conservation of tropical forests in the Anthropocene. Current Biology 29:R1008-R1020.

Evans KL, Warren PH, Gaston KJ (2005) Species-energy relationships at the macroecological scale: a review of the mechanisms. Biological Reviews 80:1-25.

Ewers RM, Banks-Leite C (2013) Fragmentation impairs the microclimate buffering effect of tropical forests. PLoS ONE 8:e58093.

Fieberg JR, Vitense K, Johnson DH (2020) Resampling-based methods for biologists. PeerJ 8:e9089.

Figueiredo MSL, Barros CS, Delciellos AC, et al (2017) Abundance of small mammals in the Atlantic Forest (ASMAF): a data set for analyzing tropical community patterns. Ecology 98:2981.

Fletcher RJ Jr, Fortin M-J (2018) Spatial Ecology and Conservation Modeling. Springer, Cham.

Fletcher RJ Jr, Green RE, Bladon EK, Atkinson PW, Phalan BT, Williams D, Visconti P, Balmford A (2025) Beyond species richness for biological conservation. Conservation Letters 18:e13124.

Frishkoff LO, Karp DS, Flanders JR, et al. (2016) Climate change and habitat conversion favour the same species. Ecology Letters 19:1081-1090.

Ganuza C, Redlich S, Uhler J, et al. (2022) Interactive effects of climate and land use on pollinator diversity differ among taxa and scales. Science Advances 8:eabm9359.

Geres LS, Richter T, Seidl R, et al. (2025) Macro- and microclimate interactively shape species diversity of multiple taxa in mountain landscapes. Ecography 2025:e07984.

Gunderson AR (2023) Habitat degradation exacerbates the effects of anthropogenic warming by removing thermal refuges. Global Change Biology 29:6157-6158.

Harrison XA (2014) Using observation-level random effects to model overdispersion in count data in ecology and evolution. PeerJ 2:e616.

Hartig F (2022) DHARMa: Residual diagnostics for hierarchical (multi-level/mixed) regression models. R package version 0.4.6.

Hasui E, Metzger JP, Pimentel RG, et al (2018) ATLANTIC BIRDS: a data set of bird species from the Brazilian Atlantic Forest. Ecology 99:497.

Hawkins BA, Field R, Cornell HV, et al. (2003) Energy, water, and broad-scale geographic patterns of species richness. Ecology 84:3105-3117.

Hijmans RJ, Cameron SE, Parra JL, et al. (2005) Very high resolution interpolated climate surfaces for global land areas. International Journal of Climatology 25:1965-1978.

Hjort J, Gordon JE, Gray M, et al. (2015) Why geodiversity matters in valuing nature's stage. Conservation Biology 29:630-639.

Hurlbert AH, Stegen JC (2014) When should species richness be energy limited, and how would we know? Ecology Letters 17:401-413.

Ismaeel A, Tai APK, Santos EG, et al. (2024) Patterns of tropical forest understory temperatures. Nature Communications 15:549.

IUCN (2012) Habitats Classification Scheme, Version 3.1. International Union for Conservation of Nature, Gland, Switzerland.

Jucker T, Jackson TD, Zellweger F, et al. (2020) A research agenda for microclimate ecology in human-modified tropical forests. Frontiers in Forests and Global Change 2:92.

Kemppinen J, Lembrechts JJ, Van Meerbeek K, et al. (2024) Microclimate, an important part of ecology and biogeography. Global Ecology and Biogeography 33:e13834.

Khaliq I, Hof C, Prinzinger R, et al. (2014) Global variation in thermal tolerances and vulnerability of endotherms to climate change. Proceedings of the Royal Society B 281:20141097.

Laurance WF, Useche DC, Rendeiro J, et al (2012) Averting biodiversity collapse in tropical forest protected areas. Nature 489:290-294.

Lima F, Beca G, Muylaert RL, et al. (2017) ATLANTIC-CAMTRAPS: a dataset of medium and large terrestrial mammal communities in the Atlantic Forest of South America. Ecology 98:2979.

Mantyka-Pringle CS, Martin TG, Rhodes JR (2012) Interactions between climate and habitat loss effects on biodiversity: a systematic review and meta-analysis. Global Change Biology 18:1239-1252. Correction (2013) Global Change Biology 19:1642-1644.

Marques MCM, Grelle CEV (eds) (2021) The Atlantic Forest: history, biodiversity, threats and opportunities of the mega-diverse forest. Springer, Cham.

McGarigal K, Wan HY, Zeller KA, Timm BC, Cushman SA (2016) Multi-scale habitat selection modelling: a review and outlook. Landscape Ecology 31: 1161-1175.

Mendes CB, Prevedello JA (2020) Does habitat fragmentation affect landscape-level temperatures? A global analysis. Landscape Ecology 35:1743-1756.

Meza-Elizalde MC, Armenteras-Pascual D (2021) Edge influence on the microclimate and vegetation of fragments of a north Amazonian forest. Forest Ecology and Management 498:119546.

Montràs-Janer T, Suggitt AJ, Fox R, et al. (2024) Anthropogenic climate and land-use change drive short- and long-term biodiversity shifts across taxa. Nature Ecology & Evolution 8:739-751.

Moura MR, Costa HC, Argôlo AJS, et al. (2017a) Environmental constraints on the compositional and phylogenetic beta-diversity of tropical forest snake assemblages. Journal of Animal Ecology 86:1192-1204.

Moura MR, Argôlo AJ, Costa HC (2017b) Historical and contemporary correlates of snake biogeographical subregions in the Atlantic Forest hotspot. Journal of Biogeography 44:640-650.

Muylaert RDL, Stevens RD, Esbérard CEL, et al (2017) ATLANTIC BATS: a data set of bat communities from the Atlantic Forests of South America. Ecology 98:3227.

Nakagawa S, Schielzeth H (2013) A general and simple method for obtaining R^2 from generalized linear mixed-effects models. Methods in Ecology and Evolution 4:133-142.

Newbold T (2018) Future effects of climate and land-use change on terrestrial vertebrate community diversity under different scenarios. Proceedings of the Royal Society B 285:20180792.

Newbold T, Hudson LN, Hill SLL, et al (2015) Global effects of land use on local terrestrial biodiversity. Nature 520:45-50.

Newbold T, Oppenheimer P, Etard A, Williams JJ (2020) Tropical and Mediterranean biodiversity is disproportionately sensitive to land-use and climate change. Nature Ecology and Evolution 4:1630-1638.

Northrup JM, Rivers JW, Yang Z, Betts MG (2019) Synergistic effects of climate and land-use change influence broad-scale avian population declines. Global Change Biology 25:1561-1575.

O'Hara RB, Kotze DJ (2010) Do not log-transform count data. Methods in Ecology and Evolution 1:118-122.

Oliver TH, Morecroft MD (2014) Interactions between climate change and land use change on biodiversity: attribution problems, risks, and opportunities. Wiley Interdisciplinary Reviews: Climate Change 5:317-335.

Parks SA, Carroll C, Dobrowski SZ, Allred BW (2020) Human land uses reduce climate connectivity across North America. Global Change Biology 26:2944-2955.

Pfeifer M, Lefebvre V, Peres CA, et al. (2017) Creation of forest edges has a global impact on forest vertebrates. Nature 551:187-191.

Prevedello JA, Winck GR, Weber MM, Nichols E, Sinervo B (2019) Impacts of forestation and deforestation on local temperature across the globe. PLoS ONE 14:e0213368.

R Core Team (2024) R: A Language and Environment for Statistical Computing. R Foundation for Statistical Computing, Vienna, Austria.

Ray-Mukherjee J, Nimon K, Mukherjee S, et al. (2014) Using commonality analysis in multiple regressions: a tool to decompose regression effects in the face of multicollinearity. Methods in Ecology and Evolution 5:320-328.

Redlich S, Zhang J, Benjamin C, et al. (2022) Disentangling effects of climate and land use on biodiversity and ecosystem services - a multi-scale experimental design. Methods in Ecology and Evolution 13:514-527.

Ribeiro MC, Metzger JP, Martensen AC, et al (2009) The Brazilian Atlantic Forest: How much is left, and how is the remaining forest distributed? Implications for conservation. Biological Conservation 142:1141-1153.

Räsänen A, Kuitunen M, Hjort J, et al. (2016) The role of landscape, topography, and geodiversity in explaining vascular plant species richness in a fragmented landscape. Boreal Environment Research 21:53-70.

Santos JP, Freitas AVL, Brown KS, et al (2018) ATLANTIC BUTTERFLIES: a data set of fruit-feeding butterfly communities from the Atlantic Forests. Ecology 99:2875.

Scharnhorst VS, Ockermüller E, Becsi B, et al. (2026) Generalist 'Winners' replace specialist 'Losers': interactions between regional anthropogenic impacts and climate warming drive long-term wild bee declines. Diversity and Distributions 32:e70151.

Schielzeth H (2010) Simple means to improve the interpretability of regression coefficients. Methods in Ecology and Evolution 1:103-113.

Senior RA, Hill JK, Edwards DP (2019) Global loss of climate connectivity in tropical forests. Nature Climate Change 9:623-626.

Senior RA, Hill JK, González del Pliego P, Goode LK, Edwards DP (2017) A pantropical analysis of the impacts of forest degradation and conversion on local temperature. Ecology and Evolution 7:7897-7908.

Smith C, Baker JCA, Spracklen DV (2023) Tropical deforestation causes large reductions in observed precipitation. Nature 615:270-274.

Souza CM Jr, Shimbo JZ, Rosa MR, et al. (2020) Reconstructing three decades of land use and land cover changes in Brazilian biomes with Landsat archive and Earth Engine. Remote Sensing 12:2735.

Souza Y, Gonçalves F, Lautenschlager L, et al (2019) ATLANTIC MAMMALS: a data set of assemblages of medium- and large-sized mammals of the Atlantic Forest of South America. Ecology 100:e02785.

Stark G, Ma L, Zeng Z-G, et al. (2023) Cool shade and not-so-cool shade: how habitat loss may accelerate thermal stress under current and future climate. Global Change Biology 29:6201-6216.

Stein A, Gerstner K, Kreft H (2014) Environmental heterogeneity as a universal driver of species richness across taxa, biomes and spatial scales. Ecology Letters 17:866-880.

Storch D, Bohdalková E, Okie J (2018) The more-individuals hypothesis revisited: the role of community abundance in species richness regulation and the productivity-diversity relationship. Ecology Letters 21:920-937.

Sunday JM, Bates AE, Kearney MR, et al. (2014) Thermal-safety margins and the necessity of thermoregulatory behavior across latitude and elevation. Proceedings of the National Academy of Sciences 111:5610-5615.

Sweeney CP, Jarzyna MA (2022) Assessing the synergistic effects of land use and climate change on terrestrial biodiversity: are generalists always the winners? Current Landscape Ecology Reports 7:41-48.

Title PO, Bemmels JB (2018) ENVIREM: an expanded set of bioclimatic and topographic variables increases flexibility and improves performance of ecological niche modeling. Ecography 41:291-307.

Tourani M, Sollmann R, Kays R, et al. (2023) Maximum temperatures determine the habitat affiliations of North American mammals. Proceedings of the National Academy of Sciences 120:e2304411120.

Trabucco A, Zomer RJ (2010) Global Soil Water Balance Geospatial Database. CGIAR Consortium for Spatial Information. Published online, available from the CGIAR-CSI GeoPortal at https://csidotinfo.wordpress.com

Trew BT, Edwards DP, Lees AC, et al. (2024) Novel temperatures are already widespread beneath the world's tropical forest canopies. Nature Climate Change 14:753-759.

Tuanmu M-N, Jetz W (2014) A global 1-km consensus land-cover product for biodiversity and ecosystem modelling. Global Ecology and Biogeography 23:1031-1045.

Tukiainen H, Bailey JJ, Field R, et al. (2017) Combining geodiversity with climate and topography to account for threatened species richness. Conservation Biology 31:364-375.

Vancine MH, Duarte K da S, de Souza YS, et al (2018) ATLANTIC AMPHIBIANS: a data set of amphibian communities from the Atlantic Forests of South America. Ecology 99:1692.

Williams JJ, Freeman R, Spooner F, et al. (2022) Vertebrate population trends are influenced by interactions between land use, climatic position, habitat loss and climate change. Global Change Biology 28:797-815.

Williams JJ, Bates AE, Newbold T (2020) Human-dominated land uses favour species affiliated with more extreme climates, especially in the tropics. Ecography 43:391-405.

Williams JJ, Newbold T (2020) Local climatic changes affect biodiversity responses to land use: a review. Diversity and Distributions 26:76-92.

Williams JJ, Newbold T (2021) Vertebrate responses to human land use are influenced by their proximity to climatic tolerance limits. Diversity and Distributions 27:1308-1323

Wood SN (2004) Stable and efficient multiple smoothing parameter estimation for generalized additive models. Journal of the American Statistical Association 99:673-686.

Table 1.1. Predictions tested in this study, with hypothesised mechanisms and alternative hypotheses.

Prediction	**Statement**	**Mechanism**	**Alternative**
H1 — LU main effects	CROP and NNCOV produce negative main effects on richness, and responses become increasingly negative as forest dependence increases from non-forest through forest-facultative to forest-specialist richness	Habitat conversion reduces viable population area and filters forest-dependent species through matrix hostility, reduced forest extent, and simplification of resources	Richness responses may be weak or non-directional if disturbance-tolerant or matrix-associated species partly offset losses of forest-dependent species at aggregate-richness scales. Land-use filtering may also act uniformly across guilds rather than scaling with forest dependence
H2 — Thermal variability moderation	LU × DTR and LU × ATR interactions are negative	Thermal variability exposure compounds canopy loss	Apparent thermal moderation reflects productivity gradients or terrain-structured land-use placement rather than climate-conditioned buffering loss; evaluated using the joint AET + TOPOD control set

Prediction	Statement	Mechanism	Alternative
H3 — Hydric stress moderation	LU × PDRY interaction is negative	Dry-season hydric stress compounds canopy loss	Confounding by water–energy gradients or terrain-structured land-use placement, evaluated using the joint AET + TOPOD control set
H4 — Peak-heat moderation and thermal-axis discrimination	LU × MAXT is expected to be negative if peak heat compounds land-use damage; comparing LU × MAXT with LU × DTR/ATR discriminates peak-heat from thermal-variability amplification	Canopy loss increases exposure to both higher peak temperatures and greater thermal variability; MAXT indexes the peak-heat axis, whereas DTR and ATR index thermal variability	Apparent LU × MAXT support may reflect geographic or land-use co-alignment rather than separable peak-heat moderation; evaluated using the joint AET + TOPOD control set and predictor-geometry diagnostics involving elevation, latitude, and forest cover

Prediction	Statement	Mechanism	Alternative
H5 — Forest-dependence scaling of climatic amplification	Negative LU × climate interactions strengthen with forest dependence, following the order forest specialist > forest facultative > non-forest	Forest specialists depend on canopy buffering and forest microclimate, so climate-conditioned land-use effects filter assemblage components in proportion to their forest dependence. Non-forest species are expected to depend less on forest-canopy buffering and therefore provide a diagnostic comparison	Amplification may act uniformly across guilds rather than scaling with forest dependence, or an apparent specialist–aggregate difference may arise from taxonomic restriction. We evaluate the first alternative using the three mutually exclusive guild responses within the same restricted universe, and the second by comparing the restricted and full aggregate anchors

All predictions are tested across response grains (10, 25, 40 km) and predictor extents (10–100 km) for both LU surrogates. Scale-consistent recovery of the predicted directions across the scale matrix reduces the likelihood that the result reflects an isolated scale-specific artefact; scale-restricted support is reported with the relevant scale domain. DTR, ATR, PDRY and MAXT represent distinct candidate climatic pathways through which land-use effects could be amplified, and H4 additionally compares absolute peak heat with thermal variability to discriminate between thermal axes.

Table 1.2. Focal land-use × climate interactions across forest-dependence strata. Primary support = grain × extent combinations (of 18) meeting the support criterion (ΔAICc > 2 and product-term structure-coefficient concordance). Controls-robust = the nested subset of primary-supported cells retaining ΔAICc > 2 after adding AET and TOPOD (n / n_primary). Controlled BW mean [bootstrap envelope] is the mean of the cell-specific controlled bootstrap medians across the primary-supported cells that retain sign stability and product-term structure-coefficient concordance under the controls, and the bracketed range is the envelope of their cell-specific 95% bootstrap intervals rather than a confidence interval for the mean. For total_full CROP × DTR that rule excludes one of the five controls-robust cells, whose controlled structure coefficient is indistinguishable from zero. Subtotal rows summarise support counts only; coefficients are not pooled across land-use surrogates or climate moderators. We use n/a for a subtotal with no primary-supported cells (n_primary = 0) and n.d. for a controlled comparison that we could not compute.

Stratum	**Interaction**	**Primary (n/18)**	**Controls-robust (n/n_primary)**	**Controlled BW mean [bootstrap envelope]**	**Interpretation / status**
total_full	CROP × DTR	5/18	5/5	-0.042 [-0.079, -0.012]	thermal-variability; controls-robust
	CROP × MAXT	4/18	4/4	-0.037 [-0.069, -0.004]	thermal-axis discrimination
	CROP × PDRY	7/18	6/7	-0.036 [-0.072, +0.005]	supported

Stratum	Interaction	Primary (n/18)	Controls-robust (n/n_primary)	Controlled BW mean [bootstrap envelope]	Interpretation / status
	NNCOV × DTR	2/18	2/2	-0.056 [-0.094, -0.020]	thermal-variability; controls-robust
	NNCOV × ATR	1/18	1/1	-0.052 [-0.081, -0.016]	thermal-variability; controls-robust
	NNCOV × MAXT	1/18	0/1	+0.012 [-0.010, +0.037]	thermal-axis discrimination (positive)
	Thermal subtotal	8/72	8/8	—	thermal-variability subtotal
	All-moderator subtotal	20/144	18/20	—	all-moderator subtotal
total_restr	CROP × PDRY	7/18	4/7	-0.031 [-0.071, +0.014]	supported
	NNCOV × DTR	1/18	1/1	-0.061 [-0.099, -0.021]	thermal-variability; controls-robust
	NNCOV × ATR	14/18	14/14	-0.064 [-0.126, -0.011]	thermal-variability; controls-robust

Stratum	Interaction	Primary (n/18)	Controls-robust (n/n_primary)	Controlled BW mean [bootstrap envelope]	Interpretation / status
	NNCOV × MAXT	1/18	0/1	+0.026 [+0.000, +0.056]	thermal-axis discrimination (positive)
	Thermal subtotal	15/72	15/15	—	thermal-variability subtotal
	All-moderator subtotal	23/144	19/23	—	all-moderator subtotal
forest-specialist	CROP × DTR	16/18	16/16	-0.094 [-0.162, -0.019]	thermal-variability; controls-robust
	CROP × ATR	17/18	17/17	-0.111 [-0.192, -0.022]	thermal-variability; controls-robust
	CROP × PDRY	18/18	0/18	-0.039 [-0.119, +0.041]	control-absorbed (joint AET + TOPOD)
	NNCOV × DTR	13/18	13/13	-0.183 [-0.284, -0.082]	thermal-variability; controls-robust
	NNCOV × ATR	16/18	16/16	-0.176 [-0.286, -0.070]	thermal-variability; controls-robust

Stratum	Interaction	Primary (n/18)	Controls-robust (n/n_primary)	Controlled BW mean [bootstrap envelope]	Interpretation / status
	NNCOV × MAXT	7/18	6/7	+0.079 [-0.005, +0.157]	thermal-axis discrimination (positive)
	Thermal subtotal	62/72	62/62	—	thermal-variability subtotal
	All-moderator subtotal	87/144	68/87	—	all-moderator subtotal
forest-facultative	NNCOV × DTR	7/18	7/7	-0.071 [-0.116, -0.025]	thermal-variability; controls-robust
	NNCOV × ATR	6/18	6/6	-0.066 [-0.114, -0.012]	thermal-variability; controls-robust
	NNCOV × MAXT	1/18	0/1	+0.022 [-0.011, +0.053]	thermal-axis discrimination (positive)
	Thermal subtotal	13/72	13/13	—	thermal-variability subtotal
	All-moderator subtotal	14/144	13/14	—	all-moderator subtotal

Stratum	Interaction	Primary (n/18)	Controls-robust (n/n_primary)	Controlled BW mean [bootstrap envelope]	Interpretation / status
non-forest	CROP × PDRY	1/18	n.d.‡	—‡	indeterminate under controls‡
	Thermal subtotal	0/72	n/a	—	thermal-variability subtotal
	All-moderator subtotal	1/144	n.d.‡	—	all-moderator subtotal

Forest-facultative shows no cropland rows because its 16 AICc-supported cropland cells are structure-coefficient-discordant (suppressor geometry; Appendix 1.S5) and fail the support criterion; interpretable facultative interactions are non-natural-cover only. Non-forest reaches primary support at a single cropland × dry-season cell. ‡ The controlled model for that cell did not converge to a positive-definite Hessian, so its controlled model comparison is not computable and its controlled bootstrap summary is withdrawn. We report its controlled status as not determined rather than as a loss of support. Its primary support is unaffected, resting on converged primary models (Appendix 1.S4.4, S4.6). Full five-stratum provenance and bootstrap detail are in Appendix 1.S5.

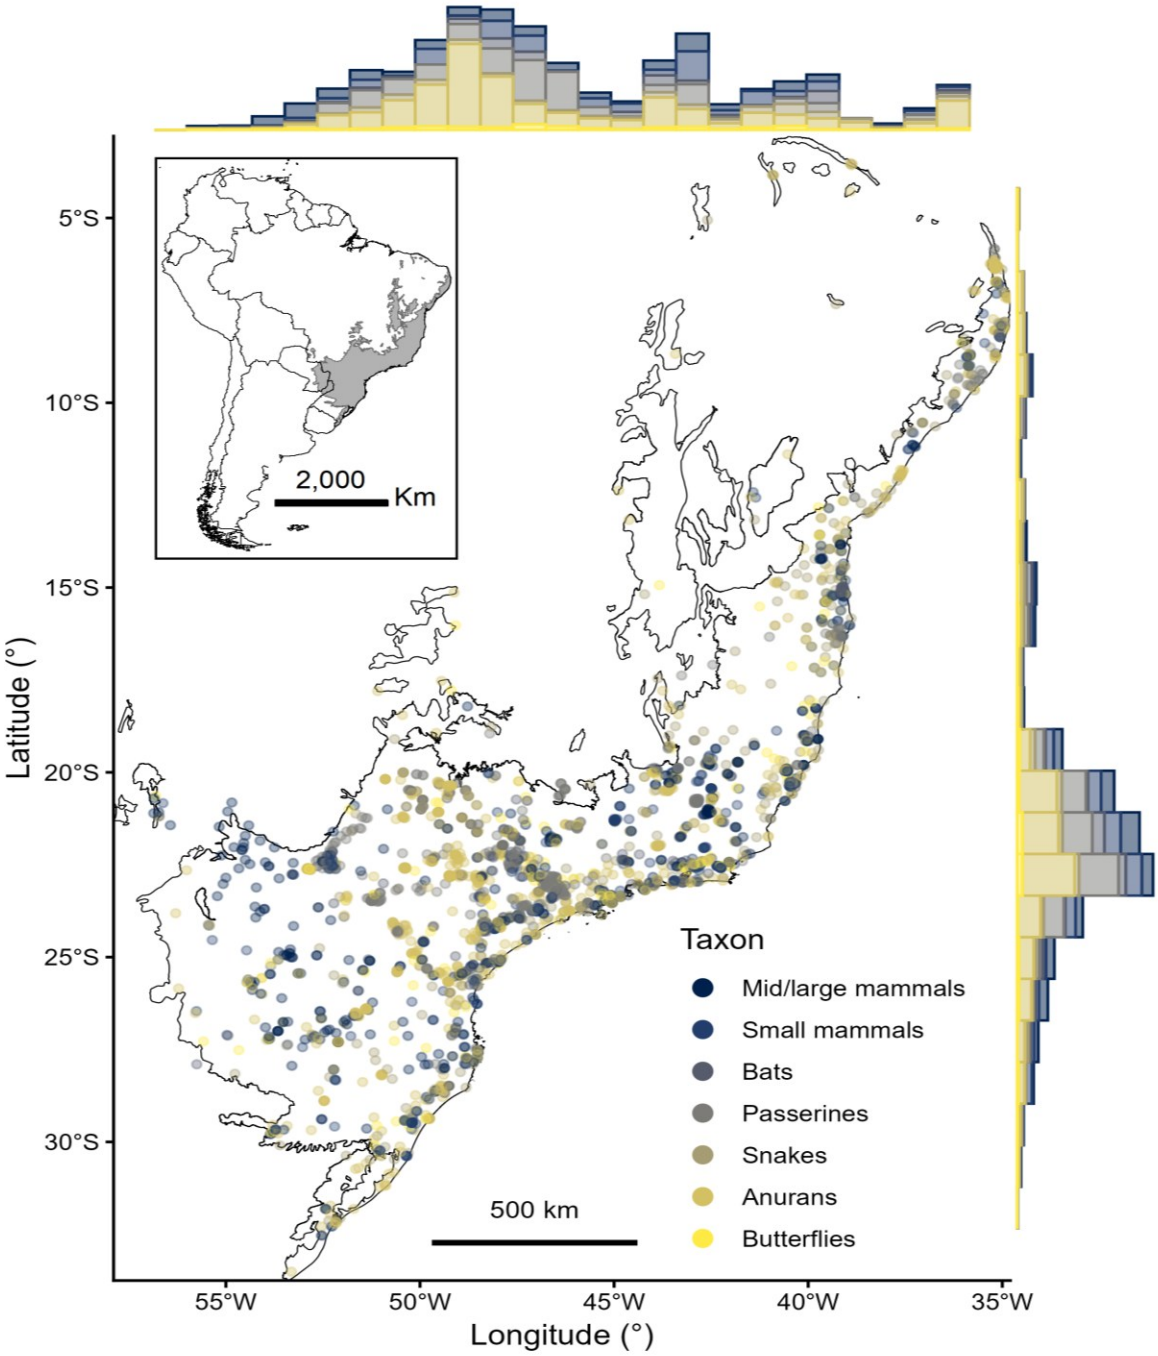


Figure 1.1. Geographic distribution of the curated species assemblages (n = 2,986) across the Atlantic Forest (AF). The marginal histograms give the sampling frequency (number of assemblages) across longitude (top) and latitude (right), stacked by animal group. The inset shows the position of the AF (shaded) within South America. These curated records form the basis of the multi-scale hexagonal grids used in the analyses; per-group sample sizes, temporal ranges, sampling methods, and the yearly sampling distribution are given in Appendix 0.S1.

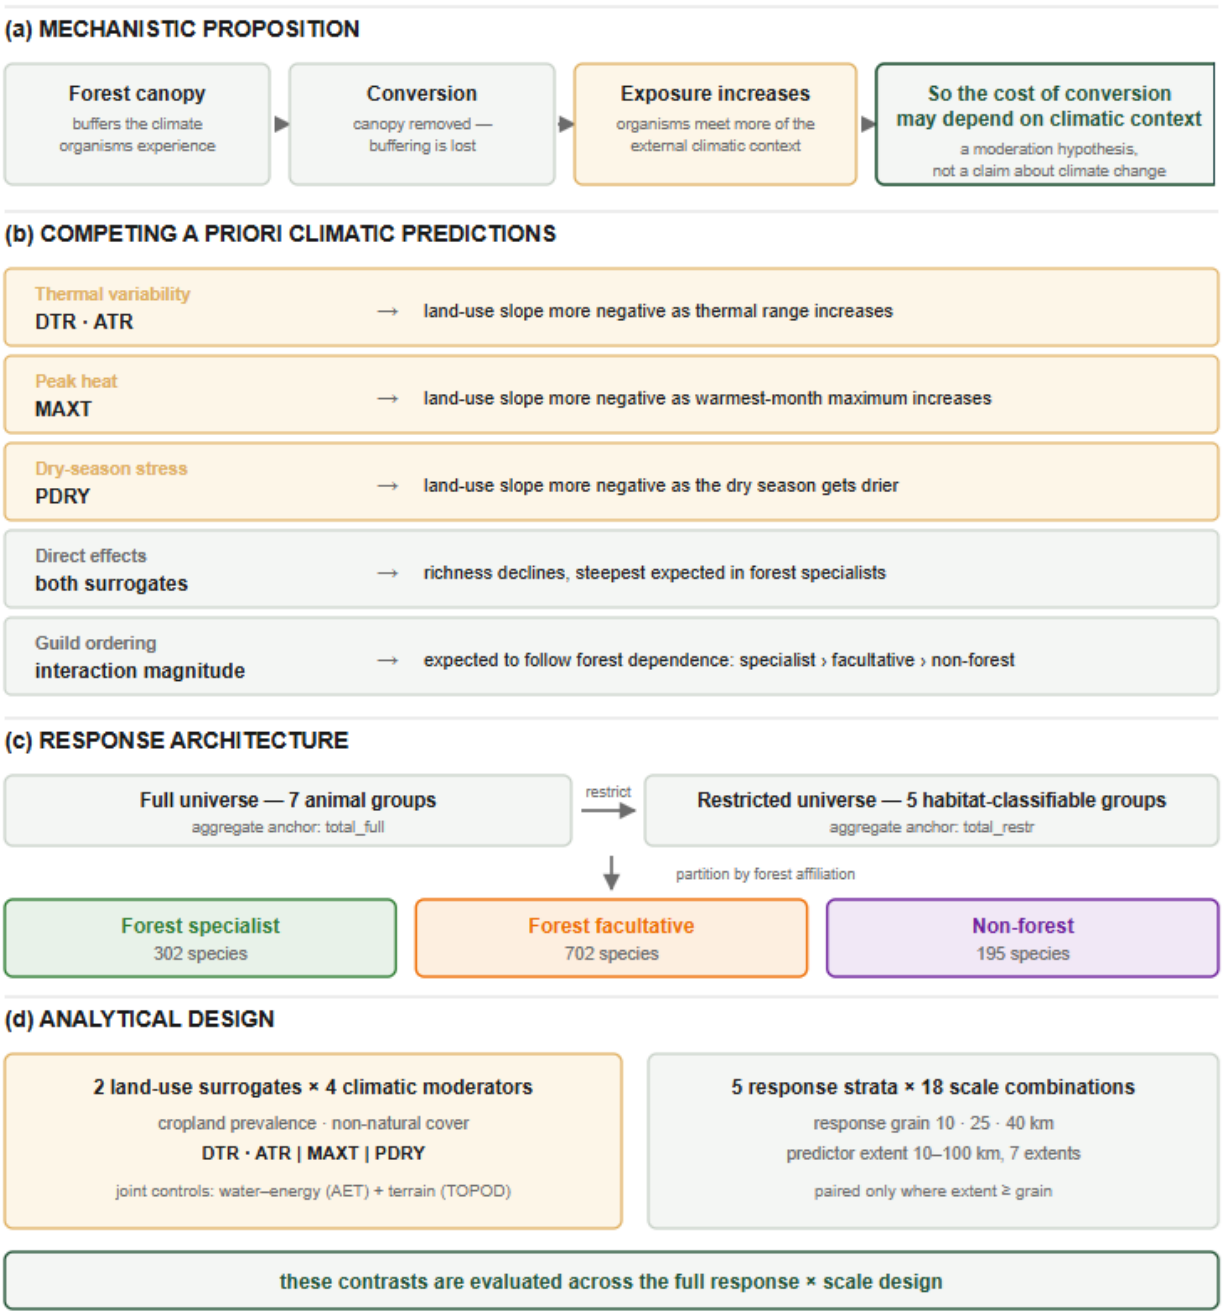


Figure 1.2. Hypothesised mechanism, competing climatic predictions, response strata and analytical design. (a) Mechanistic proposition. Intact forest canopy buffers the climate organisms experience; land-use conversion (cropland, CROP; non-natural cover, NNCOV) removes that buffering and increases exposure to external climatic conditions, so the biodiversity cost of conversion is expected to depend on climatic context. The atmospheric context itself is not altered by conversion; exposure to it is. (b) Competing a priori climatic predictions. Three exposure axes were specified before fitting and treated as competing rather than nested: thermal

variability (daily and annual temperature range, DTR and ATR), peak heat (MAXT), and dry-season stress (PDRY, inverted at assembly so that higher values denote drier dry seasons). Each predicts a more negative land-use slope at the stressed end. Two framing expectations accompany them: direct land-use effects steepest in forest specialists, and interaction magnitudes ordered by forest dependence. (c) Response strata. The full universe of seven dataset-defined groups (total_full) is restricted to five habitat-classifiable groups (total_restr), then partitioned by IUCN forest affiliation into three mutually exclusive guilds — forest-specialist (302), forest-facultative (702) and non-forest (195) species (1,199 classified). (d) Analytical design. Two land-use surrogates crossed with four climate moderators across 18 grain × extent scale combinations, with AET and TOPOD as joint controls. Support was evaluated by ΔAICc > 2 and product-term structure-coefficient concordance.

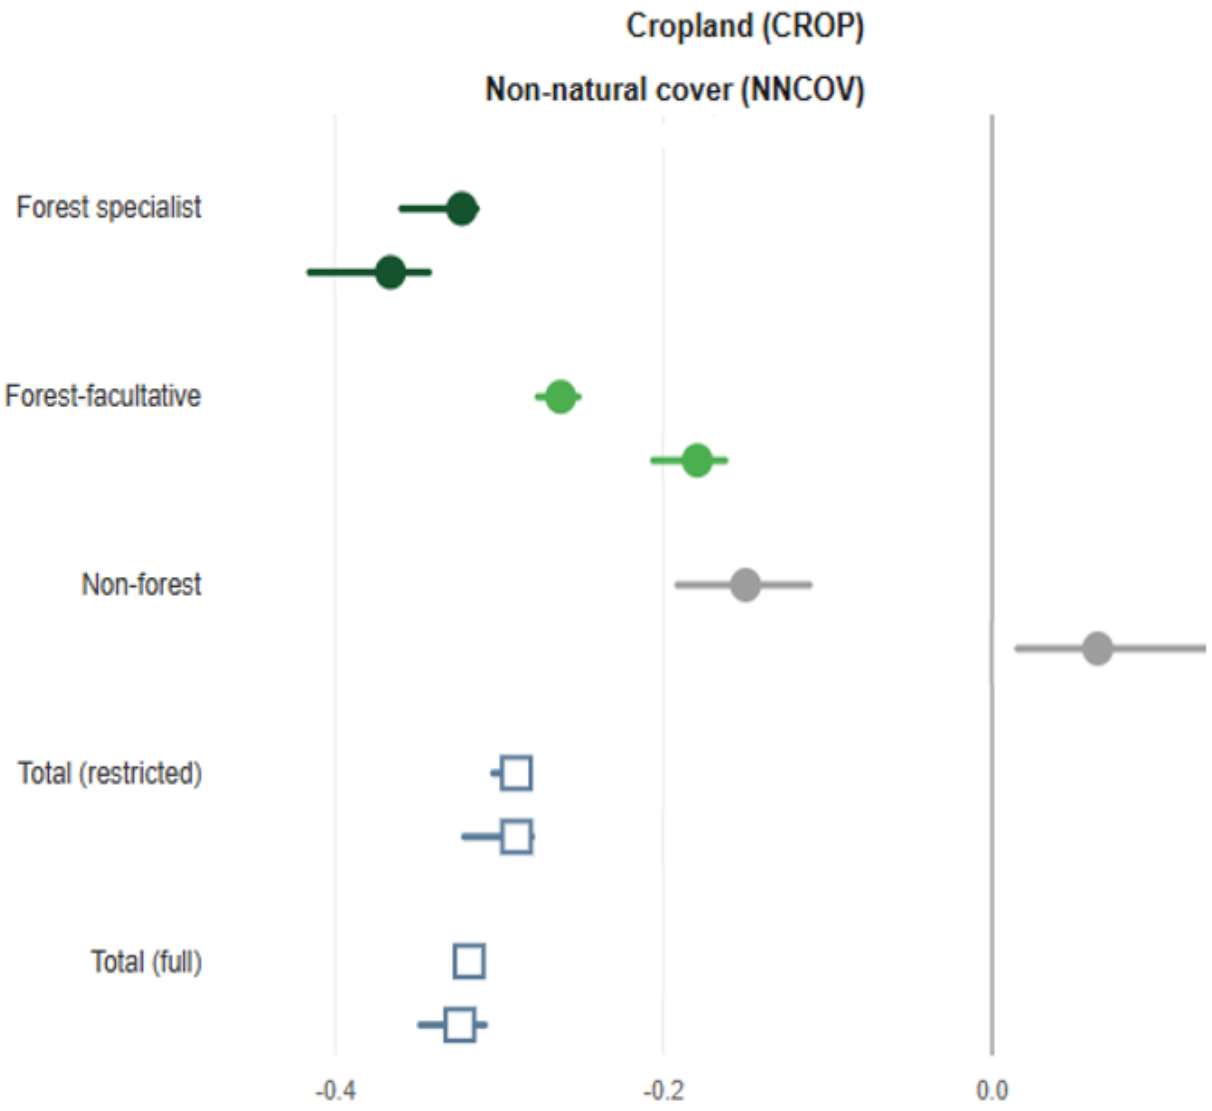


Figure 1.3. Land-use main-effect contrasts across forest-dependence strata. We show the median predicted contrast in log-richness across the shared 10th–90th percentile interval of each land-use surrogate (cropland, CROP; non-natural cover, NNCOV), estimated from adjusted main-effect models under a common random-intercept structure with the selected functional form. Filled circles are the three mutually exclusive forest-dependence guilds; open squares are the two aggregate anchors. Bars span the interquartile range of the contrast across the 18 grain × extent scale combinations (scale variability, not a confidence interval). More negative values indicate steeper richness decline across the land-use gradient. Stratum-specific random-slope fits for the forest-specialist stratum (sensitivity) are reported in Appendix 1.S4.

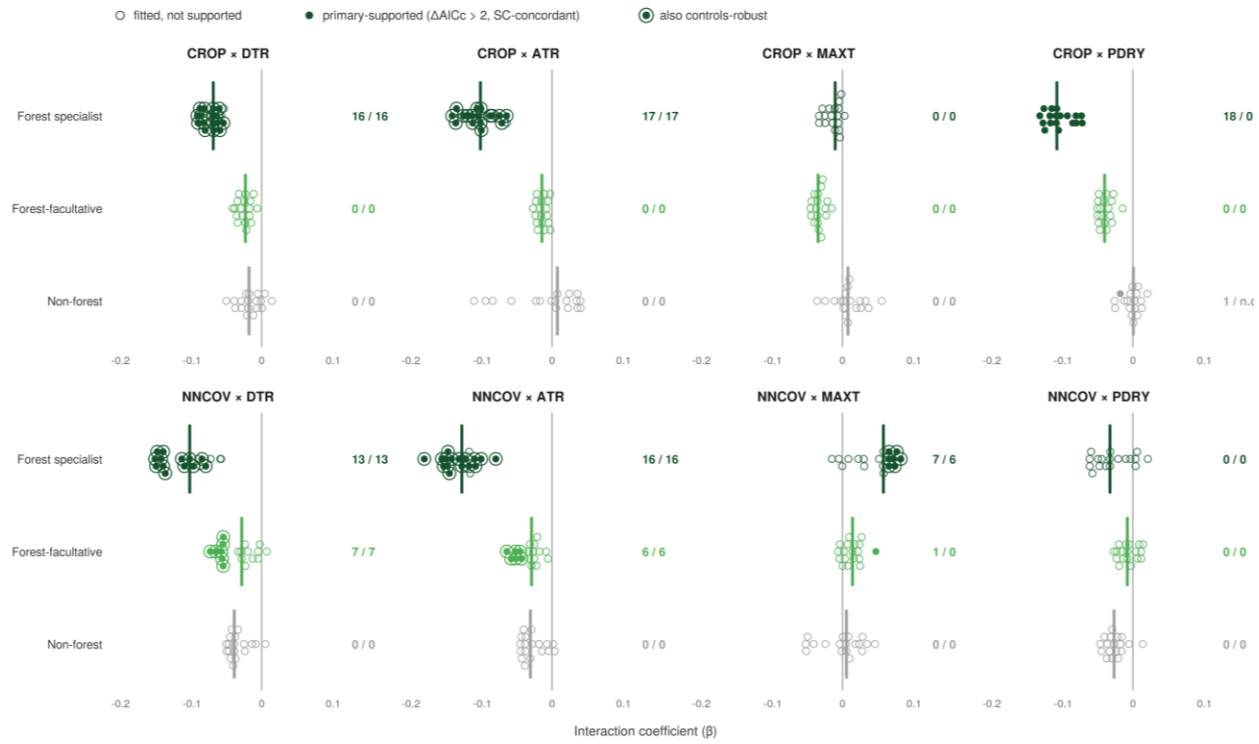


Figure 1.4. All fitted land-use × climate interaction coefficients across forest-dependence guilds. Each point represents the Tier-1 interaction coefficient for one grain × extent scale cell, with 18 cells per guild and family, across the eight combinations of cropland (CROP) or non-natural cover (NNCOV) with daily temperature range (DTR), annual temperature range (ATR), maximum temperature (MAXT), or dry-season precipitation scantness (PDRY). Open points did not meet the primary support criterion, filled points met the primary criterion of ΔAICc > 2 plus product-term structure-coefficient concordance, and filled points with an outer ring also retained ΔAICc > 2 after joint AET + TOPOD adjustment. Because CROP and NNCOV differ in definition, temporal anchor, source product, and fitted representation, their coefficient magnitudes are not directly comparable. We selected polynomial and random-slope structures independently for each response stratum and scale cell, so the plotted coefficients and median markers describe the cell-specific selected models and are not matched-specification contrasts among response strata or among scale cells. We therefore compare direction, support and control retention across strata, and we do not interpret cross-stratum or cross-cell coefficient magnitude from this figure. Aggregate anchors and full per-cell summaries are reported in Appendix 1.S5. We mark the guild median across the 18 overlapping scale cells with a vertical bar. The median summarises the fitted cells descriptively and does not estimate a pooled interaction parameter.

We report labels as primary-supported / controls-robust cells out of 18, so 18 / 0 means that all 18 cells reached primary support and none retained support after joint AET + TOPOD adjustment. We use n.d. when we could not compute the controlled comparison because the controlled fit did not converge to a positive-definite Hessian.



Figure 1.5. Scale coherence of forest-specialist thermal-variability amplification. Each panel shows one land-use × thermal-variability family across the full grain × extent design, with response grain on the vertical axis and predictor extent on the horizontal axis. Filled cells met the primary support criterion, a parsimony gain above two AICc units together with product-term structure-coefficient concordance. Open cells were fitted but did not meet it. Crossed grey cells were not fitted, because the design requires the predictor extent to be at least the response grain. Panel headers give supported cells out of the 18 valid combinations. Because response grains are nested and predictor extents overlap around many of the same assemblages, the cells are statistically dependent, so we read recurrence across the design as scale coherence rather than as independent replication, and we do not read the extent profile as identifying an ecological operating scale.

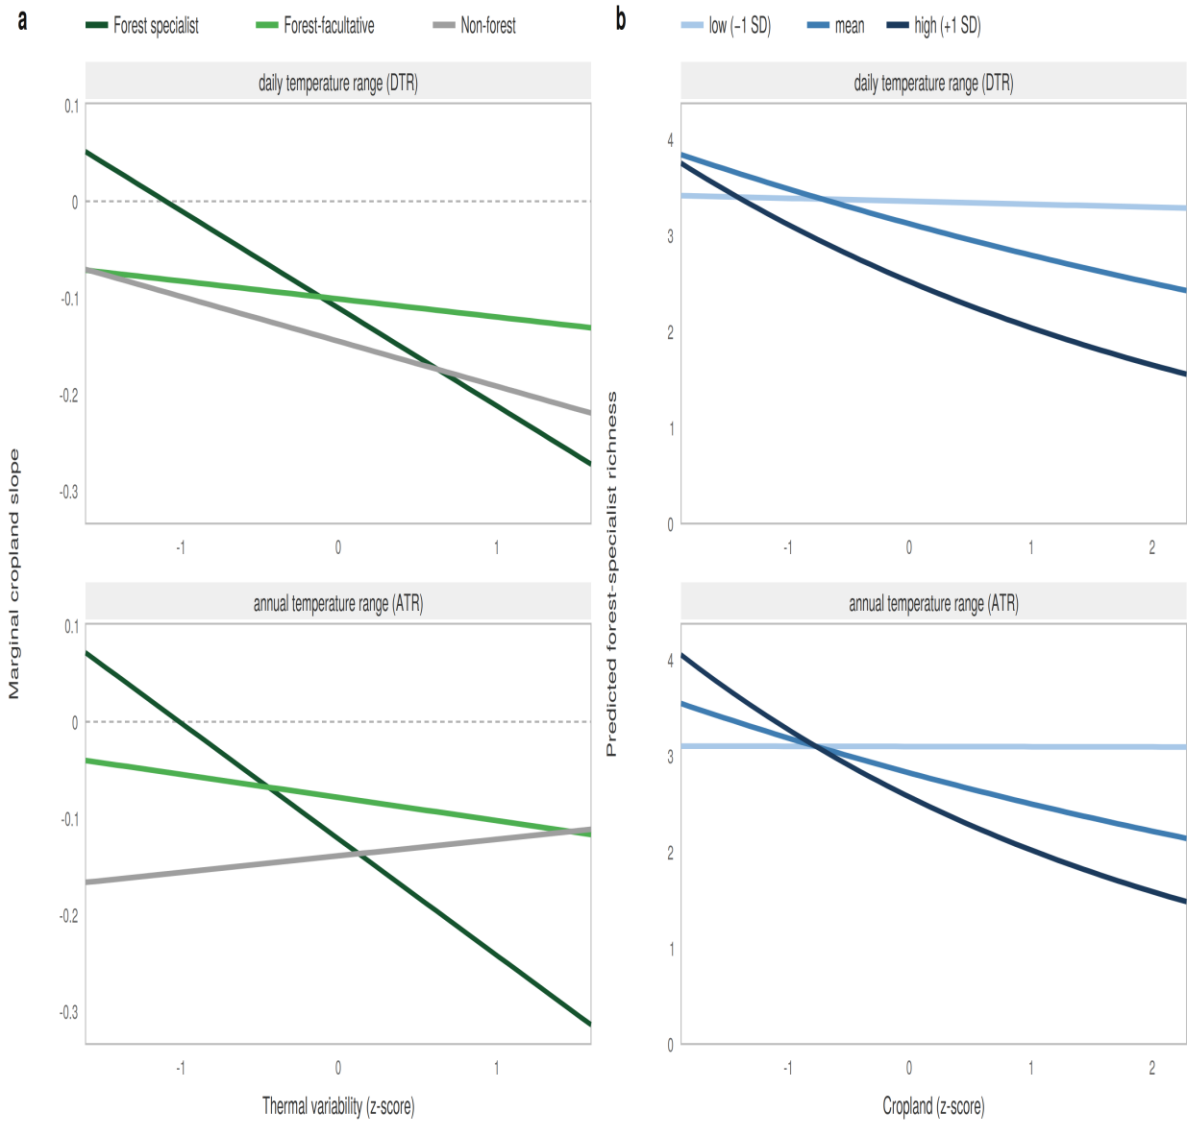


Figure 1.6. Climatic moderation of the cropland effect on the response scale, at the 25-km response grain and 25-km predictor extent, from the controlled models. (a) The marginal cropland slope across the thermal-variability gradient for each of the three mutually exclusive forest-dependence guilds. (b) Predicted forest-specialist richness across the cropland gradient at low, mean and high thermal variability. We read panel (a) for the direction and steepening of each guild's own marginal slope rather than for the size of the gap between guilds, because we selected random-slope and polynomial structures independently by stratum. Those selections coincide across the three guilds for the annual temperature range and differ for the daily temperature range. We omit intervals from panel (a) because the marginal-slope variance inherits the interaction covariance.

CHAPTER 2

Not all heterogeneity is opportunity: its biodiversity value depends on origin and metric in the transformed Atlantic Forest

Introduction

Land-use transformation is a dominant component of global environmental change and a major driver of biodiversity loss, yet its effects vary strongly among landscapes (Newbold et al. 2015; Haddad et al. 2015). This variation partly reflects the origin and ecological meaning of landscape heterogeneity. In transformed tropical forests, spatial heterogeneity can represent persistent geophysical gradients that maintain refugia and environmental opportunity, or land-use-generated vegetation texture that marks habitat subdivision, edge exposure, and matrix contrast. Distinguishing these origins is essential for predicting when landscape heterogeneity buffers biodiversity loss and when it compounds it. Environmental heterogeneity is classically expected to promote species richness by expanding niche space and enabling resource partitioning (Hutchinson 1957; MacArthur 1972; Tews et al. 2004), and positive heterogeneity–diversity relationships (HDRs) are among the most widely documented macroecological patterns (Stein et al. 2014).

Yet evidence from human-modified tropical systems is frequently inconsistent; negative or null HDRs are common precisely where heterogeneity is highest and conservation stakes are greatest (Heidrich et al. 2020; Wan et al. 2023). A central reason for this inconsistency is that heterogeneity is not one ecological quantity. Metrics can describe dispersion in continuous conditions, diversity of environmental states, spatial differentiation among sites, or configurational structure, and these operationalisations imply different mechanisms and scales of action (Stein and Kreft 2015). We therefore treat heterogeneity as a family of mechanism-linked predictors rather than as a single richness gradient. Broad-scale Neotropical evidence supports this separation: climate, topography and vegetation contribute to vertebrate richness through distinct and shared pathways, with much of the vegetation signal overlapping that of climate rather than acting independently of it (Moura et al. 2016).

In transformed forest landscapes, heterogeneity also has an origin problem. Similar values of a heterogeneity metric can arise from persistent geophysical gradients, natural habitat

diversity, secondary regeneration, agricultural mosaics, or fragmentation-driven configurational heterogeneity. These origins imply different biodiversity mechanisms — and in heavily transformed landscapes such as the Atlantic Forest, the two dominant sources, persistent geophysical variation and land-use-responsive vegetation texture, co-occur but have opposite ecological implications, so collapsing them into a single metric systematically obscures both the sign and the magnitude of HDRs (Seiferling et al. 2014; Heidrich et al. 2020).

A critical distinction in heterogeneity research is between compositional and configurational components of landscape structure (Fahrig et al. 2011). Compositional heterogeneity, the variety and proportions of cover types present, expands the environmental envelope available for niche diversification. Configurational heterogeneity, the spatial arrangement of those cover types, modulates the effective habitat area per species and the degree of connectivity among suitable patches. In intact landscapes, both components can promote diversity through complementary pathways. In fragmented landscapes, however, configurational heterogeneity increases as continuous habitat is disaggregated into smaller, more isolated patches, reducing the effective area per specialist species and elevating extinction risk (Allouche et al. 2012; Laanisto et al. 2013). Whether we interpret a vegetation structural attribute as heterogeneity or fragmentation therefore depends on the organismal group, spatial scale, and degree of land-use transformation (Tews et al. 2004; Fahrig et al. 2011). When we interpret the attribute as heterogeneity, it represents expanded niche opportunities. When we interpret it as fragmentation, it represents subdivision of usable conditions.

The area–heterogeneity tradeoff (AHTO) provides one candidate mechanism by which land-use-responsive vegetation texture may become negative rather than beneficial. As heterogeneity increases within a finite area, the effective habitat size per species declines, elevating extinction risks for species with narrow niches and limited dispersal capacity (Kadmon and Allouche 2007; Allouche et al. 2012). Microfragmentation extends this logic to the community scale, predicting that fine-grained configurational heterogeneity can reduce population viability by isolating suitable microhabitats even when total habitat composition is held constant (Laanisto et al. 2013). Negative HDRs consistent with AHTO dynamics have been documented for approximately one quarter of taxa-specific relationships in multi-taxa, multi-metric studies (Heidrich et al. 2020), and are more prevalent at smaller spatial grains where heterogeneity can reduce effective patch area rather than expand niche space (Tamme et

al. 2010). AHTO is therefore a candidate mechanistic explanation for negative vegetational interactions, not the chapter's primary question. The primary question is whether, and in which direction, environmental heterogeneity moderates the biodiversity cost of land use in a transformed tropical system — a landscape moderation question motivated by Tscharntke et al.'s (2012) conditional heterogeneity hypotheses, which treat landscape structure as a family of context-dependent pathways that can buffer, amplify, or fail to modify land-use effects depending on origin, scale, and land-use context. We use AHTO curvature as a mechanistic diagnostic for the vegetational interaction sign, not as the primary inferential target.

Historically spanning more than 150 million hectares of tropical and subtropical rainforests along the eastern coast of Brazil, extending into northeastern Argentina and eastern Paraguay (Muylaert et al. 2018), the AF has lost approximately 84% of its original cover to deforestation. The remaining forest is severely fragmented, with ~80% of remnants smaller than 50 ha, and ~80% of remaining forest lying within 1 km of a forest–matrix edge, and only ~9% within formally protected areas (Ribeiro et al. 2009; Vancine et al. 2024). Forest conversion has reorganised vegetation into mosaics of forest remnants, agriculture, pasture, plantations, and degraded matrix, generating vegetational heterogeneity as a landscape-scale signature of structural disruption rather than ecological opportunity (Laurance et al. 2012; Haddad et al. 2015). At the same time, the AF retains substantial geophysical heterogeneity — topographic and climatic gradients associated with the Serra do Mar escarpment, the Mantiqueira range, and broad latitudinal and elevational gradients — that structures environmental filtering and species turnover independently of land use. This contrast between anthropogenic and geophysical heterogeneity within the same biome, combined with large-scale multi-taxa assemblage data spanning the full biome, allows origin- and scale-contingent HDRs to be disentangled

We also tested whether land-use and heterogeneity moderation scale with forest dependence. Species differ in their reliance on continuous forest cover, canopy buffering and forest-mediated refugial conditions, so forest affiliation should govern both the cost of conversion and the value of any buffering environment (Fahrig et al. 2019; Fletcher et al. 2018). We therefore analysed forest-specialist, forest-facultative and non-forest richness as mutually exclusive components of the restricted assemblage, alongside full and restricted aggregate richness. Forest specialists should be most exposed to land-use-responsive vegetation texture and most likely to show buffering where persistent physical refugia preserve forest-mediated

persistence pathways, facultative species should respond more weakly or context-dependently, and non-forest species should respond weakly, not at all, or in the opposite direction where conversion creates open or matrix-associated habitat (Haddad et al. 2015; Chase et al. 2020). The full aggregate retains all sampled taxonomic groups, whereas the restricted aggregate uses the same taxonomic universe as the guild partition. Comparing them separates forest-affiliation structure from artefacts introduced by excluding groups without compatible guild assignments. Together these strata test whether origin-dependent moderation reflects a forest-mediated signal rather than a generic landscape gradient.

Hypotheses and predictions

This chapter asks whether the origin of a heterogeneity axis, rather than its identity as a predictor, determines how it modifies the biodiversity cost of land-use transformation. We organised the analysis around three moderation pathways. Vegetational heterogeneity tests whether land-use-responsive vegetation texture compounds habitat loss. Geophysical heterogeneity tests whether persistent terrain and climatic gradients buffer land-use effects. Productivity tests whether realised water–energy availability dampens land-use costs through a demographic pathway. We tested five a priori predictions. Together these predictions follow an origin-dependent logic. If heterogeneity effects depend on origin, land-use-responsive vegetation texture should show negative richness associations or compounding land-use interactions, whereas persistent geophysical heterogeneity should show positive richness associations or buffering interactions under the same land-use and scale framework, and these contrasts should be clearest in the forest-dependent components of richness. First, land use should show consistently negative associations with richness as the baseline biodiversity cost of transformation.

Second, geophysical heterogeneity (terrain ruggedness, TOPOD; climatic niche diversity, CLIMD) should show predominantly positive associations, because persistent abiotic gradients expand environmental space and refugial opportunity without implying recent native-cover loss (Hjort et al. 2015; Tukiainen et al. 2017; Kerr and Packer 1997). Third, vegetational heterogeneity should show negative or conditional associations in this transformed system, because texture metrics based on the Enhanced Vegetation Index (EVI) respond directly to land-cover replacement and configurational disruption (Seiferling et al. 2014). Fourth, land-use ×

vegetational heterogeneity interactions would be predominantly negative, indicating that vegetation texture compounds rather than buffers land-use damage (Kadmon and Allouche 2007; Allouche et al. 2012; Laanisto et al. 2013). Fifth, productivity and water–energy availability defined a third moderation pathway in which higher resource availability sustains larger populations and increases demographic resilience to land-use-driven habitat loss, so land-use effects on richness should be weaker in more productive landscapes — a resource-mediated buffering pathway analogous in logic to more-individuals prediction (Currie 1991; Hawkins et al. 2003; Storch et al. 2018).

We expected inference here to be geometry-limited because water–energy variables and land-use transformation are non-orthogonal in the Atlantic Forest. We therefore treated actual evapotranspiration (AET), the realised water–energy flux, as the sole downstream energy proxy and used potential evapotranspiration (PET) and the Priestley–Taylor ratio (ALPHA) only as proxy-swap robustness checks on confirmed surfaces, not as parallel moderation arms (Currie 1991; Hawkins et al. 2003). We treat vegetation heterogeneity here as an ecological signal whose origin must be interpreted, not as a direct estimate of fragmentation per se.

Our remotely sensed vegetation texture metrics capture compositional variability and configurational irregularity in greenness, but cannot separate habitat configuration from habitat amount. If fine-grained vegetational heterogeneity shows a negative association with richness distinct from the positive effects of geophysical heterogeneity, this pattern is consistent with, rather than demonstrative of, fragmentation-driven configurational heterogeneity operating through area–heterogeneity tradeoff dynamics. Our goal here is to determine whether HDRs in the Atlantic Forest are consistent with origin-dependent and scale-contingent theoretical expectations, and to identify which forms of heterogeneity buffer or compound land-use effects. We grouped predictors into mechanism-specific families with distinct ecological interpretations and translated this framework into three moderation pathways with their expected directions (Table 2.1).

Methods

Study area

The AF of South America spans tropical and subtropical rainforests along the eastern coast of Brazil, extending into northeastern Argentina and eastern Paraguay (Muylaert et

al. 2018). Now retaining approximately 16% of its original extent following centuries of agricultural expansion and urbanisation (Ribeiro et al. 2009; Vancine et al. 2024), the AF is severely fragmented, with ~80% of patches smaller than 50 ha and ~80% of forest lying within 1 km of a forest–matrix edge.

Species assemblage data

We used 2,986 georeferenced, effort-reported assemblages spanning seven animal groups across the AF between 1981 and 2018 (Figure 2.1; Appendix 0.S1). Biodiversity records were drawn primarily from the ATLANTIC series of open-access data papers published in Ecology, covering mid-large and small mammals (Mammalia; Lima et al. 2017; Souza et al. 2019; Bovendorp et al. 2017; Figueiredo et al. 2017), bats (Chiroptera; Muylaert et al. 2017), passerines (Aves: Passeriformes; Hasui et al. 2018), anurans (Anura; Vancine et al. 2018) and fruit-feeding butterflies (Lepidoptera: Nymphalidae; Santos et al. 2018), and were supplemented for snakes (Serpentes) by an Atlantic Forest snake database (Moura et al. 2017a, 2017b). We report the sampling methods and effort units for each animal group, together with group-level reference periods and gamma richness, in Table 0.S1.1. We retained only uniquely georeferenced assemblage-level records derived from standard quantitative protocols with reported effort, and restricted the working dataset to assemblages sampled between 1981 and 2018. We excluded opportunistic and presence-only records, and defined a secondary hierarchy (METHOD, n = 16) nested within taxa to account for heterogeneity in sampling protocols, effort units, and survey scope.

We classified each species by forest affiliation using species-level habitat records from the IUCN Red List, coded under the IUCN Habitats Classification Scheme (IUCN 2012) and following Donald et al. (2019); the resulting universes and their exact partition are set out in Appendix 0.S1.3. We treated species whose listed habitats fell wholly within the Forest class as forest specialists, species listing Forest together with one or more non-forest classes as forest-facultative, and species with no Forest habitat listed as non-forest. These classes describe exclusivity among listed habitat types and do not imply demonstrated habitat obligacy.

We built five response strata across two taxonomic universes. The full universe covered all seven sampled groups. The restricted universe covered the five groups whose species could be matched to habitat records, namely mid-large mammals, small mammals, bats, passerines and

anurans. We excluded snakes from the restricted universe because they entered the datasets as analysis-derived richness estimates rather than species-level records, and we excluded fruit-feeding butterflies because they were not comprehensively assessed. Within the restricted universe we modelled total richness together with the three mutually exclusive forest-affiliation guilds, which partition it exactly. We fitted all five strata through the modelling workflow. The strata differ markedly in mean richness, and coefficient precision under a Poisson likelihood scales with the mean count, so the guild strata are estimated less precisely than the aggregates. Non-forest richness carries the lowest counts and the widest intervals, and we therefore read weak results in that stratum as consistent with limited precision as well as with absence of effect (Appendix 2.S1.1).

We aggregated assemblages by METHOD into systematic hexagonal grids at three response grains: 10 km (local; ~314 km²), 25 km (intermediate; ~1,963 km²), and 40 km (broad; ~5,027 km²). We measured predictors in buffers at seven extents — 10, 25, 40, 55, 70, 85 and 100 km — and paired each response grain only with extents greater than or equal to that grain, giving seven, six and five combinations at the 10, 25 and 40 km grains and eighteen valid grain × extent combinations in total (Figure 2.2). Each buffer is centred on the centroid of the assemblages aggregated into that record, not on the geometric centre of the hexagonal unit, which serves only as an aggregation template.

We extracted and summarised predictor values within these buffers using the *sf*, *terra* and *exactextractr* R packages (Pebesma 2018; Hijmans 2022; Baston 2023). Species richness was recalculated per cell and METHOD at each grain. Site-based effort (EFFsb; assemblages per cell) was the primary effort covariate; unit-based effort (EFFub; summed original effort units) was used in sensitivity analyses (Appendix 2.S4.2). Within the 10 km grain we retained a temporally disaggregated variant (WITHyear, n = 1,939) alongside the temporally aggregated primary dataset (NOyear, n = 1,689) for temporal sensitivity assessment.

Environmental predictors

Land-use surrogates. We used two complementary land-use metrics as surrogates of land-use transformation: cropland prevalence (CROP) derived from the global 1-km consensus land-cover product of Tuanmu and Jetz (2014), using the cultivated and managed vegetation class aggregated across source maps, and non-natural land cover (NNCOV) from a high-

resolution regional product for the AF (Vancine et al. 2024Higher values of both metrics indicate greater transformation of native cover. CROP and NNCOV capture complementary dimensions: CROP isolates agriculture-specific effects; NNCOV integrates total anthropogenic transformation including urban, pasture, and infrastructure.

Vegetation heterogeneity. We grouped heterogeneity predictors into abiotic and vegetational families because they differ in origin and expected ecological interpretation. Abiotic predictors described relatively persistent geophysical and climatic gradients, whereas EVI-based predictors described contemporary variation in vegetation greenness and texture that can respond directly to land-cover replacement, forest degradation, secondary regrowth, agriculture, pasture, and plantation mosaics. This distinction follows the geodiversity literature, in which relatively persistent abiotic Earth-surface heterogeneity (topographic, geomorphic and climatic gradients) is contrasted with land-use-responsive vegetation structure as a more proximate ecological expression of land-cover state (Hjort et al. 2015; Tukiainen et al. 2017; Räsänen et al. 2016). We did not assume abiotic predictors operate independently of land-use history, because topographic and geomorphic features can direct where human activities concentrate and may therefore share explained variation with landscape composition in transformed systems (Räsänen et al. 2016). We captured vegetational heterogeneity using textural features of EVI derived from MODIS imagery at 1-km resolution for 2000–2005 (Tuanmu and Jetz 2015).

We used two complementary EVI-based metrics: the coefficient of variation of EVI (EVICV), a first-order texture measure summarising compositional heterogeneity as the normalised dispersion of EVI; and the entropy of EVI (EVIENT), a second-order texture quantity incorporating the co-occurrence and spatial arrangement of pixel values, so that higher values indicate greater disorder in the arrangement of greenness. Both indices are unitless and yield higher values where replacement and interdigitation of native vegetation with cultivated or built-up areas generate sharp structural contrasts and irregular spatial patterns (Tuanmu and Jetz 2015). EVICV captures compositional variability; EVIENT captures spatial arrangement. Spectral-texture metrics can be confounded with mean environmental state, productivity gradients, and non-random land-cover distributions along land-use gradients. We therefore evaluated their associations with richness using land-use-adjusted moderation models and structure-coefficient diagnostics, while treating realised water–energy availability (AET) as the downstream focal moderator and its component proxies as robustness checks rather than parallel

arms (Hortal et al. 2013; Carnicer et al. 2013). Their correlation is moderate and stable across the scale matrix (median r = 0.57, range 0.55–0.62), and their sampling distributions differ in shape, making them complementary rather than redundant dimensions of landscape pattern (Turner 1989; Tscharntke et al. 2012). The two also differ in their exposure to suppression, EVIENT being more strongly coupled to the land-use surrogates than EVICV, so a null result does not carry the same weight for both (Appendix 2.S2.1 for predictor definitions and provenance; Appendix 2.S2.3).

Topographic heterogeneity. We quantified topographic heterogeneity using the terrain ruggedness index (TRI) derived from the Global Multi-resolution Terrain Elevation 2010 dataset at 1-km resolution (Amatulli et al. 2018). Terrain ruggedness (TOPOD) measures the magnitude of elevational differences within an area (units in metres), capturing variation from microtopographic structures to large-scale relief. We also screened geomorphological landform diversity (GEOD), the Shannon diversity within each buffer of the ten geomorphon landform classes distributed with the same dataset (Amatulli et al. 2018), computed with the R package *landscapemetrics* (Hesselbarth et al. 2019). We did not carry GEOD into the final fitted interaction factorial, because TOPOD and CLIMD captured the retained abiotic axes with clearer interpretation.

Climatic heterogeneity. We quantified climatic heterogeneity using the Shannon diversity index ($\log_{10}$-based) of bioclimatic classes (CLIMD; unitless) derived from a multivariate clustering and ordination of temperature- and precipitation-based variables at 1-km resolution (Metzger et al. 2013). CLIMD quantifies the diversity of climatic environments by integrating variability in both thermal and hydric conditions within spatial units, providing a biologically meaningful representation of geographic variability in climate space (Coelho et al. 2023; Stein and Kreft 2015). We calculated CLIMD using the R package *landscapemetrics* (Hesselbarth et al. 2019). TOPOD and CLIMD are moderately associated (mean r = 0.67 averaged across the scale matrix; Appendix 2.S2.3 reports the full cross-scale profile, whose representative-scale value differs from any single grain–extent cell) and the association is directional, because relief generates climatic heterogeneity rather than the reverse. We therefore fit them in separate models and read concordant support as one abiotic gradient recovered by two operationalisations rather than as independent confirmation (Appendix 2.S2.3).

Productivity and water–energy predictors. We used actual evapotranspiration (AET) as the sole downstream water–energy proxy in the fitted interaction factorial. AET represents realised water–energy flux and was derived, before z-transformation, as the product of annual potential evapotranspiration from the ENVIREM dataset (PET; Title and Bemmels 2018) and the Priestley–Taylor coefficient of the Global High-Resolution Soil-Water Balance dataset (ALPHA; Trabucco and Zomer 2019), both at 1-km resolution. PET indexes atmospheric evaporative demand rather than realised biological productivity in the humid Atlantic Forest, and ALPHA represents water-balance efficiency; we therefore retained PET and ALPHA as screen-level and proxy-swap robustness variables, not as parallel land-use moderation arms. In the interaction workflow AET entered in three roles: as a focal water–energy predictor whose own richness association we estimated, as an additive control in land-use × heterogeneity robustness tests, and as the energy axis in three-way modulation models. PET and ALPHA entered only as proxy-swap checks on confirmed surfaces, testing whether energy-conditioning conclusions depended on operationalisation without fitting collinear energy proxies together. We computed structure coefficients alongside partial slopes because AET can share variance with land-use intensity in transformed landscapes, creating the same kind of suppression geometry that motivated the structure-coefficient diagnostic for vegetation-texture predictors.

Statistical modelling

We structured the analysis as a set of a priori nested contrasts rather than as an exploratory factorial screen. Land-use-responsive vegetation texture was treated as the most downstream heterogeneity signal and was therefore evaluated against both persistent terrain structure and water–energy availability, whereas geophysical heterogeneity was evaluated against water–energy availability alone, because terrain, climatic heterogeneity and productivity can share positive richness gradients. AET carried three roles in this design: a focal water–energy predictor whose own richness association we estimated, a reported land-use × energy moderation arm, and a control used to test whether land-use × heterogeneity signals survived measured water–energy structure. The land-use surrogate × scale × stratum factorial served as the robustness envelope for these contrasts, not as an independent hypothesis space. We therefore interpret persistence under conditioning as evidence that a moderation signal is not reducible to measured water–energy structure, and we do not use joint-control models to partition variance or to rank terrain against productivity.

We followed a three-tier modelling sequence: Tier 1 testing whether each LU × heterogeneity interaction is supported; Tier 2 testing robustness after AET conditioning; and Tier 3 testing whether AET modulates a supported moderation (full tier architecture in Appendix 2.S3.2). Within each tier, all model comparisons share the same spatial autocovariate, and cross-tier AICc comparisons are not made. Appendix 2.S4 summarises this sequence schematically (Figure 2.S4.1). All analyses were carried out in R (R Core Team 2024). We fitted conditional Poisson generalised linear mixed models (GLMMs) with an observation-level random effect (OLRE) to untransformed species richness counts using the *glmmTMB* package (Brooks et al. 2017). The Poisson-OLRE specification captures overdispersion through individual-level variance without requiring a negative-binomial parameterisation, which improves convergence stability under complex random-effects structures (Harrison 2014; O'Hara and Kotze 2010).

All focal predictors were z-transformed prior to fitting to place partial coefficients on a common scale and to facilitate comparison across predictor families (Schielzeth 2010). Residual adequacy, overdispersion, and spatial autocorrelation were assessed using scaled quantile residuals computed with the *DHARMa* package (Hartig 2022) and Moran's I computed with the *ape* package (Paradis and Schliep 2019); full diagnostics are reported in Appendix 2.S4.5. We used model complexity only where it corresponded to a specific inferential task: random slopes for taxon-level heterogeneity in land-use responses, polynomial terms to protect interaction tests from unmodelled baseline curvature, and interaction terms to test explicit moderation hypotheses.

Each model included one land-use (LU) surrogate and one focal heterogeneity or productivity predictor, plus sampling effort (EFFsb) and a spatial autocovariate as fixed effects, with random intercepts for TAXA and METHOD and an OLRE. The largest candidate interaction model can therefore be represented as:

$$\text{resp} \sim \text{EFFsb} + \text{LU} + \text{M} + \text{LU:M} + \text{space}$$

$$+ (1 \mid \text{METHOD}) + (1 + \text{LU} + \text{M} \mid \text{TAXA}) + (1 \mid \text{obs})$$

where EFFsb is site-based sampling effort, LU is one of the two land-use surrogates, M is the focal heterogeneity or productivity predictor, space is the residual spatial autocovariate, METHOD carries sampling-method intercepts, TAXA carries the candidate taxon-level random slopes, and the observation-level intercept accounts for overdispersion. The TAXA term shows

the fullest candidate random-slope structure; the fitted structure was selected from the four candidates described below and therefore varied among model families and scale cells. Appendix 2.S4 gives the complete specification and model-building diagnostics.

We fitted LU surrogates (CROP, NNCOV) and focal predictors (EVICV, EVIENT, TOPOD, CLIMD, and AET) as separate model families rather than in a single joint model, because each LU–predictor pair constitutes a distinct moderation hypothesis (Table 2.1), including all predictors simultaneously would render partial coefficients uninterpretable under the collinearity structure documented in Appendix 2.S2.4, and the scale factorial (three response grains, seven predictor extents) prohibits a single pooled analysis. GEOD, PET, and ALPHA were screened or used in proxy-swap diagnostics but were not carried as focal predictors in the final interaction factorial.

Spatial autocorrelation was addressed via a residual spatial autocovariate derived from a bidimensional thin-plate spline GAM fitted to Pearson residuals of the non-spatial model with the mgcv R package (Crase et al. 2012; Wood 2004). We derived the spatial term from residuals rather than fitting a coordinate smooth alongside the environmental predictors. Heterogeneity, energy and land-use gradients are themselves geographically structured, so a jointly estimated smooth competes with them for shared broad-scale variation and leaves the focal coefficients conditional on a flexible spatial surface. Deriving the term after the environmental model has been fitted confines the adjustment to dependence those predictors leave unexplained, though it identifies neither the source of that dependence nor any bias from unmeasured predictors. For all nested model comparisons, the autocovariate was held fixed from the most complex non-spatial model in the comparison set, ensuring that model comparisons compare structures differing only in the focal term (Fletcher and Fortin 2018).

Predictor geometry audit

Before fitting response models, we characterised the pairwise collinearity structure among all focal predictors, LU surrogates, and productivity covariates across all scale combinations used in the analysis (Appendix 2.S2.2). This step is required because the conventional screening criteria of pairwise Pearson $r < 0.7$ and variance inflation factors below 5 (Dormann et al. 2013) do not detect gradient-sharing problems in multi-predictor systems, where

two predictors can be moderately correlated pairwise but share a dominant latent gradient that systematically distorts additive partial coefficients.

Before fitting the focal interaction models, the geometry audit characterised three predictor families with distinct collinearity profiles, measured before fitting and reported here because they determined the diagnostic sequence (Appendix 2.S2.3): vegetational heterogeneity metrics coupled to land use in a metric-specific way, with configurational diversity (EVIENT) more strongly coupled to land use than greenness-magnitude variability (EVICV); geophysical heterogeneity metrics coupled to the land-use surrogates negatively and moderately rather than negligibly, consistent with conversion concentrating away from rugged and climatically diverse ground; and productivity proxies coupled to the land-use surrogates in opposing directions. Median values across the eighteen grain × extent combinations were NNCOV–EVIENT +0.67, CROP–EVIENT +0.62, NNCOV–EVICV +0.50 and CROP–EVICV +0.26; TOPOD–CROP −0.54, TOPOD–NNCOV −0.41, CLIMD–NNCOV −0.31 and CLIMD–CROP −0.26; and ALPHA–CROP −0.56, PET–CROP +0.54 and AET–CROP −0.33, with each of these signs holding in every combination. These coupling profiles guided the diagnostic inference sequence described below (Appendix 2.S2.3).

Additive inference and coefficient reliability

Model-component selection and hypothesis testing both use AICc, but for different purposes, and the threshold is applied differently in each. For component selection — random-slope structure and polynomial inclusion — models within $\Delta\text{AICc} \leq 2$ of the lowest-AICc reference were treated as statistically indistinguishable and the simplest member of that band was retained. For hypothesis testing, by contrast, an interaction was adopted only when it improved fit by more than two AICc units over its matched additive baseline. Keeping these uses separate prevents component selection from being conflated with evidence for the ecological interaction itself. Both applications reduce reliance on individual p-values across the scale factorial (Burnham and Anderson 2002; Burnham et al. 2011). The selection steps described below apply the appropriate form of this rule.

For each LU × predictor × scale combination, we selected the random-slope (RS) structure by comparing four nested specifications: an intercept-only baseline (v1), a model with random slopes for the land-use surrogate (v2), a model with random slopes for the heterogeneity

moderator (v3), and a model with correlated random slopes for both predictors (v4), selecting the simplest model within $\Delta AICc \leq 2$ of the lowest-AICc variant (Barr et al. 2013; Appendix 2.S4.4).

Prior to fitting interaction models, we tested quadratic terms for each LU surrogate and each focal predictor sequentially using the same AICc rule, fixing the polynomial structure before adding any cross-product term. This guards against attributing nonlinear marginal responses to interaction effects (Bar-Massada and Wood 2014; Appendix 2.S4.3).

We computed structure coefficients alongside partial slopes, where $SC = cor(\hat{\eta}, x^*)$, $\hat{\eta}$ is the fitted linear predictor and x^* the standardised predictor, for each predictor in every additive model, following the use of structure coefficients to diagnose predictor contribution and suppression under collinearity, an approach developed in landscape ecology for separating collinear habitat-loss and fragmentation effects (Smith et al. 2009; Ray-Mukherjee et al. 2014). This step was biologically necessary: because land use and vegetation texture share a transformation gradient in the AF, standard additive models can assign the negative heterogeneity–richness signal to the land-use coefficient, leaving the vegetation texture coefficient with the wrong sign — a misclassification with direct conservation implications. The SC diagnoses this redistribution by quantifying the total model-implied association between each predictor and fitted richness regardless of the sign of its partial coefficient. Predictors where SC and partial slope disagreed in sign were classified as suppressed and interpreted from univariate marginal slopes as the external reference for ecological direction. For predictors where the fitted $\hat{\eta}$ was dominated by a predictor strongly collinear with the diagnosed predictor — a situation that can itself contaminate the SC — we additionally extracted univariate slopes (models with the focal predictor only, no LU covariate) as an external reference for the marginal association. Ecological interpretation was based on the SC-plus-univariate classification (Appendix 2.S2.4), not the raw partial coefficient. This diagnostic sequence was motivated by the coupling profiles documented in the predictor geometry audit (Appendix 2.S2.3).

Interaction inference

We fitted LU × HET cross-product terms using the polynomial-adjusted additive baseline as the comparison model, with support assessed by the same AICc rule (interaction term adopted only when $\Delta AICc > 2$ in its favour over the additive-only baseline). The RS structure from the

additive layer was retained for interaction models. Where the winner RS structure included a random slope for the HET predictor (v3_het_only), interaction model fitting was attempted with that structure first; when convergence was not achieved, a fallback to the intercept-only random structure (v1_none) was applied and documented (Appendix 2.S4.4). This fallback is an estimability protocol, not a model selection decision: random HET slopes and fixed LU–HET product terms are partially confounded under v3_het_only, and v1_none is the correct specification for interaction inference in those cases.

To assess the reliability of supported interaction terms, we computed SC for each standardised product term (SC_{int} = cor($\hat{\eta}$, (LU×HET)*)), parallel to the main-effect SC. Interaction terms for which SC_{int} and the interaction coefficient disagreed in sign were classified as geometrically unstable and were not interpreted as ecological moderation, regardless of AICc support. This product-term SC screen resolved the ambiguity between interactions that reflect genuine conditional dependence and those that reflect collinearity-induced coefficient instability in the interaction layer (Appendix 2.S2.4).

Scale-coherence of supported interactions was assessed in Results: combinations with support concentrated at a single scale or lacking interpretable magnitude patterns across the grain–extent factorial were noted as scale-isolated and interpreted with caution rather than excluded by a fixed filter.

As an a priori motivated sensitivity analysis testing the alternative explanation that observed LU–HET moderation reflects energy-mediated demographic differences rather than genuine heterogeneity dependence (Table 2.1), we refitted each SC-clean supported interaction with AET added as the sole downstream energy control. PET and ALPHA entered only as targeted proxy-swap robustness checks on confirmed surfaces, not as additional controls or parallel moderation arms. For each AET-control or proxy-swap model, we recomputed the spatial autocovariate from the residuals of the corresponding non-spatial model, ensuring that energy-related spatial structure was not absorbed into the autocovariate. Coefficient stability across energy controls was assessed as a robustness check; proxy-specific shifts were interpreted in terms of the documented LU–energy coupling geometry rather than as evidence for or against energy moderation per se.

As a secondary mechanism-parsing analysis, we additionally fitted hierarchical three-way LU × HET × AET models (LU + HET + AET + LU:HET + LU:AET + HET:AET + LU:HET:AET) for each SC-clean two-way LU × HET combination, using AET as the sole downstream energy proxy; PET and the Priestley–Taylor ratio (ALPHA) entered only as proxy-swap robustness checks on the confirmed two-way surfaces, not as parallel moderation arms. The estimand was whether realised water–energy availability made heterogeneity effects under land use more favourable; the three-way coefficient direction was interpreted only when the product-term structure coefficient for the three-way term was concordant with its coefficient, and a discordant SC_triple was treated as geometrically unstable.

We validated supported surfaces using 250-iteration parametric bootstrapping, the inferential approach recommended for generalised linear mixed models where asymptotic approximations are unreliable (Bolker et al. 2009; Fieberg et al. 2020). For retained and confirmed interactions we simulated N = 250 response vectors from each fitted model, refitted the same formula to each, and extracted bootstrap distributions for the standardised partial coefficient (BW) and the structure coefficient (SC). We treated BW intervals as the coefficient-stability evidence and SC as the geometric reliability diagnostic: supported effects whose BW intervals excluded zero in the expected direction were treated as bootstrap-stable, and SC intervals were reported as stricter context rather than as an additional post-hoc gate. Sensitivity of baseline fixed-effect estimates to temporal aggregation and effort formulation is reported in Appendix 2.S4.2; interaction-layer bootstrap and robustness diagnostics are reported in Appendix 2.S5.

We decomposed each aggregate coefficient into guild contributions on the uncontrolled additive layer. For every predictor and scale combination we took the guild coefficient as the median across the additive models in which that predictor entered as a fixed effect — eight for each moderator and twenty for each land-use surrogate, spanning both land-use pairings and all four random-slope variants — weighted each guild coefficient by that guild's share of mean restricted-total richness at the response grain, and computed the cancellation fraction, a summary constructed for this analysis rather than taken from an existing method, as one minus the absolute summed contribution divided by the summed absolute contributions. Cancellation is therefore exactly zero when all three guild contributions share sign, so it measures guild sign opposition rather than effect magnitude. We report the median cancellation across the eighteen grain ×

extent combinations, and Appendix 2.S4 reports the across-model spread of the guild coefficients as a diagnostic.

*Abbreviations used in figures and tables***.** AET, actual evapotranspiration; ALPHA, Priestley–Taylor coefficient; BW, standardised partial regression coefficient; CLIMD, diversity of bioclimatic classes; CROP, cropland prevalence; EVI, enhanced vegetation index; EVICV, coefficient of variation of EVI; EVIENT, entropy of EVI; GEOD, geomorphological landform diversity; NNCOV, non-natural cover; PET, potential evapotranspiration; SC, structure coefficient; TOPOD, topographic ruggedness. Full definitions and provenance are given in Table 2.S2.1.

Results

Baseline responses, additive signs, and the sensitivity gradient

Land-use conversion reduced richness across all strata, and the reduction scaled with forest dependence. In the cropland × greenness-magnitude family the specialist cropland coefficient was −0.335 unconditioned and −0.117 under the terrain-and-energy control set, as actual evapotranspiration and terrain ruggedness absorbed shared spatial structure, and the proportion of broad-domain scale cells whose interval excluded zero fell from 100% to none. Non-forest richness moved oppositely under conditioning, its cropland association strengthening — the first instance of the specificity comparison that structures the vegetational results.

Terrain and climatic heterogeneity behaved as stable additive predictors. Specialist terrain-ruggedness and climatic-heterogeneity main effects held sign and magnitude from the unconditioned to the controlled specification in the cropland-paired models (terrain +0.408 → +0.422; climate +0.238 → +0.282), and non-forest richness reversed both. Terrain ruggedness carried the strongest marginal association in the dataset yet produced no coherent moderation surface at any scale; we treat it as an additive baseline predictor, not a moderator.

Because the guild strata carry lower mean richness than the aggregate responses, their coefficients are estimated with lower precision under the Poisson likelihood (Appendix 2.S1.1). We therefore read specialist-supported interactions as conservative where they recur despite that precision gradient, rather than ranking strata by pooled additive-effect magnitudes.

The vegetational-heterogeneity main effect is conditioning-dependent and specialist-specific: specialist greenness-magnitude variability (EVICV) was negative across three convergent pre-control diagnostics but reversed under the terrain–energy control set, which we read as suppression rather than the ecological sign; pooled and facultative EVICV were not cleanly signed once model geometry entered. The interaction below does not depend on resolving this (Appendix 2.S2.4).

Control-robust EVICV compounding and energy buffering

The specialist cropland × greenness-magnitude interaction is the central result, and unlike the main effect it is unambiguous under conditioning (Table 2.2). The negative CROP × EVICV two-way held its sign across all 15 supported specialist cells under terrain control, energy control, and both jointly. The coefficient became more negative under energy control and under joint terrain-and-energy conditioning in all 15 cells, but became less negative under terrain control alone in 14 of 15 cells. The strengthening therefore occurred under energy control and joint terrain-and-energy control but not under terrain conditioning alone. At grain 10 km and extent 70 km, the CROP × EVICV interaction coefficient changed from −0.080 before conditioning to −0.101 under joint control; all 15 supported cells had bootstrap 95% confidence intervals excluding zero. The product term is negative and control-robust in its own right, independent of the main-effect ambiguity: cropland conversion and high greenness-magnitude variability compound in their cost to forest-specialist richness (Figure 2.3).

Realised water–energy flux modulated this compounding rather than removing it. The three-way CROP × EVICV × AET term was positive and structure-coefficient-concordant across the full specialist grid (18 of 18 cells); a positive three-way term on a negative two-way means the compounding weakens as energy rises — the specialist cost of cropland at high heterogeneity is at full strength in low-energy landscapes and buffered as evapotranspiration increases. The buffering direction is robust grid-wide; it retains incremental support over the conditioned two-way in 15 of 18 cells, concentrated at fine-to-mid extents, and four of the six reported bootstrap summaries exclude zero. Energy buffering is a real but scale-bounded modulation of a surviving effect (Figure 2.4).

The interaction was scale-structured in a way sign-counting cannot capture: the extent at which the surface changes sign is ordered by forest dependence, crossing at fine extent in

specialists and progressively later in less forest-dependent strata, with non-forest inverted throughout. This crossing-point ordering, not any single threshold, is the scale signature of the result.

Two contrasts confirm the mechanism is specific to greenness-magnitude variability, not vegetation heterogeneity in general. Configurational greenness diversity (EVIENT) produced no comparable surface — structure-coefficient-concordant in one specialist cell of eighteen and no non-forest cells, attenuating under conditioning, while NNCOV × EVIENT reached support in 3 of 18 specialist cells with a positive coefficient, against the negative direction predicted by H2 — the same predictor family behaving oppositely. And non-forest richness, the specificity diagnostic, showed a positive EVICV association in both layers, positive in direction across the matrix but meeting the joint support criterion in no cell, the mirror image of the specialist response. The compounding is forest-mediated: the forest-dependent guild pays the cost, the non-forest guild gains, and the two responses oppose to the same predictor.

The second land-use surrogate corroborated the cost with a scale bound. Specialist NNCOV × EVICV was negative and supported at broad extent but structure-coefficient-discordant through the mid-scales, so under our joint criterion the corroboration holds only at broad extent; the weaker guilds again responded positively, reproducing the specialist/guild opposition.

Secondary energy moderation and conditional geophysical buffering

A second, energetic moderation axis was confirmed but subordinate. NNCOV × AET was control-robust with a forest-dependence sign-split — negative in specialists (12 of 12 cells structure-coefficient-concordant) and positive in non-forest richness (17 of 17 concordant), with specialist 10 of 12 and non-forest 15 of 17 cells showing bootstrap intervals excluding zero, and the crossing ordered as in the vegetational surface. The specialist signal survived substitution of the energy proxy by its components (34 of 36 and 32 of 36 cells), so it is not an operationalisation artefact; the non-forest signal was proxy-sensitive and is diagnostic only. This confirms an energy axis but does not support a simple productivity-buffering prediction: energy buffered the vegetational surface while the direct interaction split in opposite directions between specialists and the control.

A conditional geophysical result completed the set: specialist CROP × CLIMD was positive, fine-scale, and specialist-specific (supported in the specialist stratum against zero to two supported cells elsewhere), opposite in sign to the vegetational surface and indexing a distinct climatic axis. Under energy conditioning, all eight bootstrap-eligible cells had 95% intervals excluding zero in the positive direction. Its product-term structure coefficient is weak (0.007–0.048); we report it as a secondary, bootstrap-validated pathway, not a co-equal result, with the full surface in Appendix 2.S5.

Remaining geophysical and configurational interactions were absent rather than weak: neither surrogate produced terrain or climatic moderation outside the conditional case (NNCOV × TOPOD reached one supported cell, in full total richness; NNCOV × CLIMD reached five — two in full total, one in restricted total and two in forest-facultative richness — with neither surface supported in forest-specialist or non-forest richness), localising land-use moderation in this system to the vegetational and energetic axes. We set aside five first-tier candidates after the post-hoc directional and non-forest specificity comparisons documented in Appendix 2.S3, with the cell-level evidence resolved in Appendix 2.S5.

Process-specific aggregation limits

Having established which processes moderate land-use effects and in which guilds, we asked what aggregate richness does to them. Decomposing each total-richness coefficient into fitted-mean-weighted guild contributions gives the fraction of signal lost to cancellation when guilds oppose, reported as the median across the eighteen grain × extent combinations: cropland and non-natural cover 0% (guild contributions share sign in almost every cell, and the aggregate reports the full land-use damage); configurational diversity 0%, terrain 19%, and climate 12% (largely preserved); realised energy 37%; and greenness-magnitude variability 59%, where the aggregate coefficient collapses to near zero while the guild-level effects are strong and opposed (Figure 2.5).

The land-use surrogates are the internal validity check: because the guild contributions share sign in almost every cell, opposing in 1 of 18 cropland and 7 of 18 non-natural-cover combinations, a decomposition that manufactured cancellation should show it here first, and both surrogates return a median of zero. The 59% and 37% medians therefore do not arise from a decomposition artefact. Cancellation is process-specific — near-zero for shared land-use

damage, low for terrain and climate, and high only where forest-dependent and non-forest species oppose.

This reproduces the vegetational discriminator a third time, now in the aggregation layer: EVICV cancels 59% while EVIENT cancels 0%, because guilds oppose on greenness-magnitude variability but agree on configurational diversity across most of the scale matrix. The EVIENT contrast is scale-bounded rather than absolute: guild contributions oppose in six of the eighteen combinations, and where they do the cancellation declines monotonically with predictor extent within every grain, so the metric contrast is sharpest at broad extent and weakest at fine grain and short extent. Across the moderation surfaces, the conditioned survival, and this decomposition, the same metric-specific distinction appears — three convergent analytical layers drawn from the same data, mutually corroborating.

The opposition driving this decomposition is not a closure artefact: restricted total is not held fixed, so specialist and non-forest richness are not compositionally closed, and their opposite coefficient signs are not algebraically forced. Total richness is thus not uninformative — it reports the biodiversity cost of conversion faithfully — but it is selectively blind, losing most of the vegetational and a third of the energetic moderation signal precisely where the ecology is strongest.

Discussion

We asked whether, and through which pathways, environmental heterogeneity moderates the biodiversity cost of land use in the Atlantic Forest, testing three a priori arms (vegetational heterogeneity as a compounding axis, geophysical heterogeneity as a buffering axis, and realised water–energy availability as a resource-mediated buffer) across five response strata separating forest specialists, forest facultative species, and non-forest species. Resolving moderation by guild changes the most general conclusion. The broadest conclusion is not any single pathway but what aggregate richness can detect. We predicted that decomposing heterogeneity by origin would recover interpretable sign structure. Instead, the guild decomposition showed that sign structure also depends on which assemblage component carries the response, because aggregation reports the direct cost of conversion faithfully yet becomes selectively blind to the processes that moderate it.

Cropland and non-natural cover cancelled 0% between guilds, so total richness registers the shared cost of conversion, whereas greenness-magnitude heterogeneity cancelled 59% and realised energy 37%, because forest-dependent and non-forest species respond to those gradients in opposite directions. This identifies guild-level opposition as an additional biological source of heterogeneity–diversity relationship (HDR) inconsistency beyond differences in scale, metric choice, and measurement error (Tews et al. 2004; Stein et al. 2014; Wan et al. 2023). The mechanism is co-occurring assemblage components with oppositely-signed responses that cancel at community level. Where that literature has treated inconsistent HDR signs as a measurement or scaling problem to be resolved, our decomposition locates the inconsistency in the biology — in guild-level opposition that any community-level metric averages away.

Resolving inconsistent HDRs: origin, scale, and guild as modulating axes

The origin-and-scale account of HDR inconsistency predicts that abiotic and vegetational heterogeneity should carry opposite signs and that collapsing them obscures both. We found that pattern, and the guild strata add a third averaging axis the literature has not isolated. Geophysical heterogeneity showed stable positive additive associations across the four non-control richness layers, especially terrain ruggedness (TOPOD), with concordant partial slopes and structure coefficients, while the non-forest control reversed sign — the expected behaviour of a persistent template read against a guild that does not depend on forest structure.

The vegetational signal was not uniformly negative across the community, as a single-metric reading would imply; it was negative and control-robust specifically in forest specialists, where three pre-control diagnostics agreed, namely the univariate slope, the uncontrolled EVICV structure coefficient (−0.101), and the primary partial slope (−0.045), while pooled and facultative strata did not show a clean negative additive association once collinearity geometry entered. That specificity is the refinement. The controversy over HDR sign (Tews et al. 2004; Stein et al. 2014; Wan et al. 2023) reflects a failure to decompose heterogeneity not only by origin and scale but by response guild, because collapsing forest specialists with matrix species into one richness total averages across oppositely-signed regimes. This sharpens rather than contradicts the origin-dependent view. The negative pole is real but guild-located, not community-wide.

Our results are consistent with the broader finding that HDRs are scale-contingent and mechanistically diverse. HDRs may follow unimodal or mixed forms, shaped by a trade-off between niche diversification and the demographic constraints of spatial subdivision (Allouche et al. 2012; Laanisto et al. 2013), and negative HDRs dominate at fine grains in fragmented systems where anthropogenic heterogeneity raises habitat contrast (Fahrig et al. 2011). That non-positive HDRs are empirically common rather than anomalous is directly supported by multi-taxa work documenting monotonically decreasing relationships for roughly a quarter of taxon-specific cases in a single forest system (Heidrich et al. 2020). Our specialist result agrees with that literature and specifies the guild in which the negative relationship concentrates. The facet-specific reading these patterns support — abiotic heterogeneity as a persistent template that expands opportunity, vegetational heterogeneity as a land-use-responsive axis whose sign depends on what it indexes (Hjort et al. 2015; Räsänen et al. 2016) — is the frame we carry into the mechanism sections below.

Land-use-coupled vegetational degradation is metric-specific and forest-mediated

We predicted that land-use-responsive vegetation texture would compound rather than buffer habitat loss. We recovered that pattern for one metric and one guild, which narrows the original prediction. Cropland compounded the richness cost of high greenness-magnitude heterogeneity (EVICV) in forest specialists across all 15 supported cells. That interaction is control-robust. The CROP × EVICV interaction coefficient became more negative, from −0.080 before conditioning to −0.101 under joint terrain-and-energy conditioning at the finest grain and broad extent, with bootstrap intervals excluding zero in every supported cell. Configurational greenness diversity (EVIENT), derived from the same imagery, produced no comparable surface. The distinction is therefore metric-specific rather than a simple vegetational–geophysical split. Both metrics are vegetation-derived, and EVIENT is more strongly coupled to the land-use surrogates than EVICV, yet EVIENT does not reproduce the forest-specialist moderation or the aggregation-cancellation pattern. The relevant discriminator therefore lies in what the vegetation metrics capture and how those properties enter the biodiversity response, not simply in whether they respond to land use. We recover the same EVICV–EVIENT distinction in the moderation surfaces, their behaviour under conditioning, and the aggregation decomposition. These are convergent analytical layers from one dataset rather than independent confirmations. We therefore support the narrower conclusion that greenness-magnitude variability compounds land

use in forest specialists, rather than the broader claim that vegetation heterogeneity generally compounds land use.

We interpret this compounding surface as consistent with a fragmentation-mediated area–heterogeneity trade-off, without treating remotely sensed texture as a direct fragmentation measure. In the highly fragmented Atlantic Forest, vegetation disorder reflects the structural consequences of land-cover replacement — patch boundaries, heterogeneous canopy, proliferating edges, matrix contrast (Laurance et al. 2012; Vancine et al. 2024) — and for forest-dependent assemblages with minimum-area requirements this configurational heterogeneity reduces effective habitat area and elevates extinction risk, the mechanism predicted by area–heterogeneity trade-off and microfragmentation theory (Kadmon and Allouche 2007; Allouche et al. 2012; Tamme et al. 2010; Laanisto et al. 2013). Three qualifications from the same literature bound the inference. First, richness–heterogeneity relationships can differ among metrics and contexts within one landscape (Bar-Massada and Wood 2014), which is exactly what the EVICV/EVIENT contrast shows and why we do not generalise beyond the magnitude metric. Second, indirect proxies can produce misleading trade-off support when confounded with mean environmental state or uneven sampling along land-use gradients (Hortal et al. 2013; Carnicer et al. 2013). Third, area–heterogeneity trade-offs can carry positive as well as negative richness consequences depending on whether reduced effective area increases stochastic extinction or relaxes competitive exclusion (Ben-Hur and Kadmon 2020). The pattern is therefore consistent with, but does not demonstrate, a fragmentation-mediated mechanism.

The opposition that produces the cancellation is forest-mediated, and the non-forest stratum sharpens that reading as a specificity diagnostic. Wherever forest specialists showed a negative conditional response (EVICV compounding and the NNCOV × AET specialist signal) non-forest richness moved in the opposite direction on the same predictor, and because restricted total is not held fixed, those opposite signs are not algebraically forced by the partition. The guild expected to benefit from persistent heterogeneity is the one harmed most where heterogeneity indexes land-use-responsive texture. That is a sharper statement of the origin-dependence the chapter set out to test, now located in a specific guild contrast rather than assumed as a category. At fine grains configurational disruption dominates the vegetation signal; at broader extents vegetation heterogeneity begins to capture biogeographic variation in forest structure — the

transition from demographic constraint to niche-diversification signal anticipated by the trade-off model (Allouche et al. 2012).

Geophysical heterogeneity as an additive template, not a moderator

We predicted geophysical heterogeneity would buffer land-use effects, but we found no general geophysical buffering. Terrain acted primarily as a strong additive template without a coherent moderation surface, whereas climatic heterogeneity provided a secondary, scale-bounded specialist buffering pathway. The correction matters for how the buffering literature is applied to transformed systems. Terrain and climatic heterogeneity represent persistent physical gradients rather than land-use-responsive states, and they filter composition, maintain refugia, and support beta diversity regardless of transformation history (Hjort et al. 2015; Tukiainen et al. 2017; Stein et al. 2014; Tews et al. 2004), and TOPOD carried the strongest marginal richness association in our dataset, consistent with that template role. But terrain produced no coherent moderation surface at any scale in the guild-resolved analysis, so we read TOPOD as a persistent additive driver rather than a buffer of land-use effects.

This refines the buffering literature. Persistent heterogeneity marks high-richness terrain without necessarily reducing the marginal cost of conversion where it occurs. We do not treat the positive abiotic association as fully decoupled from land use. Terrain can direct where human activity concentrates, making abiotic heterogeneity partly an upstream cause of landscape pattern (Räsänen et al. 2016), but we did not test conversion avoidance directly here. A taxon-specific caveat also bounds the template reading. Topographic diversity raised tree richness but lowered bird richness in North American assemblages (Read et al. 2020), so the broadly positive TOPOD association is not guaranteed and warrants scrutiny of passerine responses.

Energy conditions the vegetational surface rather than buffering conversion directly

We predicted that higher realised productivity would weaken land-use costs through a demographic pathway. We use the conditioning and three-way models to answer different questions. Adding AET as a conditioner changes the estimated average CROP × EVICV interaction, whereas the positive three-way term shows that, within the conditioned surface, compounding weakens as realised energy increases. The evidence refines that prediction in two ways. First, energy buffers the vegetational surface conditionally rather than offsetting conversion in general. The CROP × EVICV × AET term was positive and structure-coefficient-

concordant across the full specialist grid, meaning the specialist cost of cropland at high heterogeneity is at full strength in low-energy landscapes and progressively buffered as evapotranspiration rises, with incremental support over the conditioned two-way retained in 15 of 18 cells, concentrated at fine-to-mid extents, and four of the six reported bootstrap summaries excluding zero. The conditioning effect is therefore supported within a bounded range of scales, where a proxy-by-proxy reading of a coarser design returned an ambiguous energy signal. Second, the direct land-use × energy relation is not a uniform buffer but a sign-split. NNCOV × AET was negative in forest specialists (12 of 12 supported cells) and positive in the non-forest control (17 of 17), so energy's clearest role is conditioning the vegetational surface, not independently offsetting conversion. Marginal proxy associations remain consistent with the water–energy framework, realised flux relating positively to richness (Currie 1991; Hawkins et al. 2003; Evans et al. 2005). Because the design carries no abundance mediator, we interpret this as energy-conditioned moderation rather than a test of more-individuals mechanisms. Distinguishing abundance-mediated persistence from compositional turnover requires orthogonalised productivity predictors or abundance data (Hurlbert and Stegen 2014), which we identify as the direct extension rather than a claim of the present design.

A secondary, conditional climatic pathway

Climatic heterogeneity was predicted to act as geophysical buffering, and cropland × climatic diversity partially confirms that prediction. It buffered specialist richness at fine scales, positive and specialist-specific, on a distinct axis from the vegetational surface and opposite in sign, with every bootstrap-eligible cell clearing zero under energy conditioning. Its product-term structure-coefficient support is weak, because taxa within the specialist guild disagree on the direction of the moderated effect. We therefore treat it as a secondary pathway rather than a co-equal result. That a climatic signal buffers where vegetational heterogeneity compounds is consistent with the geophysical-heterogeneity expectation that persistent abiotic gradients can expand environmental opportunity (Kerr and Packer 1997; Hjort et al. 2015; Tukiainen et al. 2017), while the weak product-term structure coefficient keeps the result secondary rather than co-equal. The area–heterogeneity trade-off remains the relevant alternative, because high climatic diversity within a fixed area could instead reduce effective area per climatic condition (Kadmon and Allouche 2007; Allouche et al. 2012). Either way, the result reinforces the central point that heterogeneity has no single sign. Its effect under land use depends on what it indexes.

Origin-dependent moderation and its boundary conditions

Our results do not refute landscape moderation but identify its boundary conditions in a severely transformed system. Tscharntke et al. (2012) organise moderation into multiple conditional pathways (species-pool effects, beta diversity, spillover, insurance, trait filtering) rather than a single buffering prediction, and our findings map onto that framework. The confirmed moderation is vegetational (compounding) and energetic (conditioning), while land-use-responsive vegetational heterogeneity did not buffer and instead compounded loss for the forest-dependent guild. The relevant contrast is therefore not whether moderation works or fails but which form of heterogeneity, in which guild, activates which pathway. This is consistent with threshold responses in fragmented landscapes, where demographic recovery collapses beyond a critical level of habitat loss (Swift and Hannon 2010; Haddad et al. 2015; Rybicki and Hanski 2013), so the conservation value of heterogeneity depends on its origin and on whether it remains biologically usable under land-use pressure (Swift and Hannon 2010; Laurance et al. 2012; Haddad et al. 2015; Chase et al. 2020). Compensatory dynamics offer no rescue here. Richness stability through colonisation from a regional pool requires trait-compatible colonists, adequate connectivity, and response diversity (Gonzalez and Loreau 2009; Mori et al. 2013), conditions compromised by fragmentation and trait filtering in this system, and our forest-specialist guild — the one for which vegetational heterogeneity compounds loss — is precisely the guild least able to compensate.

Caveats and direct tests

First, EVICV and EVIENT capture compositional and configurational vegetation texture from above, not the process that generated them. Vegetation texture rises through degradation, secondary regrowth, agricultural mosaics, edge proliferation or natural variation in vegetation state, and it does not separate forest amount from configuration (Tuanmu and Jetz 2015). We therefore read the CROP × EVICV surface as compounding along an observed, land-use-responsive vegetation state rather than as a test of fragmentation per se — the distinction at the centre of the debate over whether habitat loss or fragmentation drives biodiversity loss (Fletcher et al. 2018; Fahrig et al. 2019). Explicit metrics separating patch configuration from forest amount (Wang et al. 2014) would test that reading directly, and we treat the separation as the next inferential step rather than a claim established here.

Second, we treat AET as conditioning the vegetational surface rather than as evidence of demographic buffering. The positive CROP × EVICV × AET term shows that the fitted negative two-way weakens as realised water–energy flux rises, and the PET and ALPHA proxy swaps reduce the chance that this pattern depends on one energy operationalisation. Because we model richness rather than abundance, population size or demographic rates, we do not infer a more-individuals pathway from this surface. Abundance or composition data would test it.

Third, the guild comparison and the cancellation decomposition reveal response structure rather than compensation. The greenness-magnitude and energy cancellation estimates quantify how fitted guild contributions combine into weaker aggregate coefficients, but they do not show temporal replacement, colonisation balancing extinction, or demographic compensation. A gradient-dependent strengthening of that cancellation is visible in raw guild-share shifts but is confounded with taxon sampling along the same gradient, so we raise it as a hypothesis for fitted-model recomputation rather than a result. The process-specificity of the cancellation is unaffected, resting on the fitted decomposition and the cropland closure control (0% cancellation).

Fourth, the environmental layers and the assemblage records differ in temporal reference. Cropland prevalence represents the 1992–2006 consensus period, the EVI texture metrics the 2000–2005 product window and non-natural cover a 2020 land-cover state, and the water–energy layers long-term climatologies centred on 1950–1990, whereas assemblages span 1981 to 2018 (Appendix 2.S2.1). We therefore read the fitted gradients as long-term spatial contrasts in landscape context rather than as land-use histories matched in time to each assemblage, and a design resolving conversion timing against survey date would test whether the compounding surface tracks recent conversion specifically.

Finally, the ordered scale crossings show where the fitted surface changes across the matrix, not why it changes there. Grain and extent alter ecological context together with predictor averaging, buffer overlap and spatial dependence, and the recurring forest-dependence ordering shows only that the headline surface is not an isolated scale artefact. The design also substitutes space for time within a large observational framework, standard for inferring dynamics from spatial pattern but carrying well-known inferential limits (De Palma et al. 2018; Purvis et al. 2018), and we constrain all causal language above to association. Temporal

resolution is limited on the biodiversity side as well: records combining several years of survey enter on the same footing as records from a single period, because the modelling data retain reported effort and the number of consolidated assemblages but not the span over which those surveys accumulated.

Conservation implications

In transformed tropical forests, conserving heterogeneity is insufficient. Conservation must distinguish persistent geophysical heterogeneity from land-use-generated structural disorder, and must recognise that aggregate richness conceals the guild-level processes that drive loss. Geophysically heterogeneous terrain warrants attention because it supported high baseline richness, and because terrain can shape where human activity concentrates (Räsänen et al. 2016), although our results do not support treating ruggedness as an active buffer of conversion. We should not treat remotely sensed texture as a simple biodiversity proxy in degraded systems. General-purpose texture metrics capture variability irrespective of land-cover type (Tuanmu and Jetz 2015) but do not index forest configuration, and greenness-magnitude heterogeneity is negatively associated with forest-specialist richness even while geophysical heterogeneity is positively associated overall — an opposition aggregate richness masks. Because vegetation disorder in this biome tracks the fragmentation generated by land-cover replacement (Laurance et al. 2012; Vancine et al. 2024), remotely sensed texture must be read alongside land-use context, spatial scale, and response guild before informing monitoring indices or the conditional planning pathways that landscape moderation theory identifies (Tscharntke et al. 2012). Where forest-dependent and matrix-associated species coexist, guild-resolved analysis is not optional. Total richness registers the cost of conversion but masks its moderation, so setting defensible restoration targets requires resolving richness by forest affiliation rather than managing for heterogeneity per se.

Conclusions

The field-level question we opened with, whether landscape heterogeneity buffers biodiversity against land-use change, has no single answer, and that is the finding. Heterogeneity buffers, compounds, or fails to moderate depending on what it indexes and on which guild is asked. Land-use-responsive vegetation texture compounded the cost of conversion: cropland

deepened the richness decline where greenness-magnitude heterogeneity was high, and the effect concentrated in forest specialists. Persistent geophysical heterogeneity did not buffer generally. Terrain acted as a strong additive template without a coherent moderation surface, and climatic heterogeneity provided only a secondary, scale-bounded specialist pathway. Realised water-energy availability conditioned the vegetational surface rather than demonstrating demographic buffering. Aggregate richness recovered the shared cost of land use but masked most of this structure, because the guilds responded in opposing directions. The community-level richness measure through which the question is usually posed is therefore least able to reveal the guild-opposed moderation it is asked to detect. Heterogeneity is not a general safeguard against land-use change, and assessments that treat it as one, or that read moderation from aggregate richness alone, will misstate both the direction and the beneficiaries of the effect.

References

Allouche O, Kalyuzhny M, Moreno-Rueda G, Pizarro M, Kadmon R (2012) Area-heterogeneity tradeoff and the diversity of ecological communities. Proceedings of the National Academy of Sciences 109:17495-17500.

Amatulli G, Domisch S, Tuanmu M-N, et al (2018) A suite of global, cross-scale topographic variables for environmental and biodiversity modelling. Scientific Data 5:180040.

Bar-Massada A, Wood EM (2014) The richness-heterogeneity relationship differs between heterogeneity measures within and among habitats. Ecography 37:528-535.

Barr DJ, Levy R, Scheepers C, Tily HJ (2013) Random effects structure for confirmatory hypothesis testing: Keep it maximal. Journal of Memory and Language 68:255-278.

Baston D (2023) exactextractr: fast extraction from raster datasets using polygons. R package version 0.10.0.

Ben-Hur E, Kadmon R (2020) An experimental test of the area-heterogeneity tradeoff. Proceedings of the National Academy of Sciences 117:4815-4822

Bolker BM, Brooks ME, Clark CJ, Geange SW, Poulsen JR, Stevens MHH, White J-SS (2009) Generalized linear mixed models: a practical guide for ecology and evolution. Trends in Ecology and Evolution 24:127-135.

Bovendorp RS, Villar N, de Abreu-Junior EF, et al (2017) ATLANTIC small-mammal: a dataset of communities of rodents and marsupials of the Atlantic forests of South America. Ecology 98:2226.

Brooks ME, Kristensen K, van Benthem KJ, et al (2017) glmmTMB balances speed and flexibility among packages for zero-inflated data. The R Journal 9:378-400.

Burnham KP, Anderson DR (2002) Model Selection and Multimodel Inference: A Practical Information-Theoretic Approach, 2nd edn. Springer-Verlag, New York.

Burnham KP, Anderson DR, Huyvaert KP (2011) AIC model selection and multimodel inference in behavioral ecology: some background, observations, and comparisons. Behavioral Ecology and Sociobiology 65:23-35.

Carnicer J, Brotons L, Herrando S, Sol D (2013) Improved empirical tests of area-heterogeneity tradeoffs. Proceedings of the National Academy of Sciences 110:E2858-E2860.

Chase JM, Liebergesell M, Sagouis A, et al (2020) FragSAD: a database of diversity and species abundance distributions from habitat fragments. Ecology 100:e02775.

Coelho MTP, Barreto E, Rangel TF, Diniz-Filho JAF, Wüest RO, Bach W, Skeels A, McFadden IR, Roberts DW, Pellissier L, Zimmermann NE (2023) The geography of climate and the global patterns of species diversity. Nature 622:537-544.

Crase B, Liedloff AC, Wintle BA (2012) A new method for dealing with residual spatial autocorrelation in species distribution models. Ecography 35:879-888.

Currie DJ (1991) Energy and large-scale patterns of animal- and plant-species richness. The American Naturalist 137:27-49.

De Palma A, Sanchez-Ortiz K, Martin PA, Chadwick A, Gilbert G, Bates AE, Börger L, Contu S, Hill SLL, Purvis A (2018) Challenges with inferring how land-use affects terrestrial biodiversity: study design, time, space and synthesis. Advances in Ecological Research 58:163-199.

Donald PF, Arendarczyk B, Spooner F, Buchanan GM (2019) Loss of forest intactness elevates global extinction risk in birds. Animal Conservation 22:341-347.

Dormann CF, Elith J, Bacher S, et al (2013) Collinearity: a review of methods to deal with it and a simulation study evaluating their performance. Ecography 36:27-46.

Evans KL, Warren PH, Gaston KJ (2005) Species-energy relationships at the macroecological scale: a review of the mechanisms. Biological Reviews 80:1-25.

Fahrig L, Baudry J, Brotons L, et al (2011) Functional landscape heterogeneity and animal biodiversity in agricultural landscapes. Ecology Letters 14:101-112.

Fahrig L, Arroyo-Rodríguez V, Bennett JR, et al (2019) Is habitat fragmentation bad for biodiversity? Biological Conservation 230:179-186.

Fieberg JR, Vitense K, Johnson DH (2020) Resampling-based methods for biologists. PeerJ 8:e9089.

Figueiredo MSL, Barros CS, Delciellos AC, et al (2017) Abundance of small mammals in the Atlantic Forest (ASMAF): a data set for analyzing tropical community patterns. Ecology 98:2981.

Fletcher RJ Jr, Fortin M-J (2018) Spatial Ecology and Conservation Modeling. Springer, Cham.

Fletcher RJ Jr, Didham RK, Banks-Leite C, et al. (2018) Is habitat fragmentation good for biodiversity? Biological Conservation 226:9-15

Gonzalez A, Loreau M (2009) The causes and consequences of compensatory dynamics in ecological communities. Annual Review of Ecology, Evolution, and Systematics 40:393-414.

Haddad NM, Brudvig LA, Clobert J, et al (2015) Habitat fragmentation and its lasting impact on Earth's ecosystems. Science Advances 1:e1500052.

Harrison XA (2014) Using observation-level random effects to model overdispersion in count data in ecology and evolution. PeerJ 2:e616.

Hartig F (2022) DHARMa: residual diagnostics for hierarchical (multi-level / mixed) regression models. R package version 0.4.6.

Hasui E, Metzger JP, Pimentel RG, et al (2018) ATLANTIC BIRDS: a data set of bird species from the Brazilian Atlantic Forest. Ecology 99:497.

Hawkins BA, Field R, Cornell HV, et al (2003) Energy, water, and broad-scale geographic patterns of species richness. Ecology 84:3105-3117.

Heidrich L, Bae S, Friess N, et al (2020) Heterogeneity-diversity relationships differ between and within trophic levels in temperate forests. Nature Ecology and Evolution 4:1204-1212.

Hesselbarth MHK, Sciaini M, With KA, et al (2019) landscapemetrics: an open-source R tool to calculate landscape metrics. Ecography 42:1648-1657.

Hijmans RJ (2022) terra: spatial data analysis. R package version 1.6-17.

Hjort J, Gordon JE, Gray M, Hunter ML Jr (2015) Why geodiversity matters in valuing nature's stage. Conservation Biology 29:630-639.

Hortal J, Carrascal LM, Triantis KA, Thébault E, Meiri S, Sfenthourakis S (2013) Species richness can decrease with altitude but not with habitat diversity. Proceedings of the National Academy of Sciences 110:E2149-E2150.

IUCN (2012) Habitats Classification Scheme, Version 3.1. International Union for Conservation of Nature, Gland, Switzerland.

Hurlbert AH, Stegen JC (2014) When should species richness be energy limited, and how would we know? Ecology Letters 17:401-413.

Hutchinson GE (1957) Concluding remarks. Cold Spring Harbor Symposia on Quantitative Biology 22:415-427.

Kadmon R, Allouche O (2007) Integrating the effects of area, isolation, and habitat heterogeneity on species diversity: a unification of island biogeography and niche theory. The American Naturalist 170:443-454.

Kerr JT, Packer L (1997) Habitat heterogeneity as a determinant of mammal species richness in high-energy regions. Nature 385:252-254.

Laanisto L, Tamme R, Hiiesalu I, et al (2013) Microfragmentation concept explains non-positive environmental heterogeneity-diversity relationships. Oecologia 171:217-226.

Laurance WF, Useche DC, Rendeiro J, et al (2012) Averting biodiversity collapse in tropical forest protected areas. Nature 489:290-294.

Lima F, Beca G, Muylaert RL, et al (2017) ATLANTIC-CAMTRAPS: a dataset of medium and large terrestrial mammal communities in the Atlantic Forest of South America. Ecology 98:2979.

MacArthur RH (1972) Geographical Ecology. Harper and Row, New York.

Metzger MJ, Bunce RGH, Jongman RHG, Sayre R, Trabucco A, Zomer R (2013) A high-resolution bioclimate map of the world: a unifying framework for global biodiversity research and monitoring. Global Ecology and Biogeography 22:630-638.

Mori AS, Furukawa T, Sasaki T (2013) Response diversity determines the resilience of ecosystems to environmental change. Biological Reviews 88:349-364.

Moura MR, Villalobos F, Costa GC, Garcia PCA (2016) Disentangling the role of climate, topography and vegetation in species richness gradients. PLoS ONE 11:e0152468.

Moura MR, Costa HC, Argôlo AJS, et al (2017a) Environmental constraints on the compositional and phylogenetic beta-diversity of tropical forest snake assemblages. Journal of Animal Ecology 86:1192-1204.

Moura MR, Argôlo AJ, Costa HC (2017b) Historical and contemporary correlates of snake biogeographical subregions in the Atlantic Forest hotspot. Journal of Biogeography 44:640-650.

Muylaert RL, Stevens RD, Esbérard CEL, et al (2017) ATLANTIC BATS: a data set of bat communities from the Atlantic Forests of South America. Ecology 98:3227.

Muylaert RL, Vancine MH, Bernardo R, et al (2018) Uma nota sobre os limites territoriais da Mata Atlântica. Oecologia Australis 22:302-311.

Newbold T, Hudson LN, Hill SLL, et al (2015) Global effects of land use on local terrestrial biodiversity. Nature 520:45-50.

O'Hara RB, Kotze DJ (2010) Do not log-transform count data. Methods in Ecology and Evolution 1:118-122.

Paradis E, Schliep K (2019) ape 5.0: an environment for modern phylogenetics and evolutionary analyses in R. Bioinformatics 35:526-528.

Pebesma E (2018) Simple features for R: standardized support for spatial vector data. The R Journal 10:439-446.

Purvis A, Newbold T, De Palma A, et al (2018) Modelling and Projecting the Response of Local Terrestrial Biodiversity Worldwide to Land Use and Related Pressures: The PREDICTS Project. Advances in Ecological Research 58:1-35.

R Core Team (2024) R: A Language and Environment for Statistical Computing. R Foundation for Statistical Computing, Vienna, Austria.

Ray-Mukherjee J, Nimon K, Mukherjee S, Morris DW, Slotow R, Hamer M (2014) Using commonality analysis in multiple regressions: a tool to decompose regression effects in the face of multicollinearity. Methods in Ecology and Evolution 5:320-328.

Read QD, Zarnetske PL, Record S, Dahlin KM, Costanza JK, Finley AO, Gaddis KD, et al (2020) Beyond counts and averages: relating geodiversity to dimensions of biodiversity. Global Ecology and Biogeography 29:696-710.

Ribeiro MC, Metzger JP, Martensen AC, et al (2009) The Brazilian Atlantic Forest: How much is left, and how is the remaining forest distributed? Implications for conservation. Biological Conservation 142:1141-1153.

Rybicki J, Hanski I (2013) Species-area relationships and extinctions caused by habitat loss and fragmentation. Ecology Letters 16:27-38.

Räsänen A, Kuitunen M, Hjort J, Vaso A, Kuitunen T, Lensu A (2016) The role of landscape, topography, and geodiversity in explaining vascular plant species richness in a fragmented landscape. Boreal Environment Research 21:53-70.

Santos JP, Freitas AVL, Brown KS, et al (2018) ATLANTIC BUTTERFLIES: a data set of fruit-feeding butterfly communities from the Atlantic Forests. Ecology 99:2875.

Schielzeth H (2010) Simple means to improve the interpretability of regression coefficients. Methods in Ecology and Evolution 1:103-113.

Seiferling I, Proulx R, Wirth C (2014) Disentangling the environmental-heterogeneity-species-diversity relationship along a gradient of human footprint. Ecology 95:2084-2095.

Smith AC, Koper N, Francis CM, Fahrig L (2009) Confronting collinearity: comparing methods for disentangling the effects of habitat loss and fragmentation. Landscape Ecology 24:1271-1285.

Souza Y, Gonçalves F, Lautenschlager L, et al (2019) ATLANTIC MAMMALS: a data set of assemblages of medium- and large-sized mammals of the Atlantic Forest of South America. Ecology 100:e02785.

Stein A, Gerstner K, Kreft H (2014) Environmental heterogeneity as a universal driver of species richness across taxa, biomes and spatial scales. Ecology Letters 17:866-880.

Stein A, Kreft H (2015) Terminology and quantification of environmental heterogeneity in species-richness research. Biological Reviews 90:815-836.

Storch D, Bohdalková E, Okie J (2018) The more-individuals hypothesis revisited: the role of community abundance in species richness regulation and the productivity-diversity relationship. Ecology Letters 21:920-937.

Swift TL, Hannon SJ (2010) Critical thresholds associated with habitat loss: a review of the concepts, evidence, and applications. Biological Reviews 85:35-53.

Tamme R, Hiiesalu I, Laanisto L, Szava-Kovats R, Pärtel M (2010) Environmental heterogeneity, species diversity and co-existence at different spatial scales. Journal of Vegetation Science 21:796-801.

Tews J, Brose U, Grimm V, et al (2004) Animal species diversity driven by habitat heterogeneity/diversity: the importance of keystone structures. Journal of Biogeography 31:79-92.

Title PO, Bemmels JB (2018) ENVIREM: an expanded set of bioclimatic and topographic variables increases flexibility and improves performance of ecological niche modeling. Ecography 41:291-307.

Trabucco A, Zomer R (2019) Global High-Resolution Soil-Water Balance. figshare Dataset. https://doi.org/10.6084/m9.figshare.7707605.v3

Tscharntke T, Tylianakis JM, Rand TA, et al (2012) Landscape moderation of biodiversity patterns and processes - eight hypotheses. Biological Reviews 87:661-685.

Tuanmu M-N, Jetz W (2014) A global 1-km consensus land-cover product for biodiversity and ecosystem modelling. Global Ecology and Biogeography 23:1031-1045.

Tuanmu M-N, Jetz W (2015) A global, remote sensing-based characterization of terrestrial habitat heterogeneity. Global Ecology and Biogeography 24:1329-1339.

Tukiainen H, Bailey JJ, Field R, Kangas K, Hjort J (2017) Combining geodiversity with climate and topography to account for threatened species richness. Conservation Biology 31:364-375.

Turner MG (1989) Landscape ecology: the effect of pattern on process. Annual Review of Ecology and Systematics 20:171-197.

Vancine MH, Duarte K da S, de Souza YS, et al (2018) ATLANTIC AMPHIBIANS: a data set of amphibian communities from the Atlantic Forests of South America. Ecology 99:1692.

Vancine MH, Muylaert RL, Niebuhr BB, et al (2024) The Atlantic Forest of South America: spatiotemporal dynamics of the vegetation and implications for conservation. Biological Conservation 291:110499.

Wan J-Z, Wang C-J, Marquet PA (2023) Environmental heterogeneity as a driver of terrestrial biodiversity on a global scale. Progress in Physical Geography 47:912-930.

Wang X, Blanchet FG, Koper N (2014) Measuring habitat fragmentation: an evaluation of landscape pattern metrics. Methods in Ecology and Evolution 5:634-646.

Wood SN (2004) Stable and efficient multiple smoothing parameter estimation for generalized additive models. Journal of the American Statistical Association 99:673-686.

Table 2.1. Predictions tested in this study, with hypothesised mechanisms and alternative hypotheses.

Prediction	Statement	Mechanism	Alternative
H1 — Land-use main effects	Both land-use surrogates reduce richness, with the steepest declines in forest-specialist richness, intermediate declines in forest-facultative richness, and weak, inconsistent or positive responses in non-forest richness	Forest cover is the resource axis separating the guild strata. Conversion should therefore filter richness by forest dependence, with specialists losing habitat most directly, facultative species retaining partial occupancy across forest and modified cover, and non-forest species potentially gaining habitat where conversion creates open or matrix-associated conditions	Aggregate richness may conceal this gradient if disturbance-tolerant species offset losses among forest-dependent species. Land use may also filter all guilds similarly, in which case forest affiliation provides little additional explanatory leverage

Prediction	Statement	Mechanism	Alternative
H2 — Vegetational heterogeneity moderation	LU × EVICV and LU × EVIENT interactions are negative — high land use with high vegetational heterogeneity yields lower richness than the additive prediction	Vegetation texture is itself land-use-responsive. Degradation and fragmentation raise apparent compositional and configurational heterogeneity while removing the continuous cover forest-dependent species require, so high texture under high land use marks a degraded mosaic rather than usable niche diversity	Moderation may instead reflect productivity covarying with texture under cropland, or suppressor inflation where land-use–heterogeneity collinearity masks the interaction. Exposure to the second is asymmetric — EVIENT is moderately coupled to both land-use surrogates whereas EVICV is weakly coupled, so a null carries more suppression risk for EVIENT than for EVICV (Appendix 2.S2.3). The matched control set and product-term structure-coefficient diagnostics test both

Prediction	Statement	Mechanism	Alternative
H3 — Geophysical heterogeneity moderation	LU × TOPOD and LU × CLIMD interactions are positive, opposite in sign to the vegetational moderators — high land use with high geophysical heterogeneity is less depleted than land use alone	Terrain and climatic heterogeneity are persistent properties of the physical template, not land-use-responsive states, and sustain more persistence pathways under pressure through topographic refugia, reduced conversion of rugged ground, and a wider range of climatic conditions	Non-random land-use placement may confound ruggedness with forest retained on unconverted slopes. For climatic heterogeneity an area–heterogeneity tradeoff may instead reduce area per climatic condition, predicting a negative or non-monotonic interaction (Allouche et al. 2012). Matched controls and scale-coherence checks evaluate both

Prediction	Statement	Mechanism	Alternative
H4 — Energy moderation	LU × AET is positive if realised productivity buffers land-use damage. AET is the focal energy moderator; we use potential evapotranspiration and the Priestley–Taylor ratio only in proxy-swap robustness checks rather than as parallel moderation arms	Higher realised water–energy availability provides demographic subsidies that partly offset the richness cost of conversion. AET also conditions the heterogeneity interactions, so moderation surviving energy conditioning is not merely a productivity gradient	Land-use–energy coupling geometry may generate an apparent interaction through collinearity rather than resource-mediated buffering. The design carries no abundance mediator, so this is not a test of the more-individuals hypothesis, and proxy-swap survival is reported as an energy-conditioning robustness check

H5 — Forest-dependence scaling of the origin-dependent contrast	Heterogeneity moderation scales with forest dependence, and its direction depends on heterogeneity origin. Vegetational heterogeneity compounds land-use damage most strongly in forest-specialist richness, more weakly or inconsistently in forest-facultative richness, and weakly, not at all, or positively in non-forest richness. Geophysical heterogeneity buffers land-use damage most clearly in forest-specialist richness. The contrast between origins should therefore be sharpest where forest dependence	Forest specialists depend most directly on continuous forest structure, canopy buffering and refugial conditions. Land-use-generated vegetation texture marks subdivision, edge exposure and matrix contrast, so it should compound damage in proportion to forest dependence. Terrain and climatic heterogeneity predate conversion and can retain persistence pathways under land-use pressure, so they should buffer the same forest-dependent component	The guild gradient may be non-monotonic if land-use mosaics benefit facultative or non-forest species. Geophysical moderation may instead be negative under area–heterogeneity tradeoff dynamics, so direction rather than specialist concentration separates buffering from tradeoff interpretations (Allouche et al. 2012). Comparable vegetational moderation in non-forest richness would weaken the interpretation of vegetation texture as land-use-responsive disruption. Comparing full and restricted aggregate richness guards against artefacts introduced by excluding taxa without guild assignments

Prediction	Statement	Mechanism	Alternative
	is greatest and diminish where it is absent		

We evaluate all predictions across response grains of 10, 25 and 40 km and predictor extents from 10 to 100 km for both land-use surrogates. We interpret support by combining interaction direction, scale coherence and product-term structure-coefficient concordance. The central test contrasts land-use-responsive vegetation texture with persistent geophysical heterogeneity; the energy arm evaluates whether those contrasts persist after conditioning on realised water–energy availability.

Table 2.2. Focal land-use × heterogeneity and land-use × energy interactions across forest-affiliation strata. Primary support is the number of the 18 grain × extent combinations meeting the joint criterion of ΔAICc > 2 over the matched additive model together with product-term structure-coefficient concordance; for three-way terms it is the number of cells in which the triple is structure-coefficient-concordant. Conditioned support is the nested subset retaining support under the matched control set, or for three-way terms the number of cells in which the triple improves on the conditioned two-way by ΔAICc > 2. Bootstrap gives the coefficient with its 95% interval envelope and the number of cells whose interval excludes zero (B = 250); the model level differs by surface, is given in parentheses and is tabulated in Appendix 2.S5.5, and envelopes are unavailable for surfaces outside the load-bearing set. Directions are on the log scale. Coefficients are not pooled across land-use surrogates, moderators or strata. The full five-stratum support matrix is given in Appendix 2.S5.2, survival under the additive control ladder in Appendix 2.S5.3 and proxy-swap robustness in Appendix 2.S5.4; Appendix 2.S5.6 traces the load-bearing rows of this table to the scale combinations that produced them. The complete disposition across all ten tested surfaces, five strata and eighteen scale combinations is given in Figure 2.S5.1.

Stratum	Interaction	Primary (n/18)	Conditioned or incremental support	Bootstrap β [95% envelope] · cells excl. 0	Direction and scale domain	Status
Forest specialist	CROP × EVICV	15/18	15/15 controls-robust V0→V3; sign retained throughout	−0.084 [−0.157, −0.002] · 15/15 (L1)	Negative across the specialist scale surface; zero-crossing ordered by forest dependence (−0.080 → −0.101 under joint control at grain 10, extent 70 km)	Headline compounding
	CROP × EVICV × AET	18/18 SC-conc.	15/18 incremental over the conditioned two-way, at fine-to-mid extent	+0.071 [−0.010, +0.141] · 4/6 (L3)	Positive triple on the negative two-way; direction grid-wide, so the specialist cost weakens as realised energy rises	Headline energy conditioning

Stratum	Interaction	Primary (n/18)	Conditioned or incremental support	Bootstrap β [95% envelope] · cells excl. 0	Direction and scale domain	Status
	NNCOV × EVICV	9/18	not computed	−0.091 · 9/9, of which 8 negative and 1 positive (L1)	Negative at broad extent, positive at fine grain; SC-concordant at broad extent only	Scale-bounded corroboration
	NNCOV × EVICV × AET	18/18 SC-conc.	8/18 incremental over the conditioned two-way	+0.078 to +0.080 · 2/2 resampled cells (L3, L4)	Positive in all 18 cells (+0.024 to +0.136), but the two-way changes sign across the matrix, so the buffering reading holds only where that two-way is negative	Secondary energy check

Stratum	Interaction	Primary (n/18)	Conditioned or incremental support	Bootstrap β [95% envelope] · cells excl. 0	Direction and scale domain	Status
	NNCOV × AET	12/18	proxy-robust under PET and ALPHA (34/36, 32/36)	−0.082 [−0.186, +0.005] · 10/12 (L2c)	Negative; crossing ordered by forest dependence, mirrored by the non-forest control below	Secondary, sign-split
	NNCOV × EVIENT	3/18	not computed	+0.065 · 2/3, both positive (L1)	Positive where supported, against the negative direction predicted for vegetational moderation (H2); non-forest reverses sign	Diagnostic, not promoted

Stratum	Interaction	Primary (n/18)	Conditioned or incremental support	Bootstrap β [95% envelope] · cells excl. 0	Direction and scale domain	Status
	CROP × CLIMD	8/18	specialist-specific; SC weak (0.007–0.048)	+0.106 [+0.038, +0.176] · 8/8 eligible	Positive at fine-to-mid grains and extents, opposite in sign to EVICV	Secondary; supplement only
Non-forest control	NNCOV × AET	17/18	proxy-sensitive; not robust to PET or ALPHA	+0.089 [−0.014, +0.185] · 15/17 (L2c)	Positive; the mirror of the specialist sign on the same predictor	Secondary, sign-split

Surfaces not carried as focal forest-mediated moderation results are resolved in Appendix 2.S5 and Figure 2.S5.1. CROP × AET and CROP × TOPOD reached primary support only in the non-forest stratum, at 13 and 4 of 18 cells respectively, and in neither forest-specialist nor forest-facultative richness. NNCOV × TOPOD reached one supported cell, in full total richness. NNCOV × CLIMD reached five supported cells: two in full total richness, one in restricted total richness and two in forest-facultative richness, with neither surface supported in forest-specialist or non-forest richness. CROP × EVIENT reached a single supported specialist cell and was set aside on the post-hoc directional and non-forest specificity comparisons documented in Appendix 2.S3.2. The two aggregate strata are omitted here and are shown in full in Figure 2.S5.1.

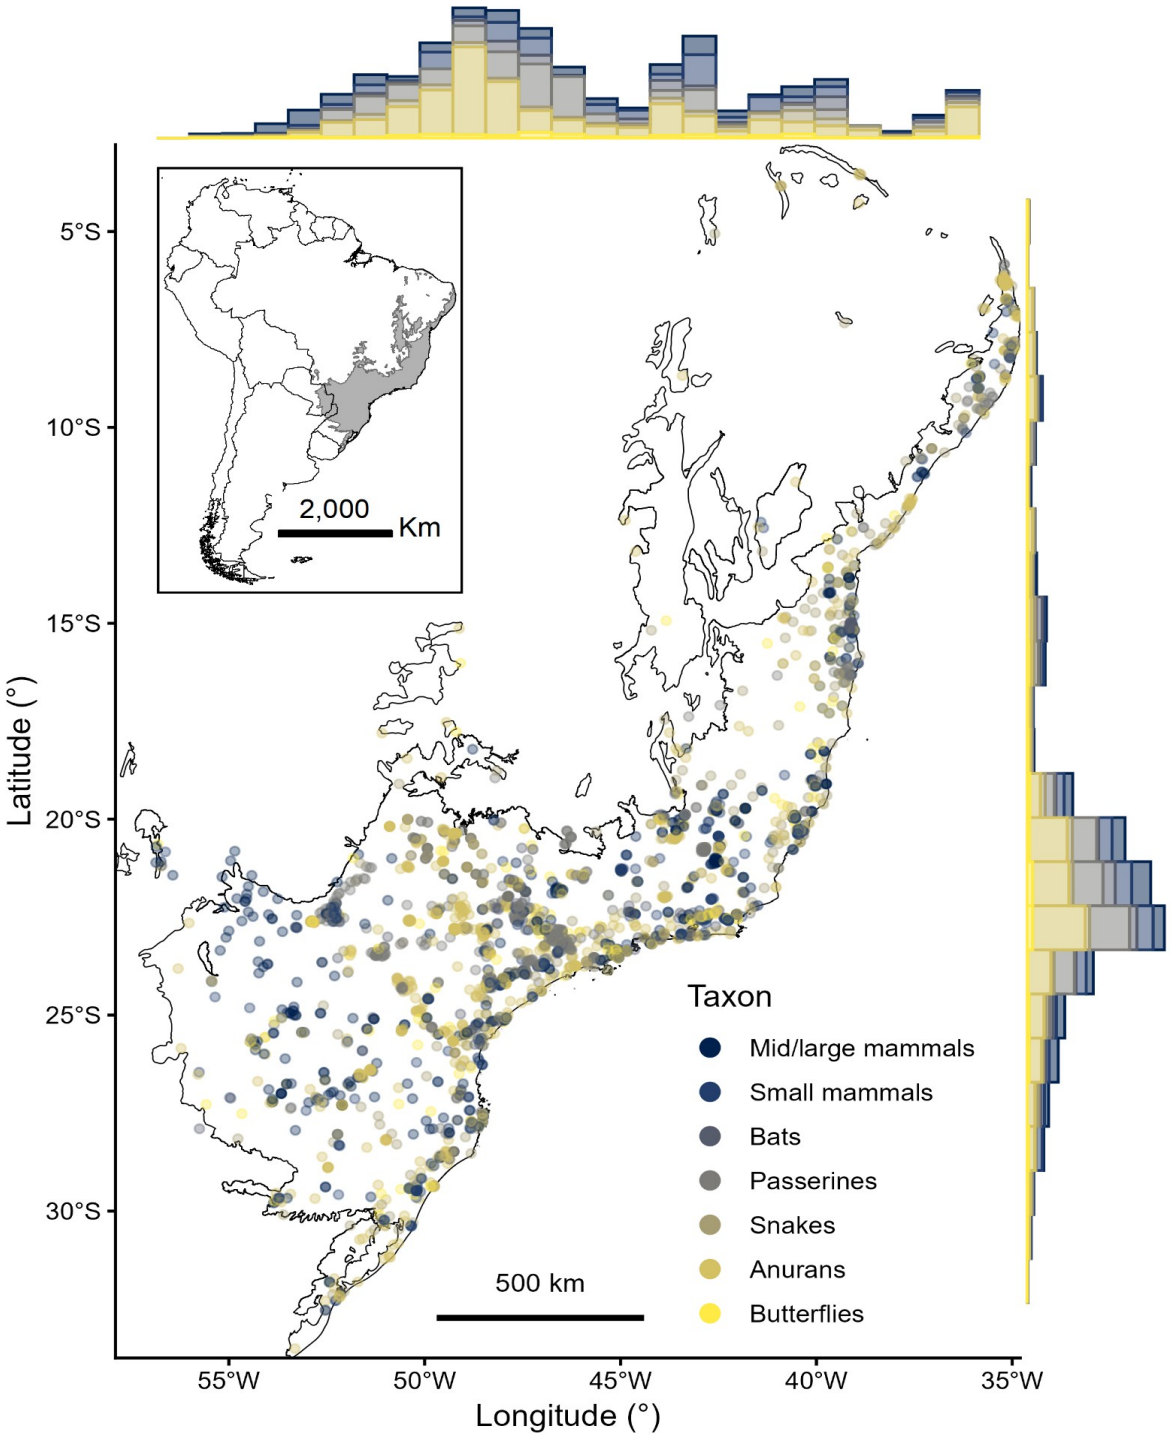


Figure 2.1. Study area and assemblages. Spatial distribution of the 2,986 curated, effort-reported species assemblages across the Atlantic Forest, prior to aggregation, coloured by animal group. Marginal histograms give the longitudinal (top) and latitudinal (right) distribution of assemblages, stacked by group. The inset locates the biome within South America, with the Atlantic Forest shaded. Scale bars are 500 km for the main panel and 2,000 km for the inset.

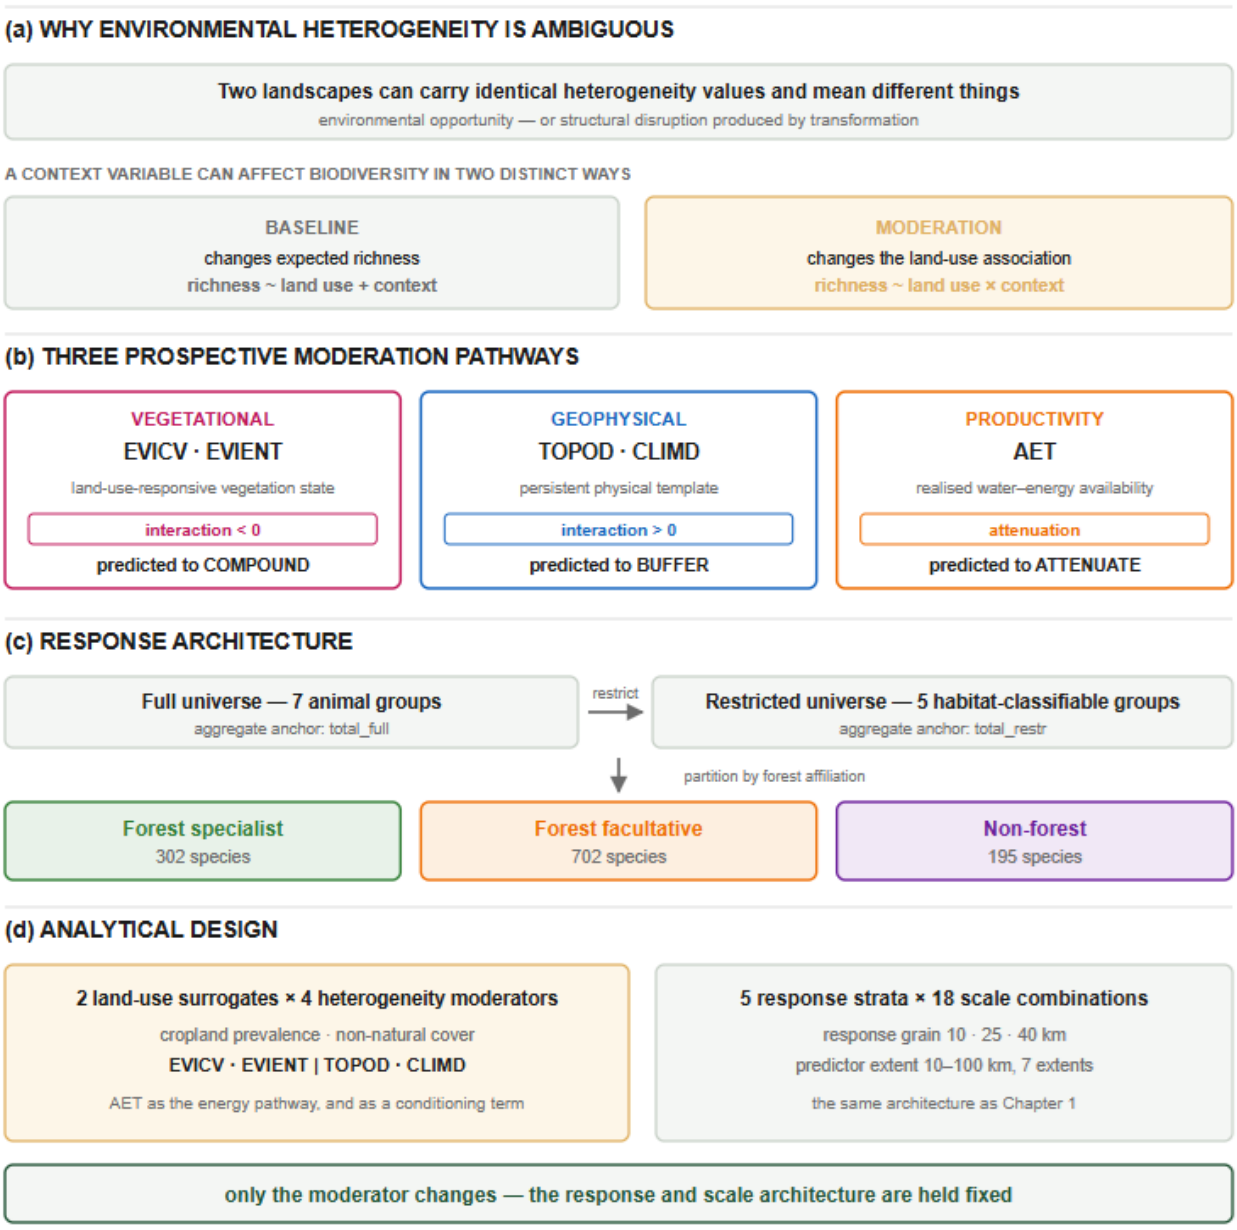


Figure 2.2. Why heterogeneity is ambiguous, prospective moderation pathways, response strata and analytical design. (a) Why environmental heterogeneity is ambiguous. Two landscapes can carry identical heterogeneity values and mean different things — environmental opportunity, or structural disruption produced by transformation. A context variable can affect biodiversity in two distinct, non-exclusive ways: it can change expected richness (a baseline effect, richness ~ land use + context) and it can change the land-use association (a moderation effect, richness ~ land use × context). This chapter tests the second. (b) Three prospective moderation pathways. Land-use-responsive vegetation texture (EVICV, EVIENT) is predicted to compound the cost of

conversion (interaction < 0); a persistent physical template (TOPOD, CLIMD) is predicted to buffer it (interaction > 0); realised water–energy availability (AET) is predicted to attenuate it, and additionally enters as a conditioner of the vegetational surface. (c) Response strata. The full and restricted aggregates, and the partition of restricted total richness into forest-specialist, forest-facultative and non-forest richness. Restricted total is a linked aggregate of the three guilds rather than an independent stratum; specialist and non-forest richness are distinct, non-overlapping components of that partition, so opposite fitted responses are not algebraically forced by a fixed total. (d) Analytical design. Two land-use surrogates crossed with four heterogeneity moderators across 18 grain × extent scale combinations, with AET as the energy pathway. The response and scale architecture are held fixed from Chapter 1; only the moderator changes. Species counts are from Appendix 0.S1, the scale design from the Methods, and the support rule from Appendix 2.S3.2.

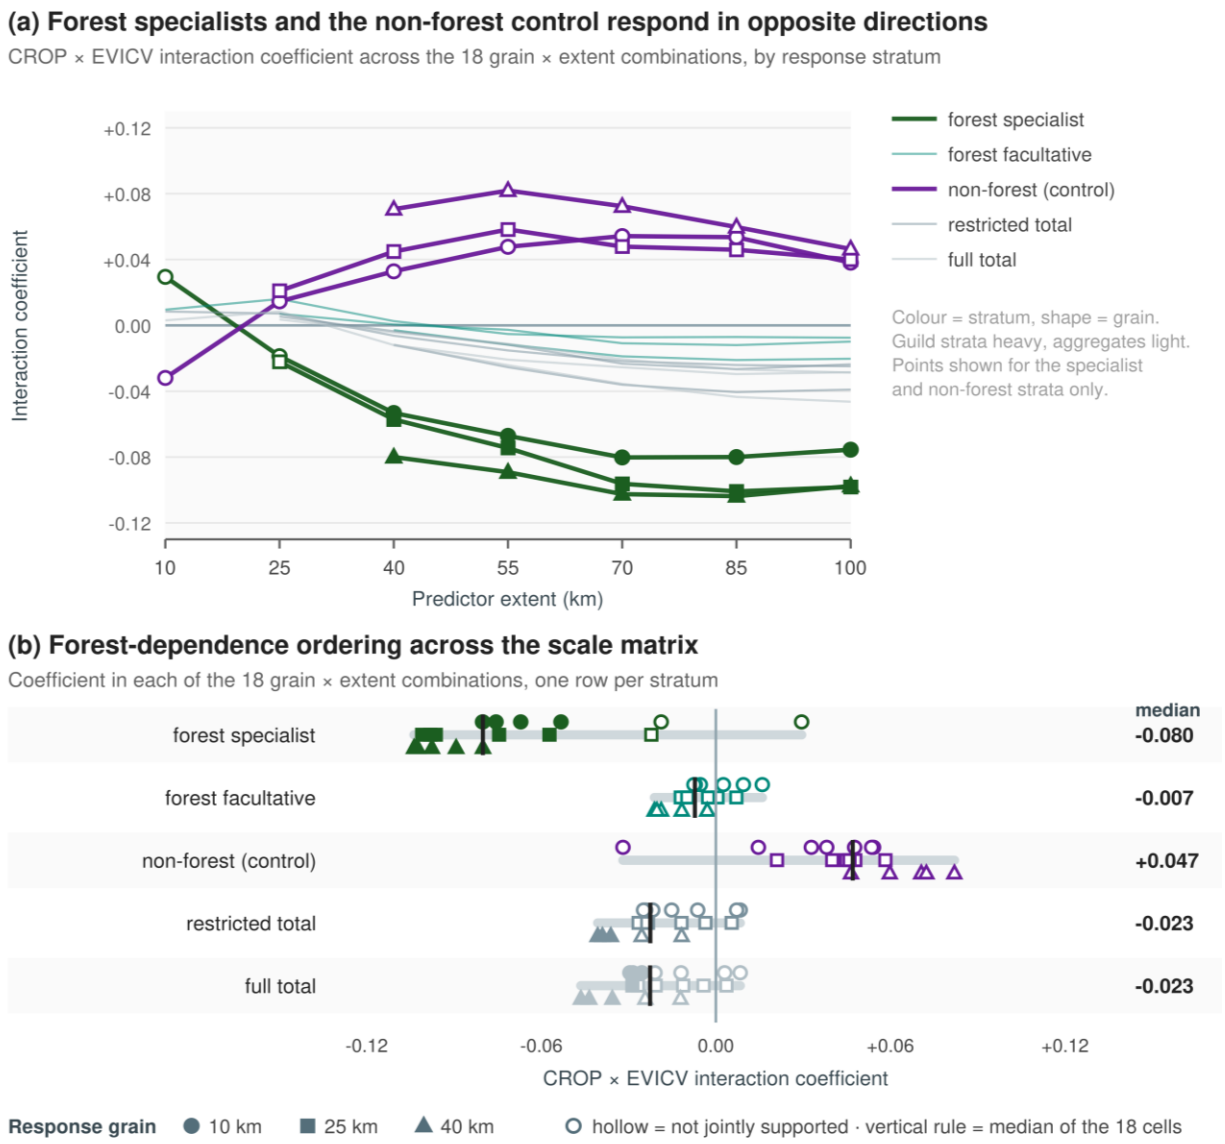


Figure 2.3. Cropland × greenness-magnitude compounding and its guild specificity. (a) The CROP × EVICV interaction coefficient across all eighteen grain × extent combinations, by response stratum: forest specialists are negative and the non-forest control positive over most of the matrix, while both aggregates sit near zero. Colour distinguishes stratum, symbol shape distinguishes response grain, and filled symbols meet the joint support rule. (b) We arrange the same coefficients by stratum and order the rows by forest dependence. Forest specialists are most negative, forest-facultative richness lies near zero, the non-forest control is positive, and the two aggregates are intermediate and close to one another. This ordering characterises the scale matrix rather than the medians alone. Forest specialists sit below the non-forest control in 17 of the 18 combinations, and the strict specialist < facultative < non-forest ordering holds in 16. The single

reversal is the finest combination, 10 km grain with 10 km extent, where the specialist coefficient is positive (+0.030) and the non-forest coefficient negative (−0.032) — the same cell in which the two-way surface runs positive in Figure 2.4. Median coefficients are −0.080 in forest specialists, −0.007 in forest facultative, +0.047 in the non-forest control and −0.023 in both aggregates; the specialist surface is jointly supported in 15 of 18 cells and the non-forest surface in none. Because the eighteen combinations are nested windows on one set of assemblages, their agreement measures scale coherence and not independent replication. Coefficients are from conditional Poisson generalised linear mixed models of untransformed species-richness counts, fitted with an observation-level random effect, and are on the log scale. From the deposited interaction-results table.

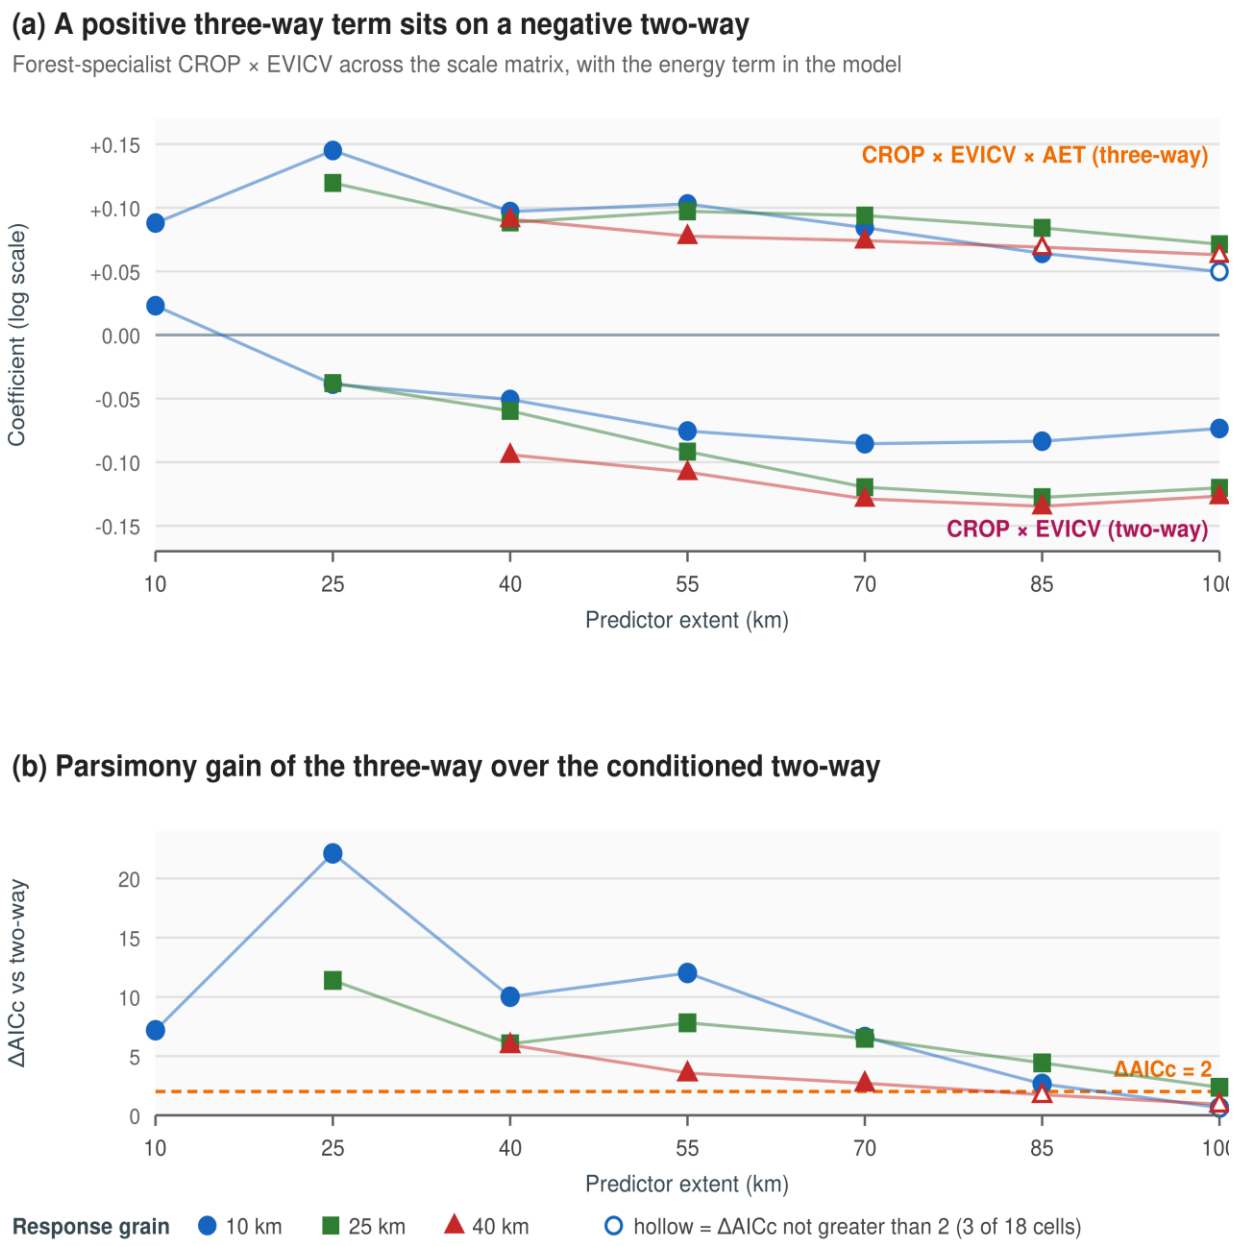


Figure 2.4. Energy conditioning of the specialist cropland × greenness-magnitude surface. (a) The two-way CROP × EVICV coefficient and the three-way CROP × EVICV × AET coefficient across predictor extent at each response grain, with the energy term in the model. A positive three-way term on a negative two-way means the specialist cost of cropland at high greenness variability is at full strength where realised evapotranspiration is low and weakens as it rises; the three-way term is positive and structure-coefficient-concordant in all 18 cells. (b) The parsimony gain of the three-way model over the conditioned two-way model. The three-way improves on

the conditioned two-way by ΔAICc greater than 2 in 15 of 18 cells, concentrated at fine-to-mid extent. Forest-specialist stratum only. The eighteen combinations overlap, so the run of supported cells across extent should be read as one coherent surface rather than as eighteen independent tests. Coefficients are from conditional Poisson generalised linear mixed models of untransformed species-richness counts, fitted with an observation-level random effect, and are on the log scale. From the deposited three-way interaction table.

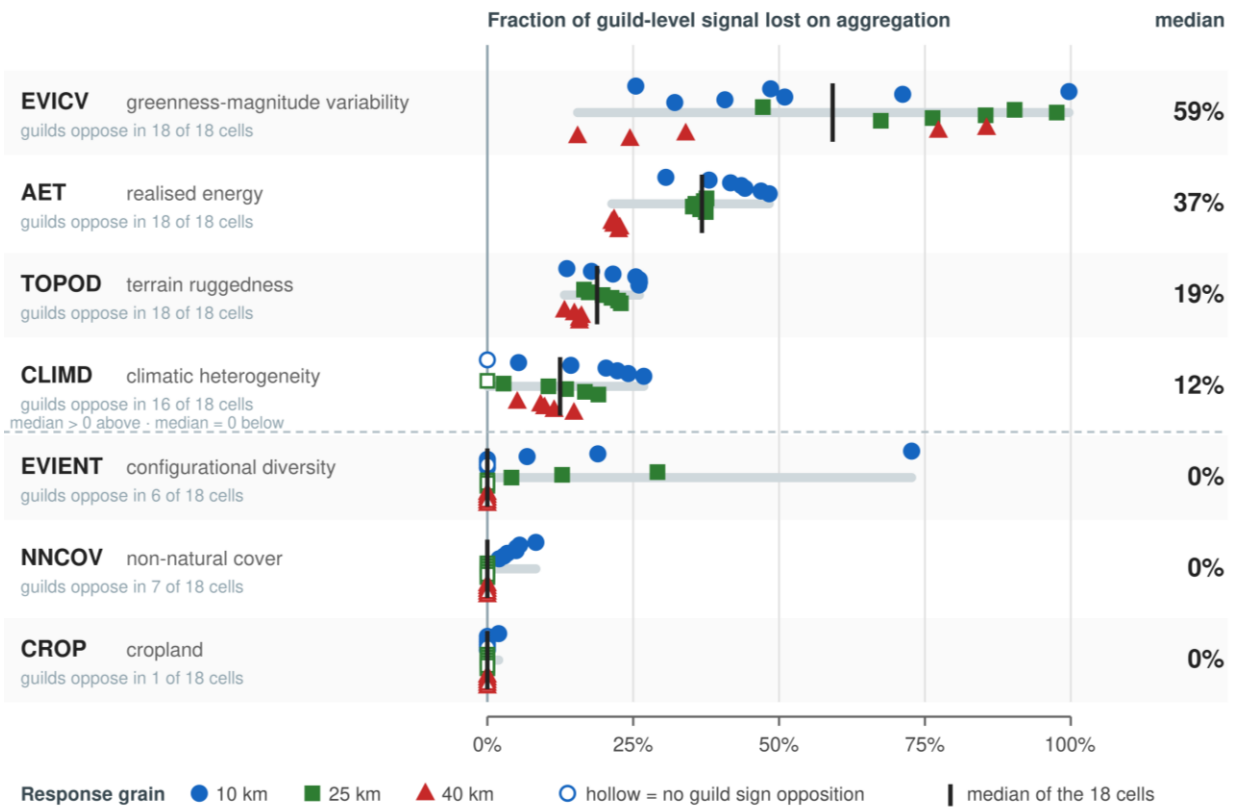


Figure 2.5. Process-specific cancellation under aggregation. For each process, the fraction of guild-level signal lost when the three forest-affiliation guilds are aggregated into restricted total richness, shown for all eighteen grain × extent combinations. The vertical rule marks the median across cells — the value reported in the Results — and the grey span the observed range. The dashed rule separates processes with a non-zero median above from those with a zero median below. Cancellation is exactly zero when all three guild contributions share sign, so hollow symbols denote cells with no guild sign opposition. Greenness-magnitude variability, realised energy and terrain ruggedness show non-zero cancellation in every cell, and their minima exceed the maxima of both land-use surrogates, which pass aggregation essentially intact. Configurational diversity is zero at the median but non-zero in six cells, declining monotonically with predictor extent within each grain. The spread within each row reflects the overlapping design of the eighteen windows rather than sampling error, so the range is not an interval estimate. Vertical position within a row separates overlapping points and carries no other meaning. The guild coefficients entering the decomposition are from conditional Poisson generalised linear mixed models of untransformed species-richness counts and are on the log scale. From the deposited aggregation-decomposition table.

CHAPTER 3

Habitat amount control is necessary but not sufficient: sign-symmetric diagnostics of fragmentation-per-se inference

1. Introduction

Land-use transformation affects biodiversity by reorganising the amount, condition, and spatial structure of native vegetation rather than by altering a single landscape property. For a given species, assemblage, or response variable, these transformations can reduce usable habitat, subdivide remaining cover, increase edge exposure, alter isolation and connectivity, and change the matrix conditions through which organisms experience the landscape (Fischer & Lindenmayer, 2007; Hadley & Betts, 2016). Broad land-use gradients can reveal where biodiversity changes, but they cannot by themselves determine whether the operative pathway is habitat loss, spatial subdivision, edge exposure, matrix transformation, vegetation reorganisation, or a response-specific mixture of these processes. Because these changes commonly occur together, fragmentation research needs a clear separation between habitat loss and fragmentation per se. Habitat loss refers to reduction in habitat amount, often represented by a land-cover proxy in landscape analyses, whereas fragmentation per se refers to spatial reconfiguration caused by the breaking apart of habitat independent of that loss (Fahrig, 2003; Hadley & Betts, 2016; Fletcher et al., 2018). Fragmentation-per-se inference therefore asks a narrower question than whether biodiversity differs among altered landscapes. It asks whether spatial configuration affects biodiversity after holding habitat amount comparable, an estimand that gained prominence once the separation of habitat loss from fragmentation per se became a major inferential goal (Ruffell et al., 2016; Fletcher et al., 2018). This estimand is conceptually clear but difficult to identify in observational landscapes because habitat loss commonly co-generates habitat amount and configuration before any model is fitted (Didham et al., 2012; Ruffell et al., 2016).

Whether habitat fragmentation affects biodiversity independently of habitat amount remains a central contested question in landscape ecology, and disagreement has been attributed to differences in spatial scale, study design, response variable, metric choice, ecological context,

and the statistical treatment of habitat amount rather than to a single empirical fact (Valente et al., 2023; Riva et al., 2024). The fragmentation debate has therefore remained unresolved partly because studies estimate different quantities under the same name and convert those differences into opposing general claims. Additive regression is the standard observational tool for isolating configuration effects (Smith et al., 2009; Fahrig, 2003), but its ecological interpretation rests on a condition stronger than model inclusion. The predictors must occupy distinct enough regions of the realised predictor space that a partial coefficient reads as movement along one axis rather than as a division of a shared landscape gradient between two model terms (Koper et al., 2007; Ruffell et al., 2016; Fletcher et al., 2023). Observational landscapes cannot guarantee this by design, because the analyst does not assign configuration independently of amount; the condition must be demonstrated in the sampled data rather than assumed from the model form (Lundberg et al., 2021).

Habitat fragmentation is a particularly acute version of this identification problem because land-use change structures both the ecological pattern to be explained and the predictors used to explain it. Land-use change can reduce cover, subdivide remaining habitat, increase edge exposure, and alter connectivity and matrix conditions through the same landscape history rather than through independently assigned treatments (Fischer & Lindenmayer, 2007; Fahrig, 2003; Didham et al., 2012; Fletcher et al., 2018). A configuration coefficient in this setting carries three claims that must be separated. The realised predictor space must contain an interpretable configuration contrast at comparable habitat amount, the separable component must have detectable statistical support, and the supported component must correspond to the ecological mechanism being invoked (Koper et al., 2007; Ruffell et al., 2016; Lundberg et al., 2021; Fletcher et al., 2023). Model inclusion alone secures none of these conditions. A configuration coefficient can therefore be non-diagnostic under any sign.

Additive models that include habitat amount and a configuration metric are therefore necessary but not sufficient. They estimate partial slopes, not automatically the ecological contrast required for fragmentation per se. A configuration coefficient approximates that contrast only when the realised predictor space contains same-amount, different-configuration comparisons, the response variable matches the biological process being interpreted, and the configuration metric represents the proposed mechanism (Dormann et al., 2013; Wang et al.,

2014; Fletcher et al., 2007; Martínez-Lanfranco, 2026a, 2026b). These requirements are often treated as implied by model specification, but they are empirical conditions that must be checked in the realised data (Martínez-Lanfranco, 2026a).

Why this condition commonly fails follows from how habitat-loss landscapes originate. Habitat loss can reduce total cover while subdividing remaining habitat into more and smaller patches, so habitat amount and patch number, a metric widely used to index patch subdivision and fragmentation (Fahrig et al., 2026), are often co-produced along one directed gradient before any landscape is sampled (Didham et al., 2012; Fahrig, 2003; Fletcher et al., 2018). Within a fixed sampling unit the coupling is partly an arithmetic identity. Mean patch size equals habitat area divided by patch number, so patch-count metrics are functionally tied to amount through the subdivision process itself (Fletcher et al., 2023). The resulting predictor space is bounded and asymmetric, and the amount–patch-number relationship can be nonlinear rather than straight-line (Andrén, 1994; Wang et al., 2014; Figure 3.1a–b). This matters for diagnosis because the standard collinearity checks, the correlation coefficient and the variance inflation factor, register only linear dependence; nonlinear coupling that survives after the linear part is removed is invisible to them (Dormann et al., 2013; Koper et al., 2007).

When predictors share an environmental gradient, a single partial coefficient no longer summarises a predictor's independent relationship with the response, and three quantities are needed to read the coefficient architecture. In an additive model of diversity on amount and a configuration metric, the structure coefficient (SC) measures how strongly a predictor aligns with the model's fitted response gradient, the beta weight measures its unique partial slope, and the semi-partial R^2 measures the unique variance it contributes (Ray-Mukherjee et al., 2014; Stoffel et al., 2021; Figure 3.1c). These can diverge in two opposite ways. A subordinate predictor can keep a large structure coefficient while its beta weight and unique variance collapse toward zero, so the coefficient is attenuated even though the predictor tracks the response — the signature of suppression; the mirror case is amplification, in which residual coupling inflates a partial slope above the predictor's weak standalone association (Smith et al., 2009; Prunier et al., 2015). Both detach the coefficient from the ecological contrast it is read as representing, and they do so under opposite coefficient signs.

Predictor geometry, the actual structure of the predictor space, is not the only way a coefficient can miss the estimand. The configuration metric is itself part of the identification problem. Patch number is a subdivision count that covaries with edge exposure, isolation, connectivity, patch size, and matrix contrast, so a patch-number coefficient can be statistically real while the mechanism attached to it stays unidentified, unless the result converges with metrics closer to the proposed pathway (Fletcher et al., 2007, 2018). The response matters in the same way. A configuration effect on one diversity component need not transfer to another, so a coefficient is interpretable only when the response measures the process being claimed.

This chapter asks what evidence is required before a configuration coefficient can be read as evidence about fragmentation per se, and answers it with four conditions that an estimate must satisfy. Two linked cases show why coefficient magnitude and sign alone are not diagnostic. In one case, the same global multi-taxa dataset produced opposed interpretations when a categorical fragmentation contrast was replaced by additive forest-cover and patch-number predictors (Gonçalves-Souza et al., 2025; Fahrig et al., 2026; hereafter GS25 and F26 respectively). The resulting near-zero patch-number coefficients were interpreted as evidence that fragmentation effects disappeared after habitat amount control. In the second case, a negative patch-number coefficient in Atlantic Forest tree total beta diversity was interpreted as evidence that fragmentation per se drives biotic homogenization (Kramer et al., 2026; hereafter K26). These signs differ, but both claims depend on whether a patch-number coefficient identifies the intended ecological contrast.

The two cases are diagnostic examples of these failures rather than an exhaustive survey, and they are deliberately sign-opposed — one a near-zero coefficient, one a negative coefficient — to show that coefficient sign alone does not settle what a configuration coefficient identifies. They are not a matched comparison: the responses, metrics and diagnostics differ between them because each case is examined against the conditions its own claim invokes, not against the other.

Figure 3.1 sets out the inferential procedure behind these conditions. Fragmentation-per-se inference first requires geometric separability, then statistical detectability, and only then ecological interpretation. I evaluate both cases using four conditions. Predictor-space adequacy asks whether the realised data contain the same-amount, different-configuration contrasts

required by the estimand, and therefore tests the geometric tier. Statistical detectability is the intervening coefficient-diagnostic tier, where suppression and amplification can distort the partial coefficient before any ecological reading. Response-estimand alignment, metric-mechanism correspondence, and domain stability then govern ecological interpretation. They ask whether the biodiversity response measures the ecological process being interpreted, whether the fitted configuration metric represents the proposed pathway rather than a broad proxy for several landscape changes, and whether the coefficient is stable across the ecological gradient where it is interpreted, while allowing that different fragmentation mechanisms may be detectable in different cover domains (Andrén, 1994; Fahrig, 2003; Fletcher et al., 2018; Fletcher et al., 2023).

The two cases reveal sign-symmetric failure. In the GS25–F26 richness case, patch number is suppressed toward a near-zero coefficient despite strong fitted-gradient alignment, a pattern that structure coefficients, beta weights, and semi-partial R^2 diagnose as suppressor attenuation (Ray-Mukherjee et al., 2014; Martínez-Lanfranco, 2026b). In the K26 beta-diversity case, patch number is amplified from weak marginal support into a significant negative partial coefficient, but the total-beta subdivision signal is narrow. It is directionally reproducible for one patch-count metric but not recovered by metrics closer to the mechanisms invoked to interpret it, and it does not persist under the authors' own alpha-corrected response diagnostic or across independent predictor extents. The coefficient sign and the failure mode are orthogonal. A null coefficient can be non-diagnostic, and a negative coefficient can be non-diagnostic, but not for the same reason. A glossary of recurrent conceptual and technical terms is provided in Appendix 3.S1.

2. Methods

2.1 Comparative diagnostic design

Fragmentation per se targets the expected difference in biodiversity between landscapes with the same habitat amount but different spatial configuration. An additive model coefficient for a configuration predictor approximates that estimand only when the realised data contain sufficient same-amount, different-configuration contrasts, the response variable measures the biological process to which the configuration mechanism is attributed, and the configuration metric represents the hypothesised ecological pathway. The distinction between estimand and

coefficient matters because these conditions are routinely assumed from model specification rather than verified from the data (Lundberg et al., 2021). A near-zero configuration coefficient can mean the ecological effect is absent, or it can mean the predictor geometry has suppressed a configuration-associated signal. A negative configuration coefficient can indicate an independent configuration effect, or it can reflect a partial-slope artefact driven by the same shared gradient that makes predictor separability difficult. Sign alone does not settle the question. Even where geometry permits a clean estimate, a partial coefficient is a direct conditional association, not a total effect. When predictors are causally coupled, part of one predictor's influence runs through the other, so the direct coefficient can understate the total ecological effect (Ruffell et al., 2016).

This design-level argument was developed in the GS25–F26 Conservation Biology comment (Martínez-Lanfranco, 2026a) and evaluated empirically in the expanded GS25–F26 reanalysis preprint (Martínez-Lanfranco, 2026b), and is generalised here as a formal estimand–coefficient mismatch across two sign-opposed cases. The first case involves a near-zero patch-number coefficient in a global multi-taxa richness dataset. The second involves a negative patch-number coefficient in Atlantic Forest tree beta diversity. Both datasets use subdivision metrics from the same patch-number family, both pass conventional linear collinearity screens, and both were used to support inferences about independent configuration effects. The diagnostic goal is not to estimate one universal fragmentation effect. It is to test whether the configuration coefficient in each case corresponds to the ecological estimand it was used to represent.

2.2 Case-study datasets and response variables

In both case studies, forest cover is treated as the habitat-amount proxy used in the source analyses. This does not imply that all forest cover is suitable habitat for all taxa or response groups. Habitat is species- or response-specific, and generic land-cover classifications can misrepresent the habitat actually available to organisms (Betts et al., 2014; Hadley & Betts, 2016). The diagnostic question is therefore conditional on the land-cover proxy available in each source analysis: whether the fitted configuration coefficient identifies fragmentation per se after controlling for that proxy.

The first case uses the global multi-taxa forest dataset assembled by GS25 and reanalysed by F26. The dataset contains 37 independent paired study comparisons and 74 landscapes. GS25

used a categorical paired contrast and reported lower local and landscape-scale diversity in fragmented landscapes. F26 replaced the categorical design with two continuous predictors measured within a 2-km buffer — forest cover percentage and number of patches — and reported near-zero patch-number coefficients across diversity components, concluding that no independent fragmentation signal remained after habitat amount control. The same underlying observations therefore produce opposite conclusions under alternative predictor representations, making the dataset analytically informative in a way that disagreements between separate datasets cannot be. The diagnostic targets the source of that divergence: whether the F26 near-zero coefficient is an ecological null or a consequence of the realised amount–configuration geometry. Analyses used the full global dataset, a high-cover subset requiring both landscapes within each study pair to exceed 50% cover (17 paired studies, 34 landscapes), and a Tropical South American subset (23 studies) as the most ecologically relevant regional domain.

The second case uses the Atlantic Forest tree beta-diversity dataset of K26. The dataset spans 95 Atlantic Forest landscapes across the full forest-cover gradient. Total beta diversity and Raup-Crick beta diversity were used as response variables. The predictor set included forest cover, number of forest patches, mean spatial distance among plant inventories, mean nearest-neighbour distance among patches (NND), and four bioclimatic variables. K26 reported that forest cover and number of patches were both significant negative predictors of total beta diversity, and interpreted the patch-number coefficient as evidence for an independent fragmentation-per-se effect on compositional dissimilarity. The diagnostic targets that specific interpretation, not the descriptive beta-diversity patterns or the spatial-distance and endemism results, which are not in dispute. Total beta diversity was treated as the primary K26 response because this is the response to which the fragmentation-per-se interpretation was attributed. Raup-Crick beta diversity was treated as an internal response diagnostic, not as a replacement response, because it asks whether the same patch-number signal persists after reducing the alpha-diversity component of total beta diversity (Chase et al., 2011).

2.3 Diagnostic conditions

Predictor-space adequacy. The fragmentation-per-se estimand requires landscapes that differ in spatial configuration while holding habitat amount approximately constant. I evaluated whether the realised predictor space contained enough such contrasts by examining the joint

habitat amount–patch number distributions, directional co-occurrence of the two predictors across study pairs or landscape gradients, the availability of high-cover, high-patch-count landscape combinations, and bidirectional generalised additive model (GAM) fits to detect nonlinear coupling that persists after the linear association is removed (Pedersen et al., 2019). Separate GAMs were fitted in both axis orientations — patch number as a function of cover, and cover as a function of patch number — to assess directional asymmetry in nonlinear coupling. For the GS25–F26 paired-study geometry, these fits were grouped by study. Standard linear collinearity metrics (Pearson r, variance inflation factor) were also reported, but are treated as necessary rather than sufficient screens for predictor separability.

Response-estimand alignment. A configuration coefficient is interpretable only with respect to the biodiversity response being modelled. Alpha richness, gamma richness, total beta diversity, alpha-corrected beta diversity, and endemism measure different ecological quantities through different mechanisms, and a coefficient supported for one response need not hold for another. I compared coefficient direction and statistical support across ecologically distinct response formulations to test whether a configuration signal was response-specific or generalised. In the GS25–F26 case, I compared the suppressor signature in richness components with its behaviour in beta diversity. In the K26 case, I compared total beta diversity with Raup-Crick beta diversity — the authors' own internal control for alpha-diversity effects — and with the endemism model as a response-specific contrast.

Metric-mechanism correspondence. Patch number and patch density measure one operationalisation of subdivision — the spatial disaggregation of remaining habitat into discrete units. Patch density was treated only as an area-standardised subdivision sensitivity, not as independent mechanism-matched evidence: at a fixed landscape extent, patch density is patch number divided by area, so it does not provide a distinct ecological signal from patch number in the same model. I therefore treated metric substitution as a mechanism test rather than a generic robustness check. In the GS25–F26 case, the diagnostic target was the behaviour of patch number as a subdivision metric embedded in the shared habitat-loss gradient. In the K26 case, the main diagnostic retained one representative metric per mechanism family — isolation, edge exposure, interior depth, connectivity, amount-decoupled arrangement, and core heterogeneity — to avoid treating redundant metrics as independent evidence. The full 19-metric audit was

used as a screening step, with the focal mechanism-congruent subset carried into the main diagnostic and the broader screen retained as a record of that selection. A fragmentation-per-se interpretation requires convergence across these metric families, not support from one patch-count operationalisation alone. For the K26 case, I additionally extracted all predictors at five circular-buffer extents: 25, 50, 100, 200, and 400 km² (radii 2,821 to 11,284 m), to test whether the patch-number total-beta association and the mechanism-metric null persisted across operational scales. Responses (total beta, Raup-Crick) and mean spatial distance were fixed from K26's published dataset. Forest cover, patch number, NND, and all mechanism-family metrics were re-extracted independently at each extent. This predictor-extent sensitivity is distinct from a spatial-resolution comparison: the response variables and landscape-level biodiversity data are unchanged, and only the operational extent of the predictor buffers varies.

Domain stability. Domain checks evaluate whether a coefficient is stable across the ecological gradient where it is interpreted, not whether fragmentation effects must always be strongest in one cover domain. Different mechanisms predict different domains of detectability. Threshold frameworks predict that configuration-associated responses may intensify under habitat scarcity, whereas extinction-filtering and specialist-persistence mechanisms predict that configuration effects may be most recoverable at intermediate-to-high cover where configuration-sensitive species remain present. I therefore treated cover-stratified analyses as coefficient-stability and mechanism-localisation diagnostics. A low-cover signal is not dismissed as artefactual simply because it occurs under habitat scarcity; rather, its fragmentation-per-se interpretation depends on whether it survives response-estimand, metric-mechanism, and domain-stability checks. These domain checks do not estimate fragmentation effects; they test whether the reported coefficient architecture persists in the ecological domain required by the proposed mechanism.

2.4 Coefficient architecture as diagnostic readout

Coefficient architecture describes the joint relationship among marginal associations, partial slopes, fitted-gradient alignment, and unique variance contributions. It is the diagnostic output of the four conditions, not a fifth inferential condition independent of them.

Coefficient architecture was evaluated using three complementary statistics. Structure coefficients (SC) quantify the alignment between a predictor and the fitted response irrespective of collinearity, standardised beta weights (BW) estimate the predictor's unique conditional contribution after accounting for correlated predictors, and semi-partial R^2 (R^2sp) quantifies the proportion of response variance uniquely attributable to that predictor. Together, these statistics distinguish fitted-gradient alignment from unique explanatory support and permit diagnosis of suppressor effects, coefficient attenuation, and coefficient amplification.

Two opposite pathologies arise when predictors share a habitat-loss gradient that prevents independent variance attribution. In suppressor attenuation, a predictor tracks the model-fitted response closely — large structure coefficient (SC; the correlation between predictor and model-fitted response) — but contributes almost no independent explanatory variance once the other predictor is included: near-zero standardised beta weight (BW) and near-zero semi-partial R^2 (R^2sp). The predictor is suppressed toward a near-zero coefficient even though its fitted-gradient alignment is strong. In partial-slope amplification, a predictor with negligible marginal association acquires a significant partial coefficient in the joint model, because conditioning on a correlated predictor reallocates shared gradient variance in a direction that inflates rather than deflates the coefficient. Both pathologies arise from the same geometric constraint; they produce opposite relationships between marginal and partial estimates. Neither is ordinary collinearity. Under simple redundancy the structure coefficient and beta weight shrink together as shared variance is absorbed, whereas suppression preserves the structure coefficient while the beta weight collapses, and amplification drives the two apart instead. The mismatch between fitted-gradient alignment and unique contribution, not the correlation itself, is the diagnostic signature.

Structure coefficients, standardised beta weights, and semi-partial R^2 (Ray-Mukherjee et al., 2014) were used to characterise coefficient architecture in the GS25–F26 case. The suppressor tripod (large |SC|, near-zero BW, and near-zero R^2sp) identifies attenuation (Smith et al., 2009). In the K26 case, coefficient architecture was assessed through the marginal-to-partial shift for patch number. A marginal correlation near zero that becomes a significant partial coefficient identifies amplification. Neither pattern, by itself, resolves whether a fragmentation effect is ecologically real. Both are warnings that the coefficient may not estimate the intended estimand.

The GS25–F26 analysis additionally used a four-step residualisation sequence (M1–M4, defined in Section 2.4) to establish that coefficient allocation between the two predictors shifts systematically as the shared gradient is reallocated, rather than reflecting a stable ecological quantity. M3, which residualises forest cover on patch number, is the diagnostic contrast. M3 establishes the direction of suppression and rules out interpreting the additive null as stable across parameterisations. It does not estimate the magnitude of a fragmentation-per-se effect, because residualisation reduces but does not eliminate the shared gradient, and because the same-amount, different-configuration contrast the estimand requires is geometrically foreclosed in this dataset regardless of reparameterisation (Koper et al., 2007).

2.5 Analytical implementation and provenance

The Case Study 1 analyses (GS25–F26) draw on Martínez-Lanfranco (2026a, 2026b). I treated the F26 additive specification as the baseline for testing whether the near-zero patch-number coefficients represented ecological nulls or suppressor-attenuated configuration signals. The case-specific workflow characterised the realised forest cover–patch-number geometry, evaluated coefficient architecture using the diagnostics defined in Section 2.4, and used the M1–M4 residualisation sequence to test whether coefficient allocation changed as the shared habitat-loss gradient was reallocated between predictors. Logistic classification and causal-hierarchy checks evaluated, respectively, whether the continuous redesign reconstructed GS25's original categorical contrast and whether the coefficient architecture was consistent with a directed ordering of habitat amount and configuration. High-cover and Tropical South American subsets tested whether the diagnosis persisted across ecologically relevant domains.

The Case Study 2 analyses (K26) draw on Martínez-Lanfranco (2026d). I treated the reported negative patch-number coefficient in total beta diversity as the focal configuration claim and evaluated whether it identified the mechanism attributed to it. The case-specific workflow characterised the forest cover–patch-number geometry, contrasted total beta diversity with Raup-Crick beta diversity as a response-estimand check, tested whether the patch-number signal converged with metrics closer to the proposed mechanisms, independently extracted the landscape predictors across five circular-buffer extents (25, 50, 100, 200, and 400 km²), and evaluated stability across the cover domain. As a predictor-geometry sensitivity rather than a primary diagnostic, I residualised number of patches against forest cover and reciprocally

residualised forest cover against number of patches to test whether the negative coefficient persisted after reducing the forest-cover-predictable component of patch number.

The chapter-level synthesis comes from applying the same four diagnostic conditions across both cases. The source analyses supply the case-specific diagnostics, whereas this chapter uses them to generalise the estimand–coefficient mismatch across sign-opposed configuration coefficients. Full scripts, processed outputs, and session information are deposited in the Zenodo archives accompanying these analyses (Martínez-Lanfranco, 2026a, 10.5281/zenodo.19546548; 2026b, 10.5281/zenodo.19546467; 2026d, 10.5281/zenodo.20635343). The diagnostics, tables, and figures in this chapter derive from those archives.

3. Results

3.1 Case Study 1. Predictor-separability failure in a global multi-taxa dataset

3.1.1 The GS25–F26 discrepancy as an inferential case study

Two analyses of the same global multi-taxa dataset reached opposite conclusions about habitat fragmentation. GS25 assembled 37 independent study pairs of continuous and fragmented forest landscapes covering multiple taxa across tropical and temperate biomes, and reported lower local (α) and landscape-scale (γ) diversity in fragmented landscapes using a categorical design contrast. Higher compositional dissimilarity (β diversity) in fragmented landscapes did not rescue richness at the landscape scale; apparent increases in β diversity disappeared once spatial distance between plots was accounted for, making distance decay rather than compensatory turnover the central beta-diversity signal. F26 reanalysed the same 37 study pairs by replacing GS25's categorical contrast with two continuous predictors, forest cover percentage and number of patches, both measured within a 2-km buffer, and obtained near-zero patch-number coefficients across all diversity components, concluding that no independent evidence for fragmentation remained once habitat amount was controlled.

When the same observations produce opposite conclusions under alternative predictor representations, the discrepancy is analytically informative in a way that disagreements between different datasets cannot be. Study systems, taxa, spatial locations, and raw observations are held

constant; only the model architecture and predictor definitions differ. The source of divergence must therefore lie in how the two analytical frameworks allocate variance between habitat amount and configuration, making this case a direct empirical test of the estimand–coefficient gap: whether an additive coefficient corresponds to the ecological contrast it is used to represent depends on the realised predictor geometry, not on model specification alone.

3.1.2 Predictor-space adequacy: design geometry before biodiversity models

The dataset comprises 74 landscapes from 37 independent studies, each contributing exactly one continuous and one fragmented landscape. At the 2-km grain, the habitat-amount distribution is right-skewed and concentrated at high cover: the median cover across all 74 landscapes is 69.9% (mean 66.9%, SD 26.0%), with 71.6% exceeding 50%. Patch number is also right-skewed (median 30, mean 48.2, SD 56.1), reflecting structural heterogeneity concentrated at intermediate cover values where the nonlinear amount–configuration relationship peaks.

The realised predictor space does not span an independent two-dimensional surface. In 35 of 37 study pairs (94.6%), the fragmented landscape had lower forest cover than its continuous counterpart. In 31 of 37 pairs (83.8%), the fragmented landscape had more patches. Both conditions held simultaneously in 31 of 37 pairs (83.8%), with a median difference in forest cover of 28.4 percentage points (continuous minus fragmented) and a median difference in patch count of 44.5 patches (fragmented minus continuous). The predictor space therefore aligns strongly with the original categorical design contrast rather than recovering an independent configuration gradient. The realised predictor space is shown in Figure 3.2.

Same-amount, different-configuration contrasts are structurally scarce in this dataset. Only 3 of 37 pairs (8.1%) had an absolute habitat-amount difference of 5 percentage points or less between landscape types, and only 2 of 37 pairs (5.4%) simultaneously met that constraint and had both landscapes in the highest habitat-amount class (77–100% cover). This scarcity is not a random sampling outcome. The structural size asymmetry underlying the paired design explains it mechanistically: the median ratio of the continuous forest area to the smallest fragment within each study pair was 1,672× (range 10× to 125,000×; 26 of 28 studies with numeric areas for both landscape types exceeded 100×). Continuous landscapes are structurally distinct from fragmented ones not only in configuration but in absolute size, making same-

amount, different-configuration contrasts geometrically constrained by the landscape-generation process that produced the study pairs.

Standard linear collinearity diagnostics do not detect this geometric constraint. The Pearson correlation between forest cover and patch number is $r = -0.674$, below the conventional $|r| < 0.70$ screening threshold, and the variance inflation factor is VIF = 1.83, well below the standard threshold of 5. Both metrics indicate that linear collinearity is not severe. GAMs fitted in both axis orientations reveal substantial nonlinear coupling that persists after the linear association is removed. In the forward orientation (patches as response, cover as predictor), the GAM R^2 is 0.486 against a linear R^2 of 0.454, yielding a nonlinear component of 0.032. In the reverse orientation (cover as response, patches as predictor), the GAM R^2 is 0.702 against the same linear R^2 of 0.454, yielding a nonlinear component of 0.248 — a 7.75-fold asymmetry in the nonlinear coupling direction. This asymmetry is consistent with the directed structure expected from habitat loss, where cover loss constrains patch-number geometry more strongly than the reverse. A predictor pair that satisfies both conventional collinearity thresholds can therefore remain embedded in a shared habitat-loss gradient with substantial residual nonlinear dependence that standard diagnostics cannot flag.

A logistic mixed-effects model of GS25's categorical landscape classification confirms that the continuous redesign reconstructs rather than dissolves the original design contrast. Forest cover and patch number together explain nearly half of the marginal variance in landscape type, and patch number is strongly aligned with that classification gradient while contributing little unique variance once cover is included. The identification problem is therefore already present in the predictor space before any biodiversity response is introduced.

3.1.3 Coefficient architecture: suppressor attenuation in the richness response

Applying the additive F26 specification (M1: diversity ~ forest cover + patch number + random intercept by study) across 18 analytical variants — three diversity metrics (observed richness, Hill number $q = 0$, Hill number $q = 2$) crossed with two pairing designs (all pairs, close pairs) for each of three diversity components (α, β, γ) — produces the diagnostic signature of cross-over suppressor attenuation in α and γ diversity. Patch number has large, consistently negative structure coefficients (α diversity: median SC = −0.725, range −0.91 to −0.61; γ

diversity: median SC = −0.747, range −0.92 to −0.66), indicating strong alignment with the model-fitted richness gradient across all variants. Yet patch-number beta weights are near zero across all 18 variants (α median BW = −0.014; γ median BW = −0.020) and the semi-partial R^2 is negligible (α median R^2sp = 0.006%; γ median R^2sp = 0.02%). The ratio of inclusive to unique variance contribution for patch number is approximately 8,800-fold for α diversity and 3,000-fold for γ diversity: patch number aligns strongly with the fitted richness gradient but contributes almost no independent explanatory variance once cover is included. The full coefficient architecture across all 18 variants is shown in Figure 3.3.

The suppressor structure arises because forest cover dominates the shared habitat-loss gradient. It has large positive structure coefficients (α median SC = +0.997; γ median SC = +0.992) and substantial positive beta weights (α median BW = +0.108; γ median BW = +0.108). The algebraic basis is directly derivable from the OLS partial-regression formula: because cover and patch number are negatively correlated (r = −0.674), cover's inclusion adds a positive contamination term to the patch-number coefficient, systematically pushing it toward zero regardless of the true ecological magnitude of the configuration effect. The tight confidence intervals around near-zero M1 beta weights across all 18 variants reflect the stability of this geometric constraint, not ecological precision about a true null.

The residualisation sequence M1 → M2 → M3 → M4 confirms that coefficient allocation shifts systematically as the shared gradient is reallocated between predictors, rather than reflecting a stable ecological quantity. Residualising patch number on cover (M2) allocates the shared gradient primarily to the cover axis; M2 retains substantial residual nonlinear coupling (GAM R^2 = 0.44 in both orientations) and does not satisfy the orthogonality criterion. Residualising cover on patch number (M3) reverses the allocation: the cover coefficient attenuates (α median BW = +0.041) and patch-number coefficients shift consistently negative for both α diversity (median BW = −0.114, 6/6 variants negative) and γ diversity (median BW = −0.101, 6/6 negative). M3 is the only parameterisation achieving near-zero residual coupling in both axis orientations (GAM $R^2 \approx 0$ in both directions), satisfying the dataset-specific orthogonality criterion. The negative-control parameterisation (M4, bidirectional residualisation) reintroduces coupling (GAM R^2 = 0.63) and produces the predicted mirror image of M1, confirming that M1's gradient allocation to cover and M4's allocation to patch number are two

sides of the same geometric coin. The 144-fold increase in patch-number semi-partial R^2 from M1 to M3 (from 0.006% to 0.86% for α diversity) shows that the configuration-associated signal is present but geometrically unavailable to the partial coefficient under the full-dataset shared gradient.

A path model fitted with the directed causal hierarchy Amount → Fragmentation → Diversity supports the same interpretation. The amount-to-fragmentation path is strongly negative, the directed specification is strongly preferred over a correlated-errors alternative, and the direct fragmentation-to-diversity path remains near zero in the full dataset. The path model is therefore the structural-equation analogue of the near-zero M1 beta weight: both indicate geometric suppression from the shared habitat-loss gradient rather than an absent configuration-associated signal.

Beta diversity shows no equivalent suppressor structure under M1. Structure coefficients for patch number in β diversity range widely, including positive values, and beta weights remain near zero under both M1 and M3 without the systematic sign recovery observed in richness responses. This response-dimension specificity is consistent with β diversity responding to processes — spatial distance among plots, distance decay of similarity, habitat heterogeneity — that are not reducible to the amount–configuration partial-regression geometry producing suppressor attenuation in the richness components. The suppressor diagnosis does not apply to a response variable that GS25 themselves showed was primarily structured by spatial distance rather than the landscape-type contrast. This distinction carries forward to Case Study 2, which uses beta diversity as the response dimension in which a negative patch-number coefficient is reported.

3.1.4 Domain stability: high-cover robustness of the suppressor diagnosis

The suppressor diagnosis is not an artefact of the low-cover, severely deforested tail of the habitat gradient. Restricting the analysis to a study-level high-cover subset — requiring both the continuous and fragmented landscape within each study pair to exceed 50% cover — retains 17 studies (34 landscapes; Tier 1, full suppressor tripod applicable). The paired structure of the original GS25 design is preserved within each retained study. Linear coupling strengthens in this subset ($R^2 = 0.634$ vs. 0.454 in the full dataset), while the reverse nonlinear component increases

to 0.259 — a 16.2-fold asymmetry compared with 7.75-fold in the full dataset. The geometric coupling is therefore more pronounced, not less, at high cover: the upper portion of the cover gradient is exactly where progressive subdivision generates the steepest nonlinear increase in patch number relative to habitat area.

The suppressor signature persists intact within this subset. Under M1, patch number retains large negative structure coefficients (α median SC = −0.951; γ median SC = −0.953) alongside consistently negative beta weights (α median BW = −0.077, 6/6 variants negative; γ median BW = −0.060, 6/6 negative). Under M3, which satisfies the orthogonality criterion within the subset, patch-number coefficients remain negative across all variants (α median BW = −0.122; γ median BW = −0.099), marginally more negative than in the full dataset. The path model for the ≥50% cover subset produces a more negative direct Fragmentation → Diversity path (α = −0.22, 6/6 variants negative; γ = −0.17, 6/6 negative) than the full dataset. These patterns are consistent with the configuration-associated signal being more recoverable in the portion of the landscape gradient where forest-dependent taxa are most likely to persist and where configuration variation is real rather than conflated with overall landscape degradation.

3.1.5 Tropical South America: ecological conditionality of the diagnostic result

The Tropical South American subset provides the most relevant regional domain for the GS25–F26 diagnostic. The subset comprises 23 studies and satisfies the within-subset residual-uncoupling criterion under M3. Under the additive specification M1, the suppressor signature persists: patch number retains a large negative structure coefficient for α diversity (SC = −0.869, 6/6 variants negative) while its beta weight remains close to zero (BW = −0.049), reproducing the full-dataset suppressor tripod within the tropical regional component alone. Under M3, patch-number coefficients are negative across all six analytical variants for both α (median BW = −0.113) and γ diversity (median BW = −0.143). The tropical-only check excluding five temperate studies (n = 32) reproduces the same additive null (M1 median α BW = −0.002) and the same diagnostic recovery under M3 (6/6 variants negative), indicating that the pattern is not driven by the temperate minority of the global dataset.

More diagnostically informative than the coefficient direction is the behaviour of the habitat amount predictor under M3 in Tropical SA specifically. Under the additive baseline M1,

amount has a large positive structure coefficient (SC = +0.950 for α diversity), consistent with the full-dataset result. Under M3, once the shared habitat-loss gradient is reduced through residualisation of amount on patch number, the amount structure coefficient collapses toward zero or reverses sign (M3 SC = −0.091 for α; −0.257 for γ). This reversal is absent in the global ≥50% cover subset, where amount SC remains positive under M3 (α = +0.36; γ = +0.34). The contrast is biome-specific, not a general property of high-cover landscapes: within the Tropical SA regional component, patch number carries most of the fitted residual diversity alignment after the shared gradient is reduced, and the residualised amount predictor adds nothing or works in the opposing direction.

This result should be interpreted as a bounded ecological diagnostic. The subset is thin relative to a dedicated biome-level study and lacks explicit forest-specialist response strata, multi-scale habitat definitions, and multiple fragmentation metric families. Its contribution is narrower. The global additive null remains non-diagnostic in the regional setting where extinction-filtering expectations are most relevant to tropical forest fragmentation inference, and the recoverable configuration-associated direction is consistently negative once the shared gradient is reduced. The M3 coefficients reported here are interpreted as diagnostic evidence about gradient reallocation direction, not as effect-size estimates of fragmentation per se. The ecological mechanism implied by this regional pattern — the differential persistence of configuration-sensitive specialists under progressive deforestation — defines the basis for a response-surface framework in which magnitude, scale-dependence, and taxon-specificity can be addressed with an appropriately independent design. These Case 1 readouts are summarised in the first comparison column of Table 3.1.

3.2 Case Study 2. Response-estimand and metric-mechanism failure in Atlantic Forest tree beta diversity

The diagnostics presented here do not dispute that K26 documented meaningful Atlantic Forest beta-diversity patterns, spatial-distance structure, or endemism associations. The diagnostic claim is narrower: the negative patch-number coefficient in total beta diversity does not, by itself, identify an independent fragmentation-per-se effect on compositional dissimilarity. These are different claims, and establishing the stronger one requires evidence not fully provided

by the reported analyses. The corresponding Case 2 readouts are summarised in the second comparison column of Table 3.1.

3.2.1 Why K26 is the linked second case

K26 analysed tree beta diversity across 95 Atlantic Forest landscapes spanning the full forest-cover gradient of this hotspot. Using an additive model that included forest cover, number of forest patches, mean nearest-neighbour distance among patches (NND), mean spatial distance among plant inventories, and four bioclimatic variables, the authors reported that both forest cover and number of patches were significant negative predictors of total tree beta diversity. They concluded that the study successfully separated the effects of habitat fragmentation, habitat amount, and spatial distance on compositional change, making this one of the few observational studies in the Atlantic Forest to claim explicit fragmentation-per-se identification in a beta-diversity response.

K26 interpret the negative patch-number coefficient through generalist proliferation, edge exposure, habitat degradation, anthropogenic disturbance, specialist loss, and interaction between fragmentation and local degradation. Their earlier review supports this pathway in tropical and subtropical plant communities, where biotic homogenization and differentiation are frequently associated with land-use change, fragmentation, and edge-related processes (Kramer et al., 2023). The diagnostic question is therefore whether the fitted patch-number coefficient identifies that pathway, not whether biotic homogenization is plausible. This distinction matters because patch number is a subdivision metric, while the proposed mechanism involves matrix, edge, forest-condition, disturbance, and community-composition pathways that are not directly measured by patch count (Fletcher et al., 2007; Wang et al., 2014).

This case is analytically linked to Case Study 1 through two structural connections. Patch number belongs to the same subdivision-metric family as the patch-number predictor in the GS25–F26 reanalysis — algebraically coupled to habitat amount through the fixed-extent subdivision identity, and expected to generate geometric coefficient distortions under shared habitat-loss gradients. Beta diversity is precisely the response dimension that the GS25–F26 analysis identified as structured by spatial distance rather than the amount–configuration gradient: GS25 found that higher compositional dissimilarity in fragmented landscapes

disappeared once spatial distance was controlled, and the reanalysis showed no suppressor structure for beta diversity under the additive GS25–F26 specification. K26 apply a subdivision metric to the response dimension that prior global analysis had already shown to be governed by spatial processes, and then use a negative partial coefficient to infer fragmentation per se. The question is whether that inference survives the four-condition framework.

3.2.2 Predictor-space adequacy: nonlinear threshold geometry in the K26 dataset

Forest cover and number of patches in the K26 dataset have a Pearson correlation of r = −0.60, below the conventional |r| < 0.70 screening threshold, and variance inflation factors below 3 in the best reported model — both satisfying standard linear collinearity criteria. The linear diagnostics do not, however, capture the threshold structure of the amount–configuration relationship in this dataset. Landscapes in the low and middle forest-cover tertiles carry similar high mean patch counts (~144–154 patches), while the high-cover tertile collapses to ~42 patches. A GAM of patch number on forest cover confirms that this relationship is nonlinear (effective degrees of freedom > 1, F-test $p < 0.05$). The high-cover, many-patches cell — the region of predictor space required for same-amount, different-configuration contrasts — is structurally empty, for the same reason it is foreclosed in the GS25–F26 dataset: the land-use process that increases patch number by subdividing remaining forest simultaneously reduces total cover, making landscapes with both high cover and high patch count geometrically rare in habitat-loss systems. The K26 predictor-space geometry is shown in Figure 3.2c–d, alongside the corresponding GS25–F26 geometry in Figure 3.2a–b.

This threshold structure has a direct consequence for what the additive model can identify. The same-amount, different-configuration contrast required for fragmentation-per-se inference corresponds to variation across the horizontal axis of the predictor space at a fixed cover value. The tertile collapse from ~150 to ~42 patches between mid and high cover means that within the portion of the gradient where forest-dependent taxa are most likely to persist — at higher cover — the predictor space provides the least configuration variation. Standard linear diagnostics pass because the linear projection of the predictor cloud does not reveal the bounded geometry of the joint distribution; the GAM coupling confirms what the tertile structure illustrates directly.

3.2.3 Coefficient architecture: partial-slope amplification in total beta diversity

The coefficient architecture in the K26 dataset produces the opposite diagnostic pattern from Case Study 1. In Case Study 1, patch number had a large marginal association with the fitted richness gradient but a near-zero partial coefficient — suppressor attenuation, where including the dominant predictor deflates the subordinate's partial slope. In the K26 dataset, patch number has a near-zero marginal correlation with total tree beta diversity ($r = -0.12$) yet achieves a significant standardised partial coefficient in the joint model ($\beta = -0.267$, $p = 0.012$). This is partial-slope amplification. A negligible bivariate association becomes significant once a correlated predictor is included, because including forest cover reallocates shared gradient variance in a direction that inflates the patch-number coefficient rather than deflating it.

Amplification and attenuation are not the same failure, but they share the same root cause. Both arise when predictors remain embedded in a shared habitat-loss gradient that prevents independent variance attribution. In attenuation, the dominant predictor absorbs shared variance and the subordinate collapses toward zero. In amplification, the relationship between the two predictors and the response is such that conditioning on one inflates the other's apparent contribution — a form of distortion that is equally invisible to standard collinearity screening and equally incompatible with the independent-identification claim. The marginal-to-partial shift from $r = -0.12$ to $\beta = -0.267$ in the K26 patch-number predictor is the coefficient-architecture signature of this second pathology. A near-zero marginal association that becomes significant in an additive model is not a hallmark of predictor separability; it is a warning that shared gradient structure may be driving the partial coefficient rather than an independent ecological effect of configuration on compositional dissimilarity.

3.2.4 Response-estimand alignment: the Raup-Crick internal control

K26 report two beta-diversity analyses using different response definitions. Total beta diversity (Bray–Curtis dissimilarity) is the primary response to which the fragmentation-per-se interpretation was attributed. The authors additionally report Raup-Crick beta diversity as an internal methodological control, justifying this index on the grounds that standard dissimilarity measures are influenced by alpha-diversity patterns whereas Raup-Crick accounts for variation in alpha diversity among sampling sites. Under the alpha-corrected Raup-Crick response, neither forest cover nor number of patches retains significance in any selected model variant. Mean

spatial distance among inventories is the only consistently supported predictor (standardised effect = +0.31, $p < 0.01$; K26 Appendix 3.S7). The internal control provided by the authors themselves therefore shows that the total-beta associations with forest cover and patch number are response-specific: they appear in total beta diversity but disappear under the estimand designed to remove alpha-diversity mediation. The Raup-Crick result does not show that fragmentation cannot affect beta diversity through richness or abundance structure. It shows that the reported patch-number association is contingent on the total-beta response and does not generalise to the alpha-corrected diagnostic response. This Raup-Crick failure is not specific to the original additive model: as the metric and predictor-extent triangulation shows, no alternative configuration metric — including independent patch-count measures, nearest-neighbour isolation, edge metrics, functional connectivity, amount-decoupled arrangement metrics, and core-area heterogeneity — recovers a Raup-Crick signal across any tested predictor extent.

The residualisation sequence provides the Case Study 2 analogue of the M1–M4 diagnostic in Case Study 1, but only for the forest-cover–patch-number pair. Residualising patch number on forest cover (K2) reduced the total-beta patch-number coefficient from −0.267 to −0.213 while retaining support, indicating that the total-beta subdivision signal is not merely a reconstruction failure or a trivial cover proxy. However, the same residualised patch-number term was not supported under Raup-Crick, and patch number was also unsupported when forest cover was reciprocally residualised on patch number (K3; $\beta = -0.111$, $p = 0.188$). This predictor-geometry sensitivity therefore strengthens rather than rescues the attribution problem: the patch-count signal is real enough to persist in total beta diversity, but it remains response-contingent and does not identify fragmentation per se without response-estimand, metric-mechanism, and domain convergence. The separate predictor-extent audit asks a different question: whether mechanism-matched metrics recover the signal across operational scales.

3.2.5 Metric-mechanism correspondence and domain stability

Mean nearest-neighbour distance among patches, one of K26's own retained configuration metrics, captures interpatch isolation rather than subdivision. Patch number and nearest-neighbour distance are empirically decoupled in K26's deposited predictors ($r = -0.04$) and in the independent extraction ($r = -0.11$), so the supported patch-number coefficient cannot be read as one side of a coupled subdivision–isolation gradient. Under both total beta diversity

and Raup-Crick models, NND adds no explanatory support: its partial coefficient is not significant in any variant. The reported configuration signal is therefore carried specifically by patch subdivision, not by patch isolation. The retained isolation metric most directly linked to dispersal-mediated fragmentation mechanisms does not recover the reported patch-number coefficient.

The endemism analysis reported by K26 provides a further response-specific contrast. Forest cover and temperature seasonality are the supported predictors and are near-orthogonal (VIF = 1.02), whereas patch number attenuates to zero rather than amplifies when habitat amount is controlled — the opposite of the total-beta coefficient architecture. The endemism result therefore rests on a simpler and more stable predictor relationship than the total-beta patch-number attribution, and it further localises the patch-number signal to total beta diversity rather than to a general configuration effect across biodiversity estimands.

The cover-stratified analysis is not a test against threshold theory. The low-cover concentration of the NP signal is compatible with the threshold prediction of stronger configuration-associated responses under habitat scarcity. The diagnostic question here is different: whether the reported global additive coefficients are stable across the realised cover gradient and whether the low-cover signal survives the independent checks required for fragmentation-per-se attribution. They do not. Restricting the analysis to landscapes with ≥50% forest cover collapses both effects: forest cover $\beta = -0.027$ ($p = 0.936$) and patch number $\beta = +0.002$ ($p = 0.995$). Mean spatial distance remains the only supported predictor. The full-dataset negative coefficients are therefore domain-specific rather than cover-general. This domain-specificity is not sufficient for fragmentation-per-se attribution because the low-cover NP signal remains confined to total beta diversity, is not recovered by the retained isolation metric, and does not persist under Raup-Crick beta diversity.

The diagnostic conditions converge on the same interpretation. The amount–configuration space does not provide the high-cover, high-subdivision contrasts required for a clean fragmentation-per-se estimate. The coefficient architecture shows amplification rather than independent identification, because a weak marginal patch-number association becomes a supported partial coefficient only after conditioning on forest cover. The response check shows that the patch-number signal is specific to total beta diversity and does not persist under the

alpha-corrected Raup-Crick response. The metric check gives the same result: isolation, connectivity, edge exposure, interior depth, and amount-decoupled arrangement metrics do not recover the signal across the tested predictor extents. Domain stability also fails because the coefficient is concentrated in low-cover landscapes and collapses at ≥50% cover. The dominant failure in the K26 case is over-attribution of a response-specific patch-count association to fragmentation per se, not absence of a total-beta signal.

3.2.6 Extended metric and predictor-extent triangulation

The K26 patch-number coefficient lacks the convergence expected of a fragmentation-per-se signal across mechanism-matched metrics and predictor extents. I evaluated this by retaining one representative per mechanism family: nearest-neighbour distance for isolation, edge coverage for edge exposure, distance inside forest for interior depth, high-gap functional connectivity for connectivity, CLUMPY and NLSI for amount-decoupled arrangement, and core-area variability for core heterogeneity. The full 19-metric audit was used to screen sensitivity metrics that were redundant, amount-coupled, or unstable under the realised predictor geometry. The metric-by-extent coefficient and structure-coefficient patterns for the focal mechanism-congruent subset are shown in Figure 3.4; the broader metric screen is retained as provenance for the triangulation audit.

Independent extraction of K26's landscape predictors from a 30 m binary forest layer at matched extents confirmed that the total-beta patch-number signal is real and directionally reproducible, not a geospatial reconstruction failure. Forest cover was recovered at $r = 0.949$, patch number at $r = 0.874$, and nearest-neighbour distance at $r = 0.823$. Independent patch number retained the same negative total-beta direction and recovered approximately 74% of the published effect size. The signal is present but narrow.

The metric pattern fails to triangulate the proposed mechanism. The proposed homogenization pathway involves matrix condition, edge exposure, habitat degradation, disturbance history, generalist proliferation, and specialist loss, none of which is directly measured by a patch count (Kramer et al., 2023, 2026). A patch-count coefficient can therefore be real while the matrix and forest-condition pathways required by the interpretation remain unidentified. If patch number were standing in for those pathways, metrics closer to edge

exposure, isolation, connectivity, or amount-decoupled spatial arrangement should provide convergent support (Fletcher et al., 2007; Wang et al., 2014; Fletcher et al., 2018; Fletcher et al., 2023). No such convergence was detected. Pure arrangement metrics, including CLUMPY and the normalised landscape shape index, are non-significant for both total beta and Raup-Crick across all tested predictor extents. CLUMPY is especially informative because it is nearly orthogonal to forest cover but still shows no detectable response signal. Across five circular-buffer extents from 25 to 400 km², patch-number support was scale-sensitive, whereas representative non-collinear mechanism-matched metrics remained unsupported at every extent and for both response indices. The mechanism-triangulation failure is therefore not an artefact of K26's specific operational extent. The surviving result is a response-specific association with one patch-count operationalisation of subdivision, not convergent evidence across mechanisms, metrics, and predictor extents that fragmentation per se drives biotic homogenization.

Structure coefficients reinforce this interpretation. In the original total-beta model, patch number has coherent coefficient architecture: its beta weight and structure coefficient are aligned in sign and magnitude, with NP structure coefficients approximately −0.25 to −0.30 across specifications, indicating that the total-beta NP signal is not an artefact of suppressor geometry. The signal is therefore real in the fitted model. Mean spatial distance, however, has substantially larger fitted-gradient alignment across total-beta specifications and dominates the Raup-Crick fitted gradient across all metric, source, and extent checks, with spatial-distance structure coefficients approximately +0.83 to +0.95 across specifications. Forest cover shows sign-reversal behaviour under Raup-Crick, indicating suppressor geometry in the alpha-corrected response. The K26 case has a configuration signal, but it is too response-specific, too metric-specific, and too scale-sensitive to identify fragmentation per se. A path model testing whether the total-beta NP signal is mediated by alpha richness is consistent with richness-sensitive structure but does not provide unambiguous evidence for richness mediation. It is not relied on as a main result.

3.3 Cross-case synthesis

3.3.1 Two failure modes from one metric family

The two case studies do not produce the same diagnostic failure. That difference is the chapter's first cross-case result. In the GS25–F26 richness case, patch number is attenuated

toward a near-zero partial coefficient despite strong fitted-gradient alignment. The structure coefficient is large and negative, the beta weight collapses, and the semi-partial R^2 is near zero. This is the suppressor signature: the metric tracks the fitted gradient but contributes negligible unique variance because habitat amount absorbs the shared component. In the K26 beta-diversity case, patch number is amplified from a near-zero marginal correlation with total beta diversity ($r = -0.12$) into a significant standardised partial coefficient ($\beta = -0.267$, $p = 0.012$). Partial-slope amplification is the opposite signature: the metric carries little univariate alignment with the response, but the joint model promotes it through residualised geometry. The same subdivision-metric family therefore produces opposite coefficient pathologies under the two response dimensions: attenuation in richness models and amplification in total beta diversity. The residualisation diagnostics also differ in scope. In GS25–F26, M1–M4 reallocate the full amount–patch-number gradient. In K26, K2/K3 test the forest-cover–patch-number pair, while the predictor-extent audit asks whether the total-beta patch-count signal generalises across operational scales. Both diagnostic routes leave the identification problem intact.

3.3.2 The response dimension is the analytical hinge

The two cases also share an asymmetry that the additive coefficient cannot signal. In Case 1, the suppressor signature is present in alpha and gamma richness but absent under beta diversity, where mean spatial distance dominates. In Case 2, the amplification signature is present in total beta diversity but absent under Raup-Crick, where again mean spatial distance is the only consistently supported predictor. The Raup-Crick internal control in the K26 dataset confirms what the GS25 beta result implies: the configuration coefficient supported in total beta diversity does not transfer to the alpha-corrected response. Beta diversity is therefore not a generic biodiversity index that can absorb a coefficient interpretation derived from richness. It tracks spatial-assembly and alpha-diversity structure that the additive amount-configuration model is not designed to identify.

3.3.3 Estimate–estimand mismatch as the unifying frame

Both case studies fail for a reason more general than collinearity, suppressor structure, or metric ambiguity. In each case, the statistical estimate is well-defined and reproducible. The estimate is the partial slope of patch number in the fitted additive model. The ecological target is

fragmentation per se: an effect of spatial configuration at comparable habitat amount on a response and metric aligned to a proposed mechanism. The two coincide only when the predictor space contains the relevant same-amount, different-configuration contrasts, the response measures the ecological quantity the mechanism is supposed to drive, and the metric matches the proposed pathway. In Case 1, the predictor space is structurally inadequate: same-amount pairs are scarce (3/37) and 31/37 paired studies re-express the original directional contrast rather than orthogonalise it. In Case 2, the predictor space is also inadequate, the response check shows that the patch-number signal is contingent on total beta diversity, and the metric-scale check fails: the signal is directionally reproducible for patch number, but it is scale-sensitive and not recovered by mechanism-matched metrics across 25–400 km² predictor extents. The unifying frame is whether the estimate identifies the estimand, not whether the additive coefficient is null, negative, or positive. In both cases, it does not — but for different reasons, and with opposite coefficient signs. That is the chapter's central contribution. Table 3.1 collects these diagnostic readouts for both cases, with the cross-case inference in the final column.

4. Discussion

4.1 Amount control is necessary, but not sufficient

Additive habitat amount control addresses one source of confounding in fragmentation inference, but including both predictors in a model does not demonstrate that the realised predictor space supports the ecological contrast the coefficient is used to represent. The two cases analysed here show this from opposite coefficient signs. In the GS25–F26 case, a near-zero patch-number coefficient was interpreted as evidence that fragmentation effects disappear once habitat amount is controlled. In the K26 case, a negative patch-number coefficient was interpreted as evidence for an independent configuration effect on compositional dissimilarity. The diagnostic framework shows that neither interpretation can be sustained on the basis of the coefficient alone.

Additive modelling is not the problem. Additive models remain appropriate when the data support their estimand, and conventional collinearity thresholds are a necessary screen. They do not prove that predictor separability has been achieved. In both datasets here, standard

diagnostics passed while nonlinear geometric coupling remained. The estimand–coefficient gap persisted under both coefficient signs.

This reframes what it would mean to move the fragmentation debate forward. Valente et al. (2023) argued that the debate has become locked in partly because studies differ in design, scale, metric choice, habitat definition, ecological context, and biodiversity response, yet these differences are often converted into opposing general claims about fragmentation. The present diagnostics make that problem operational. The relevant contrast is not habitat amount versus fragmentation as competing positions. It is whether a given coefficient identifies the fragmentation-per-se estimand under the realised predictor geometry, response variable, metric, domain, and scale.

4.2 Subdivision metrics are structurally vulnerable to shared habitat-loss geometry

Patch number and patch density are generated jointly with habitat loss, making their partial coefficients especially sensitive to how the additive model allocates a shared degradation gradient between predictors. Within a fixed landscape extent, forest conversion increases patch count through progressive subdivision of remaining habitat while simultaneously reducing total cover, co-generating the two predictors along a shared trajectory before any landscape is sampled. This coupling has a directed structure — cover loss constrains patch-number geometry more strongly than the reverse — which is why nonlinear GAM coupling is asymmetric in both datasets and why standard linear diagnostics, which project the predictor cloud onto a single axis, cannot detect it.

The K26 case makes the metric problem especially clear. The only metrics that approach or reach support are patch-count subdivision metrics. These are precisely the metrics least able to isolate mechanism because they can covary with remnant number, remnant size, edge exposure, matrix condition, disturbance history, and degradation without directly measuring the pathways Kramer et al. invoke (Kramer et al., 2023, 2026). Metrics closer to the proposed pathways, including isolation, connectivity, edge exposure, and amount-decoupled spatial arrangement, do not converge on the result (Fletcher et al., 2007; Wang et al., 2014; Fletcher et al., 2018; Fletcher et al., 2023). The pattern is stronger than a single failed alternative metric. Mechanism triangulation fails because support is concentrated in the ambiguous subdivision metric, while

metrics with cleaner mechanism interpretation return null support. A coefficient produced in this configuration cannot, on its own, identify the proposed homogenization mechanism. The scale audit reinforces this conclusion. At no predictor extent from 25 to 400 km² does a non-collinear mechanism-matched metric recover the total-beta association, confirming that the triangulation failure is not contingent on K26's operational scale. A patch-number null is therefore axis-specific rather than evidence of configuration invariance; the same logic is illustrated by a corridor-reachability analysis in which the patch-number interaction was weak but a reachability metric showed a strong cover-dependent response (Martínez-Lanfranco, 2026c).

The same structural vulnerability appears in the GS25–F26 case, but with the opposite coefficient manifestation. There, patch number also indexes subdivision within a shared habitat-loss gradient, yet the additive model allocates the shared gradient primarily to habitat amount, attenuating rather than amplifying the patch-number coefficient. The two cases therefore do not show that subdivision metrics become non-diagnostic in only one direction. They show that the same metric family can be pushed toward either attenuation or amplification depending on the response dimension and the realised predictor geometry.

Defining fragmentation per se as the residual configuration signal left after habitat amount control would not solve the attribution problem. Without convergence across response, metric, and scale, that residual category remains unidentified rather than mechanistically interpretable.

The consequences for coefficient interpretation are specific to subdivision metrics. An isolation metric such as mean nearest-neighbour distance is not algebraically coupled to habitat amount through the subdivision identity and does not carry the same structural vulnerability. In the K26 dataset, NND adds no support in any model variant, whereas the subdivision metric generates the reported negative coefficient in total beta diversity. That contrast does not make isolation ecologically unimportant in Atlantic Forest tree communities. It shows that the supported signal is carried by subdivision geometry, the most structurally coupled component of the predictor set.

4.3 Attenuation and amplification are sign-symmetric failures from the same metric family

The central cross-case result is that the same subdivision-metric family produces opposite coefficient pathologies. In the richness case, patch number shows the suppressor signature, with a large structure coefficient, near-zero beta weight, and near-zero semi-partial R^2. In the beta-diversity case, patch number shows the amplification signature, with a near-zero marginal association, significant negative partial coefficient, and response-specific fitted-model support. Both pathologies arise from shared habitat-loss geometry that prevents independent variance attribution, but they manifest at opposite ends of the marginal-to-partial relationship.

The K26 case also shows why amplification alone is not the full diagnosis. The negative partial coefficient is real enough to persist after residualising patch number on forest cover in some total-beta specifications, and it is supported by structure-coefficient diagnostics indicating coherent fitted-gradient alignment for total beta. The coefficient is present in total beta diversity, but it remains response-specific and mechanism-unidentified. It does not survive Raup-Crick under the tested predictor-geometry sensitivities, and it is not recovered by metric families or predictor extents closer to the mechanisms implied by biotic homogenization. This distinction is essential because a real coefficient can still be the wrong estimate for the ecological target.

This sign symmetry is the chapter's central cross-case contribution. Fragmentation debates often treat coefficient sign as the primary evidence. Under that framing, a null configuration coefficient supports habitat amount control, whereas a negative coefficient identifies fragmentation per se. The comparison here shows that neither inference follows from sign alone. A null coefficient can reflect suppression of a configuration-associated signal. A negative coefficient can reflect amplification by the same shared-gradient geometry. What distinguishes the two cases is the relationship between the marginal association and the partial slope. In GS25–F26, the partial coefficient collapses below the fitted-gradient alignment. In K26, it inflates above a weak marginal association. The K26 coefficient and the F26 null are mirror-image cases because the same patch-number family misses the estimand in opposite ways.

The symmetry is diagnostic, not evidentially identical. In Case Study 1, the main inference is a rejection of the null reading. A near-zero additive coefficient cannot be treated as evidence that configuration-associated richness responses are absent. In Case Study 2, the inference evaluates a positive identification claim. A supported negative coefficient identifies

fragmentation per se only if it converges across response, metric, scale, and domain. The two cases therefore share a coefficient-identification failure, but the evidential burden differs between them.

4.4 Beta diversity requires response-estimand verification, not coefficient transfer

The beta-diversity comparison is the analytical hinge between the two cases. Total beta diversity carries alpha-diversity structure and spatial-distance effects that are not reducible to the amount–configuration contrast, and the configuration signal that appears in total beta disappears under both the alpha-corrected Raup-Crick response and the GS25–F26 framework where it was absent as a suppressor signal from the outset. In both cases, spatial distance rather than subdivision governs the residual beta-diversity response. The low-cover concentration of the K26 NP signal is ecologically compatible with threshold expectations, but threshold compatibility does not substitute for response-estimand verification. The signal does not persist under alpha-diversity correction and is not recovered by the retained isolation metric.

The expanded K26 diagnostics strengthen this point because the response-estimand failure is not confined to the original additive model with patch number as the configuration metric. It holds across representative mechanism-matched metrics and across predictor extents from 25 to 400 km². The total-beta signal may reflect biologically meaningful homogenization or degradation-associated compositional structure, but the alpha-corrected diagnostic response does not support an independent configuration effect under the tested metric and scale alternatives.

The response dimension that GS25 showed to be structured by spatial distance is the same response dimension K26 used to infer fragmentation per se. The Raup-Crick internal control in the K26 dataset removes forest cover and patch number from support while retaining spatial distance as the only consistent predictor. The two cases therefore raise a parallel response-estimand warning from different datasets. Total beta diversity remains ecologically meaningful, but a total-beta patch-count coefficient cannot be treated as a fragmentation-per-se effect without response, metric, and scale convergence.

Spatial distance should be treated as an interpretable spatial-assembly signal. Beta-diversity patterns can reflect different assembly mechanisms once species-pool sampling effects are accounted for, and Atlantic Forest tree composition is structured by both environmental and

spatial components (Myers et al., 2013; Zwiener et al., 2020). Spatial distance among inventories, patch subdivision, patch isolation, edge exposure, and functional connectivity are distinct quantities. Support for one cannot be transferred to the others without a mechanistic bridge (Fletcher et al., 2007; Wang et al., 2014; Fletcher et al., 2018; Fletcher et al., 2023).

4.5 Ecological context and the Tropical South American landing

The richness result shifts the diagnostic from beta-diversity response alignment to ecological domain. The Tropical SA result is consistent with extinction-filtering expectations. Under progressive tropical deforestation, disturbance-tolerant generalists persist across the degradation gradient while range-restricted forest specialists are disproportionately lost as cover declines below viability thresholds (Henle et al., 2004; Banks-Leite et al., 2014). The assemblages remaining along the gradient are therefore filtered rather than random. At low cover, disturbance-tolerant taxa are expected to dominate because many forest specialists have already been lost. At intermediate-to-high cover, enough configuration-sensitive species may remain for spatial subdivision to structure residual richness once the shared amount gradient is reduced. The diagnostic recovery under M3 is consistent with that expectation. The high-cover subset is therefore not only a robustness check. It is an ecologically informed detectability check, asking whether the configuration-associated signal appears in the portion of the gradient where configuration-sensitive taxa are expected to remain present and where spatial subdivision can still be expressed as a biological mechanism rather than only as a legacy of severe habitat loss (Andrén, 1994; Villard & Metzger, 2014).

The Tropical SA subset is diagnostic rather than an Atlantic Forest effect estimate. It is thin relative to a dedicated regional study, lacks explicit forest-specialist response strata, and does not include multi-scale metric families or explicit habitat thresholds. Its value is the contrast it creates with the K26 case. In Case Study 1, the configuration-associated richness signal becomes more recoverable in the high-cover subset and the Tropical SA domain. M3 coefficients become more negative, nonlinear asymmetry increases (16.2× vs. 7.75×), and habitat amount no longer retains positive fitted-gradient alignment in Tropical SA. In Case Study 2, the patch-number signal concentrates in low-cover landscapes and collapses at ≥50% cover. The asymmetry reflects the response dimension. Richness responses can retain configuration-sensitive specialists where forest cover remains adequate. Total beta diversity can instead

register broader landscape degradation. This pattern points to response surfaces rather than single additive coefficients. The steepest configuration responses are expected at intermediate-to-high cover where specialists persist and configuration constrains their dynamics, and flatter responses are expected where specialists have already been filtered out or habitat remains non-limiting.

That contrast should not be over-interpreted. It is diagnostic motivation rather than direct identification, and it does not supply the multi-metric or multi-scale triangulation that a direct fragmentation-per-se estimate would require.

4.6 From diagnostic failure to response-surface inference

The two diagnostic cases motivate a response-surface approach in which the configuration effect varies across the habitat-amount gradient rather than entering as a fixed partial slope. The estimand is most nearly approached where configuration-sensitive species persist and configuration constrains their dynamics. Where filtering has already occurred, or where habitat is abundant enough to be non-limiting, the additive model can return a null or ambiguous coefficient regardless of ecological reality. Diagnosing why coefficients fail to identify the estimand specifies what a response-surface analysis must demonstrate. The diagnosis is falsifiable rather than open-ended. It predicts that the negative configuration signal residualisation recovers here should also surface in landscapes where habitat amount and configuration are separable by design, so a stable near-zero coefficient across such uncoupled datasets would point to a genuine ecological null rather than suppression. Even where amount and configuration are better separated, pooled coefficients can remain non-diagnostic when responses and mechanisms vary with species sensitivity, matrix conditions, and the cover domains in which configuration is expected to matter (Ewers & Didham, 2006; Henle et al., 2004; Villard & Metzger, 2014).

The shared failure in both cases is an estimate–estimand mismatch. The estimate is simply a partial slope for patch number in an additive model with habitat amount. That estimate identifies fragmentation per se only when the realised data contain the relevant same-amount contrasts, the response matches the mechanism, and the metric represents the pathway being interpreted. In both cases analysed here, those conditions fail, and they fail under opposite coefficient signs. A single additive coefficient therefore cannot resolve habitat-amount

invariance against configuration-sensitive mechanisms when those mechanisms are expected to vary across cover, response dimension, and ecological filtering history. These diagnostic cases therefore motivate a response-surface research programme that asks how configuration effects vary across habitat amount, ecological context, response dimension, metric choice, and scale.

5. Conclusions

Additive habitat amount control is a necessary step toward fragmentation-per-se inference, but it is not a sufficient diagnostic. In the two cases analysed here, the same subdivision-metric family produced opposite coefficient pathologies under shared habitat-loss geometry. Patch number was attenuated toward a near-zero richness coefficient in the global multi-taxa dataset despite strong fitted-gradient alignment — suppressor attenuation driven by the dominant habitat-loss gradient allocating shared variance to habitat amount. Patch number was amplified from a negligible marginal association to a significant negative partial coefficient in Atlantic Forest tree total beta diversity — partial-slope amplification that does not survive the alpha-corrected internal control provided by the authors themselves.

These outcomes show that coefficient sign alone cannot resolve fragmentation per se. A null coefficient can reflect geometric suppression rather than ecological absence. A negative coefficient can reflect amplification and mechanism ambiguity rather than independent identification. The K26 case also shows that even a real, directionally reproducible patch-count signal is not enough when it remains response-specific, scale-sensitive, and mechanism-unidentified. Fragmentation-per-se inference requires demonstrated predictor-space adequacy, response-estimand alignment, metric-mechanism correspondence, and domain-specific stability. Rather than estimating fragmentation effects, the diagnostic framework applied here establishes what evidence is and is not present in the realised datasets, and clarifies that valid fragmentation-per-se identification would require convergence across the predictor geometry, the response variable, the configuration metric, and the spatial scale at which the proposed mechanism is supposed to operate.

References

Andrén, H. (1994). Effects of habitat fragmentation on birds and mammals in landscapes with different proportions of suitable habitat: a review. Oikos, 71, 355–366.

Banks-Leite, C., Pardini, R., Tambosi, L. R., Pearse, W. D., Bueno, A. A., Bruscagin, R. T., Condez, T. H., Dixo, M., Igari, A. T., Martensen, A. C., & Metzger, J. P. (2014). Using ecological thresholds to evaluate the costs and benefits of set-asides in a biodiversity hotspot. Science, 345, 1041–1045.

Betts, M. G., Fahrig, L., Hadley, A. S., Halstead, K. E., Bowman, J., Robinson, W. D., Wiens, J. A., & Lindenmayer, D. B. (2014). A species-centered approach for uncovering generalities in organism responses to habitat loss and fragmentation. Ecography, 37, 517–527.

Chase, J. M., Kraft, N. J. B., Smith, K. G., Vellend, M., & Inouye, B. D. (2011). Using null models to disentangle variation in community dissimilarity from variation in α-diversity. Ecosphere, 2(2), art24.

Didham, R. K., Kapos, V., & Ewers, R. M. (2012). Rethinking the conceptual foundations of habitat fragmentation research. Oikos, 121, 161–170.

Dormann, C. F., Elith, J., Bacher, S., Buchmann, C., Carl, G., Carré, G., Marquéz, J. R. G., Gruber, B., Lafourcade, B., Leitão, P. J., Münkemüller, T., McClean, C., Osborne, P. E., Reineking, B., Schröder, B., Skidmore, A. K., Zurell, D. and Lautenbach, S. (2013). Collinearity: a review of methods to deal with it and a simulation study evaluating their performance. Ecography, 36, 27–46.

Ewers, R. M., & Didham, R. K. (2006). Confounding factors in the detection of species responses to habitat fragmentation. Biological Reviews, 81, 117–142.

Fahrig, L. (2003). Effects of habitat fragmentation on biodiversity. Annual Review of Ecology, Evolution, and Systematics, 34, 487–515.

Fahrig, L., Galán-Acedo, C., Edwards, B. P. M., Habrich, A. K., Martin, A. E., Albert, C. H., Arroyo-Rodríguez, V., Binley, A. D., Cole, J. R., Lockhart, J. K., Nielsen, S. E., Riva, F., Schmoeller, M. P. R., & Tscharntke, T. (2026). Why controlling for habitat amount is critical for resolving the fragmentation debate. Conservation Biology, e70245.

Fischer, J., & Lindenmayer, D. B. (2007). Landscape modification and habitat fragmentation: a synthesis. Global Ecology and Biogeography, 16, 265–280.

Fletcher, R. J., Ries, L., Battin, J., & Chalfoun, A. D. (2007). The role of habitat area and edge in fragmented landscapes: definitively distinct or inevitably intertwined? Canadian Journal of Zoology, 85, 1017–1030.

Fletcher, R. J., Reichert, B. E., & Holmes, K. (2018). The negative effects of habitat fragmentation operate at the scale of dispersal. Ecology, 99, 2176–2186.

Fletcher, R. J., Betts, M. G., Damschen, E. I., Hefley, T. J., Hightower, J., Smith, T. A. H., Fortin, M.-J., & Haddad, N. M. (2023). Addressing the problem of scale that emerges with habitat fragmentation. Global Ecology and Biogeography, 32, 828–841.

Gonçalves-Souza, T., Chase, J. M., Haddad, N. M., Vancine, M. H., Didham, R. K., Melo, F. L. P., Aizen, M. A., Bernard, E., Chiarello, A. G., Faria, D., Gibb, H., de Lima, M. G., Magnago, L. F. S., Mariano-Neto, E., Nogueira, A. A., Nemésio, A., Passamani, M., Pinho, B. X., Rocha-Santos, L., … Sanders, N. J. (2025). Species turnover does not rescue biodiversity in fragmented landscapes. Nature, 640(8059), 702–706.

Hadley, A. S., & Betts, M. G. (2016). Refocusing habitat fragmentation research using lessons from the last decade. Current Landscape Ecology Reports, 1, 55–66.

Henle, K., Davies, K. F., Kleyer, M., Margules, C., & Settele, J. (2004). Predictors of species sensitivity to fragmentation. Biodiversity and Conservation, 13, 207–251.

Koper, N., Schmiegelow, F.K. and Merrill, E.H., 2007. Residuals cannot distinguish between ecological effects of habitat amount and fragmentation: implications for the debate. Landscape Ecology, 22(6), 811-820.

Kramer, J. M. F., Zwiener, V. P., & Müller, S. C. (2023). Biotic homogenization and differentiation of plant communities in tropical and subtropical forests. Conservation Biology, 37, e14025.

Kramer, J. M. F., Zwiener, V. P., Camana, M., de Lima, R. A. F., & Müller, S. C. (2026). Effects of deforestation and fragmentation on tree beta diversity and endemism across landscapes in the Atlantic Forest biodiversity hotspot. Global Change Biology, 32, e70855.

Lundberg, I., Johnson, R., & Stewart, B. M. (2021). What is your estimand? Defining the target quantity connects statistical evidence to theory. American Sociological Review, 86, 532–565.

Martínez-Lanfranco, J. A. (2026a). Habitat amount control is necessary but not sufficient to resolve the fragmentation debate. (Conservation Biology, accepted). EcoEvoRxiv: https://doi.org/10.32942/X2KT05

Martínez-Lanfranco, J. A. (2026b). Habitat-loss-driven predictor coupling limits inference about the independent effects of configuration in additive habitat-amount models: implications for the fragmentation debate. (Ecosphere, in review). bioRxiv: https://doi.org/10.64898/2026.04.12.718042

Martínez-Lanfranco, J. A. (2026c). A patch-number null does not rule out conditional configuration effects. Conservation Letters, 19, e70063.

Martínez-Lanfranco, J. A. (2026d). Configuration coefficients can mislead fragmentation-based conservation inference from beta diversity. (Conservation Biology, in review). EcoEvoRxiv, https://ecoevorxiv.org/repository/view/14049/

Myers, J. A., Chase, J. M., Jiménez, I., Jørgensen, P. M., Araujo-Murakami, A., Paniagua-Zambrana, N., & Seidel, R. (2013). Beta-diversity in temperate and tropical forests reflects dissimilar mechanisms of community assembly. Ecology Letters, 16, 151–157.

Pedersen, E. J., Miller, D. L., Simpson, G. L., & Ross, N. (2019). Hierarchical generalized additive models in ecology: an introduction with mgcv. PeerJ, 7, e6876.

Prunier, J. G., Colyn, M., Legendre, X., Nimon, K. F., & Flamand, M. C. (2015). Multicollinearity in spatial genetics: separating the wheat from the chaff using commonality analyses. Molecular Ecology, 24, 263–283.

Ray-Mukherjee, J., Nimon, K., Mukherjee, S., Morris, D., Slotow, R., & Hamer, M. (2014). Using commonality analysis in multiple regressions: a tool to decompose regression effects in the face of multicollinearity. Methods in Ecology and Evolution, 5, 320–328.

Riva, F., Koper, N., & Fahrig, L. (2024). Overcoming confusion and stigma in habitat fragmentation research. Biological Reviews, 99, 1411–1424.

Ruffell, J., Banks-Leite, C., & Didham, R. K. (2016). Accounting for the causal basis of collinearity when measuring the effects of habitat loss versus habitat fragmentation. Oikos, 125(1), 117–125.

Smith, A. C., Koper, N., Francis, C. M., & Fahrig, L. (2009). Confronting collinearity: comparing methods for disentangling the effects of habitat loss and fragmentation. Landscape Ecology, 24, 1271–1285.

Stoffel, M. A., Nakagawa, S., & Schielzeth, H. (2021). partR2: Partitioning R^2 in generalized linear mixed models. PeerJ, 9, e11414.

Valente, J. J., Gannon, D. G., Hightower, J., Kim, H., Leimberger, K. G., Macedo, R., Rousseau, J. S., Weldy, M. J., Zitomer, R. A., Fahrig, L., Fletcher, R. J., Wu, J., & Betts, M. G.

(2023). Toward conciliation in the habitat fragmentation and biodiversity debate. Landscape Ecology, 38, 2717–2730.

Villard, M.-A., & Metzger, J. P. (2014). Beyond the fragmentation debate: a conceptual model to predict when habitat configuration really matters. Journal of Applied Ecology, 51, 309–318.

Wang, X., Blanchet, F. G., & Koper, N. (2014). Measuring habitat fragmentation: an evaluation of landscape pattern metrics. Methods in Ecology and Evolution, 5, 634–646.

Zwiener, V. P., Padial, A. A., & Marques, M. C. M. (2020). The mechanisms explaining tree species richness and composition are convergent in a megadiverse hotspot. Biodiversity and Conservation, 29, 799–815.

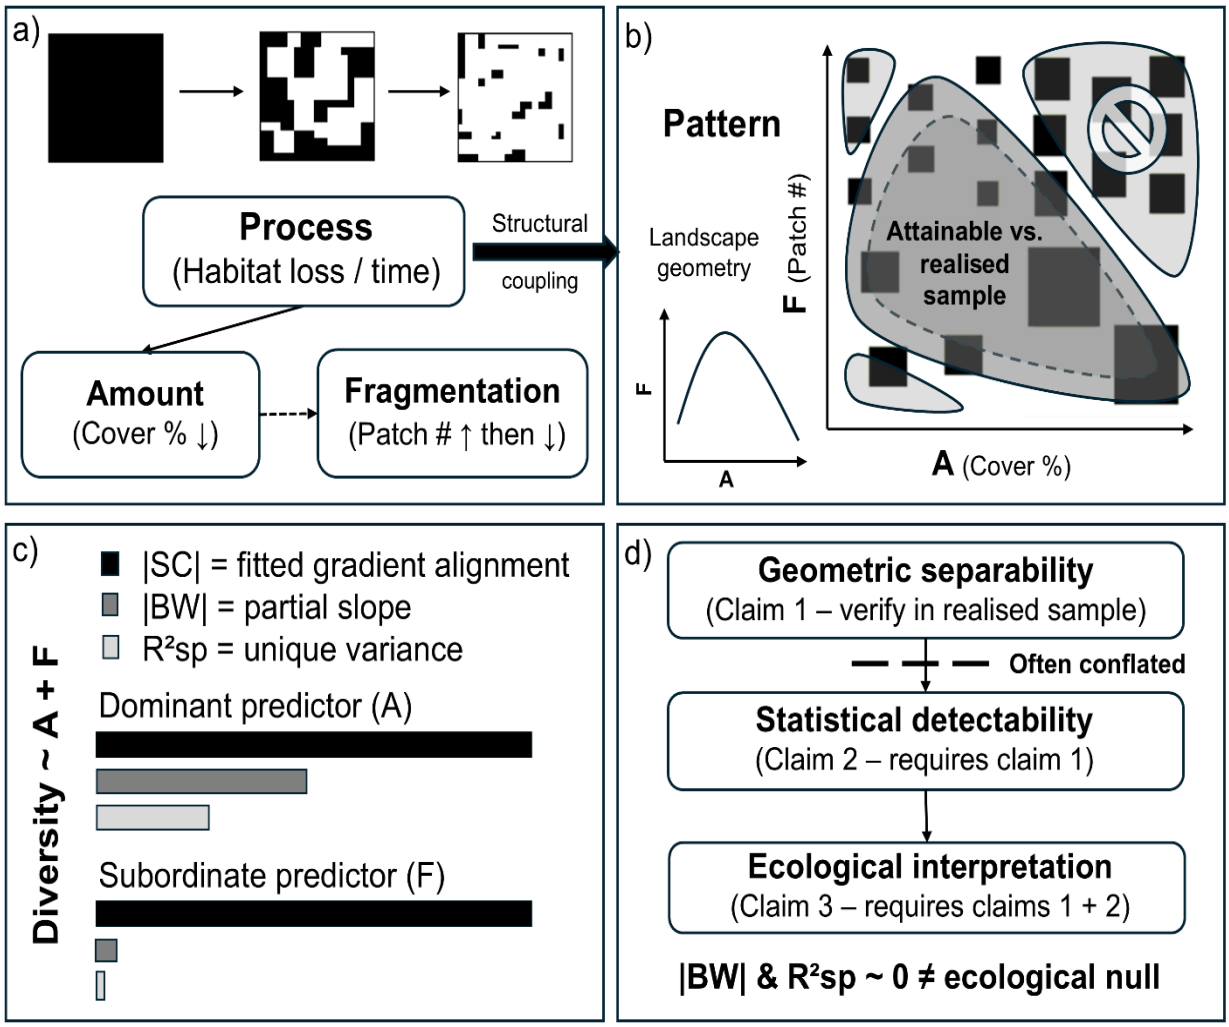


Figure 3.1. Diagnostic architecture for fragmentation-per-se inference. Habitat loss can co-generate habitat amount and spatial configuration along a directed, asymmetric gradient before any landscape is sampled; the dashed arrow marks derived coupling rather than an independent causal pathway (a). This produces a bounded realised predictor space in which same-amount, different-configuration landscapes may be scarce, and high-amount, many-patch combinations may be strongly constrained or poorly sampled (b). In an additive model of diversity on amount and configuration, the structure coefficient measures signed alignment with the fitted response gradient, the beta weight measures the unique partial slope, and the semi-partial R^2 measures unique variance; a subordinate predictor can align strongly with the gradient while contributing little unique variance (c). Independent interpretation of an additive configuration coefficient therefore requires more than including both predictors: the realised predictor space must supply the same-amount contrast, the response must match the proposed process, the metric must

represent the proposed mechanism, and the coefficient must be stable or explicitly localised across the domain where it is interpreted. Panel (d) places this requirement in an inferential order, separating geometric separability, statistical detectability, and ecological interpretation, so that ecological interpretation depends on prior support for the geometric and statistical claims. Near-zero beta weights and semi-partial R^2 do not by themselves constitute an ecological null. Figure adapted from Martínez-Lanfranco (2026b).

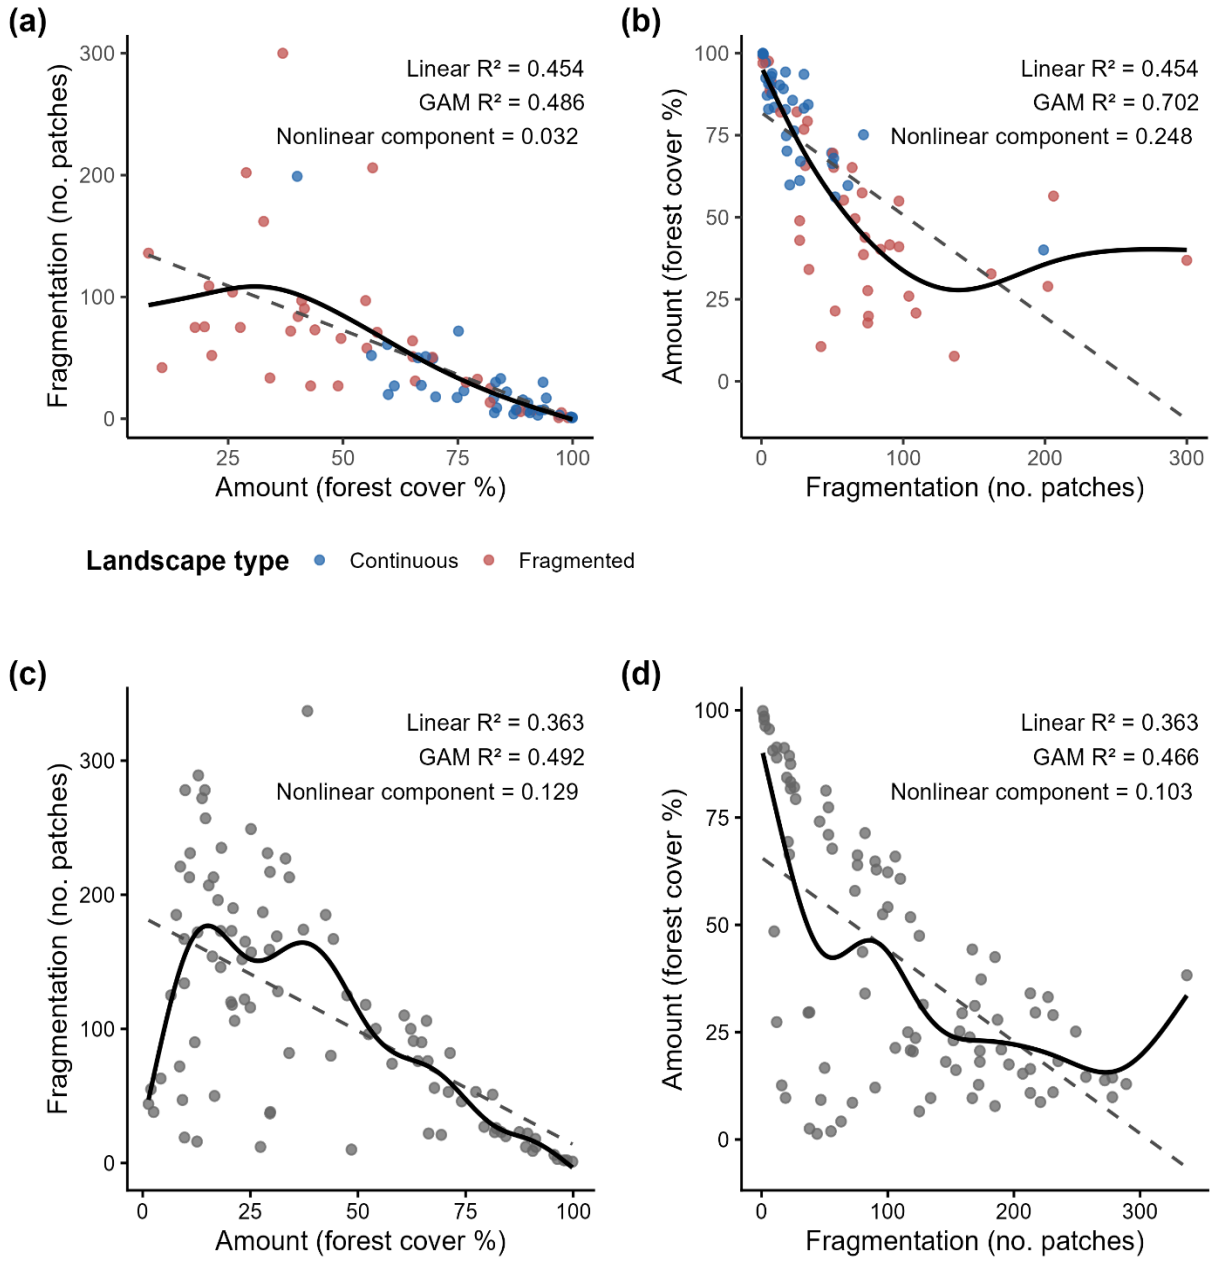


Figure 3.2. Realised habitat amount–patch number geometry in the two diagnostic case studies. Panels show bidirectional nonlinear coupling between habitat amount and patch number. In the GS25–F26 dataset, panels (a) and (b) show patch number as a function of forest cover and forest cover as a function of patch number, respectively, with points coloured by landscape type (blue, continuous; red, fragmented). In the K26 dataset, panels (c) and (d) show the corresponding bidirectional coupling across Atlantic Forest landscapes. Dashed lines show linear fits and solid lines show GAM smooths; text annotations report the linear R^2, the GAM R^2, and the nonlinear

component, defined as GAM R^2 minus linear R^2. The comparison illustrates that conventional linear collinearity diagnostics can understate realised predictor coupling, because habitat amount and patch-count subdivision occupy a bounded, nonlinear predictor space rather than an independent two-dimensional surface. Figure adapted from Martínez-Lanfranco (2026a,d). The smooths are shown to characterise the shape of the realised predictor space rather than as fitted estimates for interpretation.

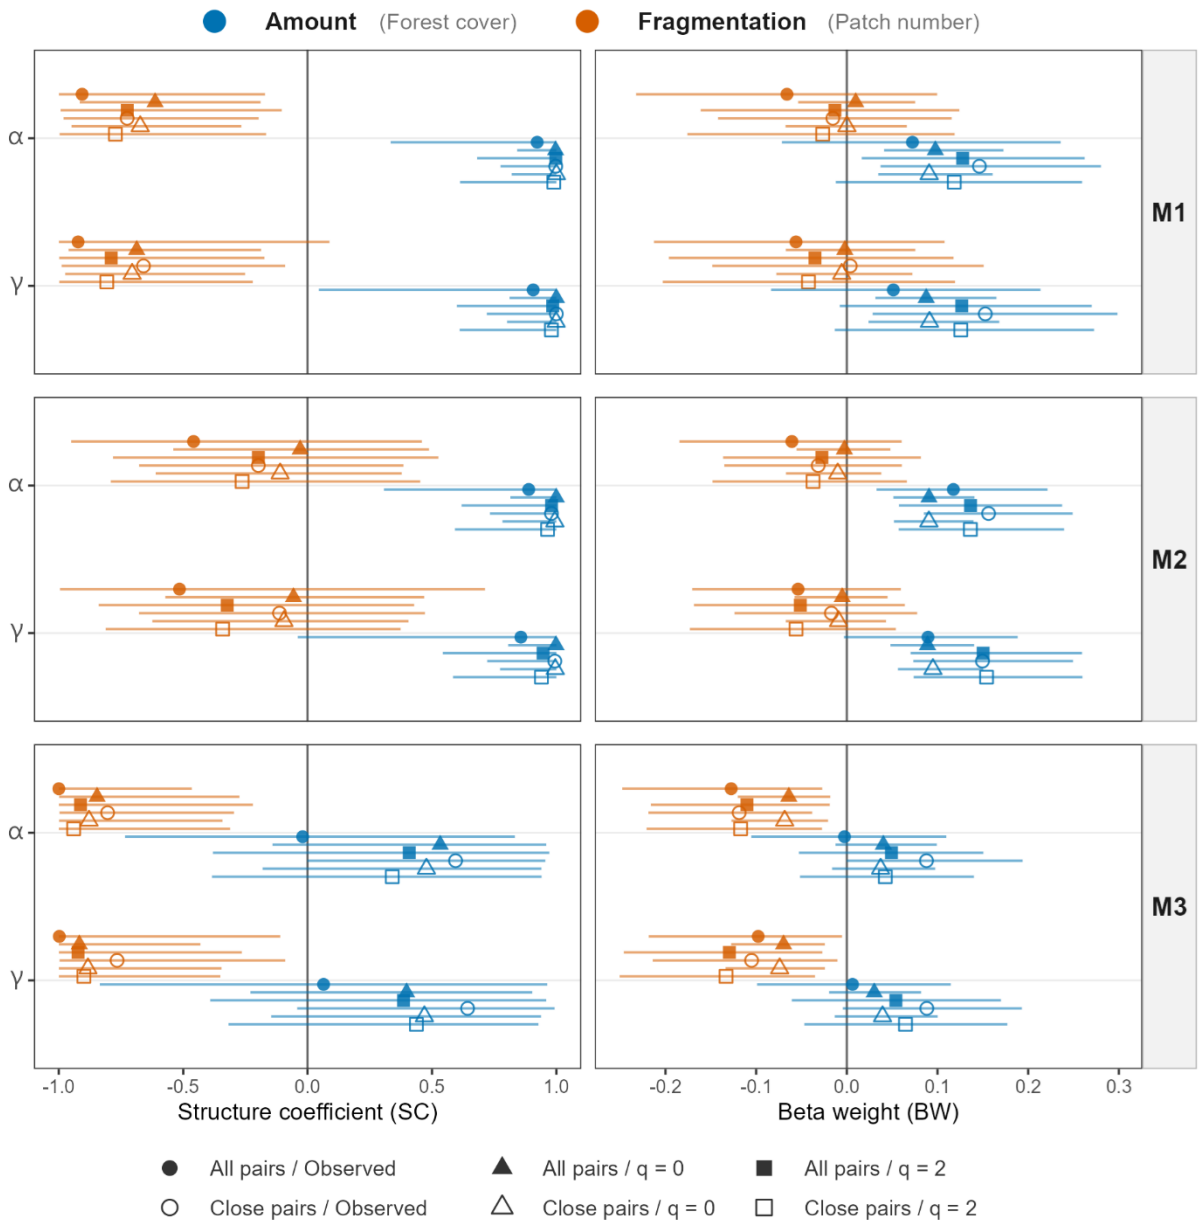


Figure 3.3. Coefficient architecture of habitat amount and patch number across 18 analytical variants (α and γ diversity) in the GS25–F26 case. Each row block is one parameterisation (M1–M3). The left column plots structure coefficients (SC) and the right column standardised beta weights (BW) for patch number (orange) and habitat amount (forest cover, blue). Point shapes encode the pairing design and diversity metric (see key); horizontal lines are 95% bootstrap confidence intervals. The diagnostic contrasts are: under M1, the patch-number structure coefficient is large and negative across all variants while its beta weight stays near zero — the suppressor signature, geometrically invariant rather than variant-specific; under M3, structure

coefficient and beta weight align in sign for patch number, indicating gradient reallocation rather than a sampling artefact. The habitat-amount beta weight is positive across all variants and parameterisations and shrinks under M3 as shared variance is redistributed to the patch-number predictor. Semi-partial R^2 and the full 18-variant outputs, including β diversity, are deposited in the Zenodo archive accompanying Martínez-Lanfranco (2026b).

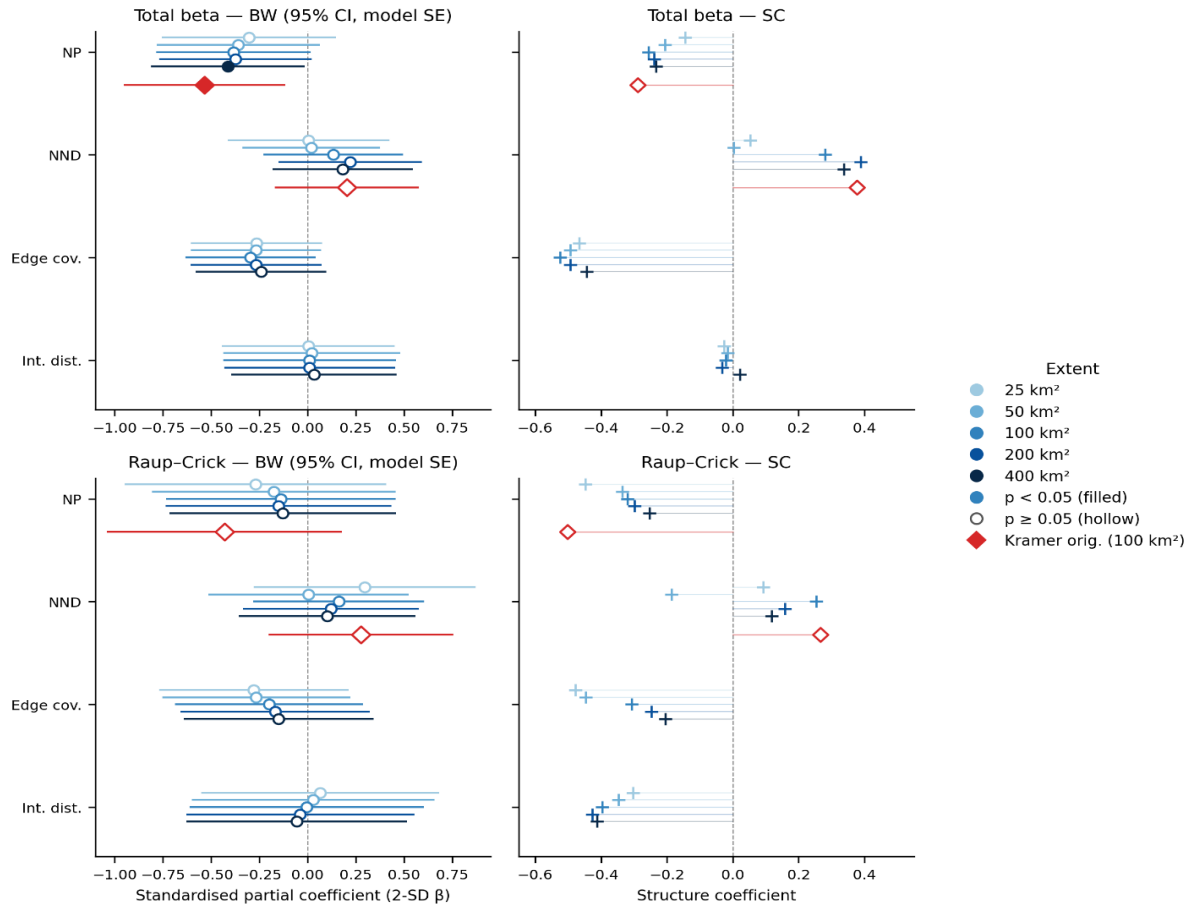


Figure 3.4. Mechanism triangulation for fragmentation-per-se inference from beta diversity in the K26 Atlantic Forest dataset. Each row is one response: total beta diversity (top) and Raup-Crick dissimilarity (bottom). Left panels show standardised partial coefficients (BW, 2-SD β, with 95% confidence intervals from the model standard error); right panels show structure coefficients (SC). Both are given for a focal subset of mechanism-congruent metrics: patch number (NP), nearest-neighbour distance (NND), edge coverage (Edge cov.), and interior distance (Int. dist.) across five operational extents (25–400 km², shaded light to dark). In the coefficient panels, filled symbols denote $p < 0.05$ and open symbols $p \geq 0.05$. Red diamonds show K26's original deposited predictors refit at 100 km². SC is the descriptive correlation between each standardised predictor and the model's fitted linear predictor — an index of fitted-gradient alignment, not a hypothesis test or a mechanism ranking. Patch-number support is confined to a narrow set of extents and is not matched by the mechanism-specific metrics, none of which recovers an equivalent signal at any extent for either response. Figure adapted from Martínez-Lanfranco (2026d).

Table 3.1. Cross-case diagnostic synthesis. Summary of the four conditions required to identify a fragmentation-per-se effect and their diagnostic outcomes in the GS25–F26 and K26 cases — predictor-space geometry, coefficient architecture, response–estimand alignment, and metric–mechanism correspondence. The table maps the chapter's central argument: an additive patch-number coefficient alone does not identify fragmentation per se, and its interpretation depends on convergence across the four diagnostic conditions.

Diagnostic readout	**Case 1 — GS25–F26 (alpha and gamma diversity)**	**Case 2 — K26 (beta diversity)**	**Cross-case inference**
Predictor-space adequacy	Fails: 31/37 pairs re-express the original directional contrast; same-amount contrasts structurally scarce (3/37)	Fails: nonlinear threshold geometry; high-cover, high-patch cell structurally empty	Linear collinearity screens are insufficient for predictor separability
Coefficient architecture	Suppressor attenuation: large SC, near-zero BW and R^2sp (~8,800-fold ratio)	Partial-slope amplification: near-zero marginal r → significant partial β. Structure coefficients show NP is real for total beta, while spatial distance dominates fitted-gradient alignment under Raup-Crick.	Opposite pathologies from the same metric family: sign and failure mode are orthogonal

Diagnostic readout	Case 1 — GS25–F26 (alpha and gamma diversity)	Case 2 — K26 (beta diversity)	Cross-case inference
Response-estimand alignment	α/γ richness: suppressor signature present. β diversity: absent — governed by spatial distance	Total beta: NP supported and directionally reproduced. Raup-Crick: HC and NP lose support, while spatial distance remains. K2/K3 show that residual NP persists only for total beta, not under Raup-Crick or reciprocal residualisation.	In these cases, beta diversity tracks spatial and alpha-diversity structure more strongly than an independently identified configuration gradient
Metric-mechanism correspondence	Patch number = subdivision metric, not generic fragmentation. Algebraically coupled to amount	Independent NP directionally recovers the total-beta signal, but NND, edge exposure, connectivity, interior-depth, and amount-decoupled arrangement metrics do not recover convergent support across tested predictor extents. PD is a subdivision-rescaling	Subdivision, isolation, edge exposure, connectivity, spatial arrangement, and spatial distance are not interchangeable. Fragmentation-per-se inference requires the metric to match the hypothesised mechanism.

Diagnostic readout	Case 1 — GS25–F26 (alpha and gamma diversity)	Case 2 — K26 (beta diversity)	Cross-case inference
		sensitivity, not independent mechanism-matched evidence. The supported signal is response-specific patch-count subdivision, not a convergent fragmentation mechanism.	
Domain stability	≥50% cover and Tropical SA: suppressor structure amplified, recoverable negative direction	Low-cover NP signal compatible with threshold expectations but remains total-beta-specific; ≥50% cover: HC and NP collapse, spatial distance persists	Domain checks locate where coefficient signals arise; they do not validate mechanism alone. Low-cover localisation can be threshold-compatible, but response-estimand and metric-mechanism checks remain necessary.

CHAPTER 4

Reframing the habitat fragmentation debate around the inferential targets predicted by ecological theory

1. The debate and its structure

Research on habitat fragmentation has produced one of the most persistent disagreements in landscape ecology and conservation biology (Miller-Rushing et al. 2019). Some studies report negative effects of spatial subdivision independent of habitat loss. Others find weak, neutral, or positive associations (Fletcher et al. 2018; Fahrig et al. 2019). Disagreement has been attributed to differences in spatial scale and study design, habitat definition, fragmentation metrics, and the difficulty of separating configuration from amount statistically in empirical landscapes (Ewers and Didham 2006; Didham et al. 2012; Hadley and Betts 2016; Fletcher et al. 2023). These concerns explain why empirical studies disagree but have not supplied a common inferential target for resolving the disagreement. The debate has therefore remained caught in a demonstration–counter-demonstration cycle for more than three decades (Andrén 1994; Valente et al. 2023a). The same global multi-taxa dataset has recently supported opposed conclusions under different implementations of amount-controlled fragmentation inference (Gonçalves-Souza et al. 2025; Fahrig et al. 2026).

These methodological diagnoses share a deeper foundation. They were formulated within a fragmentation-per-se paradigm (Fahrig 2003; Fahrig et al. 2019) that clarified the target of inference but also directed empirical practice toward amount-controlled contrasts among present-day landscapes. Fragmentation as a process is the breaking apart of habitat through time, whereas fragmentation as pattern is the spatial configuration of the remaining habitat mosaic at a given time (Fahrig 2003; Fletcher et al. 2026). Fragmentation per se is the independent effect of spatial configuration on biodiversity at a given habitat amount, defined by comparisons among landscapes that differ in arrangement but not in habitat amount (Fahrig 2003, 2013, 2017; Fahrig et al. 2019). The paradigm therefore improved conceptual precision while leaving unresolved whether the amount-controlled contrast the field adopted is the right inferential target.

In empirical landscapes shaped by habitat loss, the processes that generate these patterns couple habitat amount and configuration along the same gradient, making snapshot contrasts difficult to align with an independent fragmentation effect (Didham et al. 2012; Fletcher et al. 2018). Amount-controlled contrasts address that coupling, but the additive structure they impose constrains inference to a single configuration effect evaluated at the sample mean of habitat amount. Current empirical practice therefore estimates that quantity (Smith et al. 2009; Fahrig et al. 2026), one defined narrowly relative to what competing ecological theories predict.

1.1 The ecological conditionality problem

The conditional nature of the theoretical question has been explicit since the debate's early formulation. Fragmentation effects were already shown to emerge mainly below a cover threshold of approximately 10–30%, implying that the ecological consequence of configuration depends on where a landscape sits along the habitat-amount axis (Andrén 1994). The statistical implication was stated directly when Fahrig (2003) argued that the threshold "should result in a statistical interaction effect between habitat amount and habitat fragmentation per se." Configuration effects were later predicted to be most detectable at intermediate cover (Villard and Metzger 2014). Competitive metacommunity models identified a sign-reversal mechanism in which the fragmentation effect is positive at high cover and negative at low cover (Rybicki et al. 2020; Zhang et al. 2024). Extinction filtering predicts fading fragmentation effects as cover declines and sensitive species are lost (Betts et al. 2019). These predictions converged across three decades of distinct theoretical traditions, yet their analytical implication has not been adopted as standard practice (McGarigal and Cushman 2002; Hadley and Betts 2016; Riva et al. 2024).

The HAH is the limiting exception. It predicts that the fragmentation effect is invariant once habitat amount is controlled (Fahrig 2013, 2021; Saura 2021). The other frameworks considered here predict that fragmentation sensitivity changes along the habitat-amount gradient and may reverse sign, with the form of the change differing among mechanisms. A single average fragmentation effect cannot distinguish these alternatives.

The disagreement is therefore not simply over whether configuration matters, but over which conditional response surface best characterizes a given system. When the fragmentation effect changes along the gradient, an additive estimate becomes partly a function of where the

sampled landscapes happen to sit rather than of the underlying ecology alone. The same conditional surface can therefore yield different average fragmentation effects in studies positioned at different gradient regions (Smith et al. 2009; Duncan and Kefford 2021; Spake et al. 2023).

1.2 A response-surface reframing

Progress requires identifying the inferential target the debate has been trying to characterize. That target is the conditional biodiversity response surface. Recent calls for reconciliation have identified the absence of a common conceptual framework as a key impediment to moving from demonstration to investigation (Valente et al. 2023a), while earlier gradient-based work showed that habitat configuration should matter most only in particular regions of the habitat-amount gradient (Villard and Metzger 2014; Zhang et al. 2024). The missing step has been to make habitat amount not one moderator among many, but the organizing axis along which the competing frameworks make distinguishable predictions. To our knowledge, existing syntheses have not jointly organized these frameworks around a common conditional target, derived their gradient-specific predictions, and connected those predictions to the statistical structure required to distinguish them empirically. They differ in whether configuration ever becomes independently consequential along the habitat-amount gradient, and in where and how that conditionality is biologically expressed. These alternatives are ecological hypotheses before they are statistical contrasts, and their implied response surfaces must be specified before model coefficients can resolve them.

The habitat-amount gradient can be divided into response-specific regions of persistence, configuration sensitivity, and empirical recoverability. Empirical threshold syntheses showed that configuration effects often become detectable only below particular cover ranges, with the numerical position varying among systems (Andrén 1994). Spatially explicit theory showed that habitat loss can alter persistence through extinction thresholds and connectivity loss (Fahrig 2002). Percolation models provide a structural reference point because the critical percolation threshold, p_c, marks the cover value at which habitat cells first form a connected cluster under a specified landscape-generation model, although p_c is not itself a biological threshold (Gardner et al. 1987; With and Crist 1995). For any response state variable R, ET_R denotes the lower persistence boundary, and FT_R denotes the upper configuration-saturation boundary above which

configuration no longer limits the response. The interval between ET_R and FT_R is the response-specific configuration-sensitive window (W_R) in which configuration effects can be biologically expressed (Fig. 4.1). Empirical recovery additionally requires sampled landscapes within W_R to contain sufficient configuration contrast. This sampled portion of W_R is the recoverability window. These boundaries vary with response variable, taxon, spatial scale, matrix, configuration metric, and landscape-generation process.

This notation previews the four frameworks developed in Section 2 and the limiting-case predictions (Fig. 4.1; Table 4.1). These schematic surfaces are canonical cases for comparison, not discrete categories into which real systems must fall. Real datasets may express different signatures across taxa, configuration metrics, response variables, and gradient positions within the same landscape sample. The corresponding response-surface and marginal-effect predictions are formalized in Section 3 and visualized in Fig. 4.2.

This reframing changes the object of the debate. Except under the HAH, fragmentation is not predicted to have one gradient-invariant effect once habitat amount is controlled. It is predicted to have effects that appear, reverse, or fade depending on where taxa and landscapes sit along the habitat-amount gradient, and on which configuration metrics and response variables are analyzed. The task is therefore not to assess whether fragmentation effects are positive, negative, or null, but to identify which conditional surface a system expresses and where that surface can be recovered.

2. Ecological mechanisms generating conditional fragmentation effects

The conditional framework applies to any ecological response variable Y for which a fragmentation metric F and a cover measure A can be jointly estimated at a biologically meaningful spatial scale. "Cover" denotes the landscape-level measurement used to quantify A; "habitat" refers to its species-specific ecological sense, because a cover type constitutes habitat only for species whose requirements it meets (Fahrig 2013).

Four main theoretical frameworks have organized much of the fragmentation debate and define the limiting cases considered here (Table 4.1; Fig. 4.1). Each proposes a distinct ecological mechanism and a distinct prediction about how fragmentation effects change along the habitat-amount gradient. All predict conditional fragmentation effects except the HAH,

which predicts invariance once habitat amount is controlled. The frameworks disagree about the direction, shape, and gradient position of that conditionality.

2.1 Habitat amount hypothesis and configuration invariance

Fahrig (2013, 2015) proposed that biodiversity at a sample site responds primarily to the total amount of habitat within the local landscape surrounding that site, with spatial configuration exerting little independent influence once amount is controlled. The local landscape is defined empirically at the radius at which habitat amount most strongly predicts the response, and this scale must be estimated rather than assumed a priori. The mechanism applies at any level of organization because local habitat amount determines the species pool available regardless of spatial arrangement. Under HAH, the fragmentation effect is zero, or near zero, everywhere along the gradient. HAH is the amount-only limiting case in Fig. 4.1. An extinction threshold may arise from habitat loss, but no fragmentation threshold is expected. Saura (2021) formalized the HAH in a way that corresponds here to the joint restriction $\beta 1 = 0$ and $\beta 3 = 0$ in the interaction framework. This makes the HAH a specific falsifiable case rather than a default assumption. Evidence bearing on this broader invariance claim is mixed. A vote-counting review of 381 effects across 118 studies found no subset of conditions with predominantly negative responses (Fahrig 2017), and a model-based synthesis of 35 species-density studies found non-significant interactions between habitat amount and patch size, isolation, or fragmentation (Watling et al. 2020; De Camargo et al. 2018). Controlled experiments, by contrast, found steeper species–area relationships in isolated than in connected treatments at fixed habitat amount, with the divergence widening over 15 years (Haddad et al. 2017a). These results address different aspects of the HAH prediction set and are not automatically commensurable as tests of a single estimand. Because the HAH is defined relative to the local landscape, its predictions are scale dependent and most directly tested when A is measured at the empirically determined scale of effect for the focal taxon (Jackson and Fahrig 2015).

2.2 Threshold and scarcity models below cover thresholds

Andrén (1994) synthesized bird and mammal studies across a cover gradient and proposed that fragmentation effects intensify specifically when landscape cover falls below approximately 10–30%. Above the threshold, habitats remain structurally connected and subdivision imposes little additional cost. Below it, patches become small, isolated, and edge-

dominated, and spatial configuration amplifies extinction risk beyond what habitat loss alone would predict. The connectivity threshold has a formal theoretical basis in spatially explicit metapopulation dynamics, where patch isolation drives disproportionate extinction at low habitat amounts (Hanski 2015; Fahrig 2002). The threshold prediction was later translated into observational study design (Fahrig 1998; Villard and Metzger 2014). The 10–30% range is not fixed. Percolation theory predicts structural disconnection at approximately 59% cover under random habitat loss, but empirically observed thresholds are lower because land clearing is non-random, producing more spatially aggregated loss patterns than random models assume (Gardner et al. 1987; With and Crist 1995).

Threshold-like fragmentation responses arise through multiple pathways, including connectivity breakdown, edge proliferation, interior habitat collapse, and local patch-viability failure. Their relative importance depends on species traits and the habitat-amount range studied (Villard and Metzger 2014). Threshold dynamics should therefore be read as a family of structural and demographic breakpoints rather than a single universal low-cover rule. Fragmentation per se can raise ET_R by reducing colonization rates or increasing edge-mediated matrix mortality (Fahrig 2002, 2003), and p_c provides a structural reference point that can help determine FT_R but is not equivalent to it. Fahrig (2003) made the statistical implication explicit. If threshold dynamics operate, the effect of fragmentation per se should depend on habitat amount, producing a habitat amount × fragmentation interaction. Smith et al. (2011) report this directional signature on fragmentation-sensitive bird richness across five landscape radii, with the fragmentation effect intensifying at low cover as threshold dynamics predict. This places threshold/scarcity dynamics in the lower portion of the configuration-sensitive window, where configuration effects are expected before persistence failure truncates the response (Fig. 4.1).

2.3 Metacommunity coexistence and sign reversal across the gradient

A mechanistic explanation for positive fragmentation effects was formalized by Rybicki et al. (2020). At high habitat amounts, fragmentation can promote coexistence by breaking competitive dominance. Large patches support strong competitors, small patches support superior colonizers, and the landscape mosaic maintains both. The mechanism is a colonization–competition coexistence (C–C) tradeoff. As habitat amount declines, isolation overwhelms the coexistence-promoting effect and fragmentation becomes detrimental. Zhang et al. (2024)

identified slope reversals consistent with this prediction across 71 systems and five taxonomic groups (Riva and Fahrig 2022, 2023). These reversals do not require positive and negative effects to have symmetric magnitudes. The crossover should shift with species traits, dispersal, mortality, competition strength, and the sampled habitat-amount domain (Rybicki et al. 2020; Zhang et al. 2024). Whether these reversals reflect C–C or another mechanism generating similar richness patterns depends on beta-diversity decomposition and assemblage composition history. The fragmentation effect is positive or neutral at high cover and negative at low cover. This sign-changing response occurs within the configuration-sensitive domain (Fig. 4.1). Because the mechanism operates through compositional differentiation among patches, the C–C tradeoff also predicts higher turnover in fragmented landscapes at intermediate-to-high cover, declining at low cover as competitive dominants are lost.

2.4 Extinction filtering and fading configuration effects

In landscapes with long histories of habitat loss, fragmentation-sensitive species may already have been progressively removed from assemblages by prior extinction, leaving communities dominated by matrix-tolerant taxa that do not respond strongly to configuration regardless of cover level (Betts et al. 2019). Extinction filtering is the community-level consequence of sensitive taxa progressively crossing species-specific ET_R values as habitat amount declines. Fragmentation per se can raise individual ET_R values and accelerate the process, but the trajectory is driven primarily by habitat loss (Fahrig 2002; Villard and Metzger 2014). The taxa remaining after filtering often have broad habitat associations and high dispersal capacity, conferring resilience to fragmentation effects (Henle et al. 2004; Weeks et al. 2023; Püttker et al. 2020). Extinction filtering is therefore a historical corollary describing how prior compositional pruning alters the detectability of configuration effects in the present (Balmford 1996). This interpretation assumes that habitat associations remain stable enough for lost taxa to be identified as fragmentation-sensitive from traits or history, and extinction debt can attenuate filtering signals where sensitive species have not yet disappeared (Kuussaari et al. 2009). The conditional relationship is one in which fragmentation effects are most negative at high cover, where sensitive species remain, and weaken toward zero as those species are progressively lost. Null results from heavily deforested systems should not be interpreted as ecological invariance without first characterizing assemblage filtering history. The beta-diversity signature is nestedness, meaning ordered species loss, rather than turnover. This places extinction filtering

toward the higher-cover portion of the recoverable domain, where sensitive taxa remain, with the signal fading as species-specific extinction thresholds are crossed (Fig. 4.1).

In the statistical framework of Section 3, filtering corresponds to a negative fragmentation × habitat amount interaction when biodiversity is the response and higher F values denote greater fragmentation per se. The fragmentation effect should be most negative where habitat amount remains high enough for sensitive taxa to persist, and should weaken toward zero as habitat loss progressively filters those taxa from the assemblage. Across 19,432 vertebrate species, the risk consequences of recent forest loss were disproportionately strong in landscapes retaining high initial forest cover (Betts et al. 2017). Historical anthropogenic deforestation further attenuates impacts of current land-use change, with longer deforestation histories associated with trait profiles shifted toward greater deforestation tolerance (Hua et al. 2024).

2.5 Shared conceptual structure and mixed regimes

Despite their ecological differences, all four frameworks can be characterized by asking how the configuration effect behaves along the habitat-amount gradient. Under the HAH, it stays near zero across the entire gradient. Under threshold/scarcity dynamics, it becomes harmful within W_R after cover falls below FT_R but before ET_R truncates the response. Under C–C coexistence, it switches sign within the coexistence portion of W_R. Under extinction filtering, it is strongest where sensitive taxa remain and fades as species-specific ET_R crossings accumulate. The fragmentation debate can therefore be reframed as a question about how the configuration effect varies with habitat amount. Each hypothesis yields a different response-surface geometry relating ecological context to the biodiversity consequences of spatial subdivision (Fig. 4.2).

These regimes are not necessarily mutually exclusive in real landscapes. A system with a long history of fragmentation may show extinction-filtering signatures at mid-to-high cover while threshold dynamics operate at low cover among the remaining habitat-sensitive taxa. Table 4.1 should therefore be read as a set of limiting cases rather than discrete categories. Mixed regimes are expected when mechanisms operate in different portions of the cover gradient or on different response components. Threshold dynamics may persist among remaining sensitive taxa while extinction filtering simultaneously weakens the aggregate signal. The C–C tradeoff may generate turnover and positive richness effects where competitive structure remains intact, while filtering progressively removes the taxa carrying that signal at lower cover. Identifying the

dominant process requires explicit evaluation of the response variable, the beta-diversity decomposition, assemblage history, and the spatial scale at which taxa perceive habitat.

Response variables differ in what they can diagnose. Population-level responses can show whether configuration affects persistence, colonization, extinction, or demography, but they cannot by themselves separate mechanisms that depend on community responses such as the C–C tradeoff or extinction filtering. Community-level responses such as species richness, occupancy, or aggregate diversity indices can reveal mechanism signatures, but only when decomposed. Total richness can be invariant while specialist composition or beta-diversity components are not. The richness response combines turnover and nestedness contributions that the four frameworks separate (Anderson et al. 2011; Baselga 2010; Socolar et al. 2016). Tests of these frameworks at the community level should report richness alongside the decomposed beta-diversity components, and conclusions should be reported as response-specific, scale-specific, and metric-specific rather than as global statements about fragmentation effects on biodiversity. For single-species responses, C–C has no direct analogue as formulated, although positive edge effects, competitive release, resource complementation, connectivity and source–sink dynamics, or response truncation can generate similar response-surface geometries at the population level (Fahrig et al. 2019).

3. From ecological mechanism to statistical identifiability

The competing frameworks imply different response-surface geometries. The statistical question is whether those geometries can be distinguished in data. Ecological theory specifies the form of the conditional relationship; statistical models determine whether that relationship can be evaluated empirically. A coefficient can be statistically valid while still failing to represent the ecological effect the analysis intends to measure. Defining the target quantity is what connects statistical evidence to theory (Lundberg et al. 2021). When theory predicts that the fragmentation effect depends on habitat amount, empirical tests require models that allow the slope of fragmentation to vary along the gradient. Models that constrain this variation impose the invariance component of the HAH without evaluating it against the alternatives. The interaction between fragmentation and habitat amount is therefore the minimum linear statistical structure required to distinguish among the competing predictions (Fahrig 2003; Saura 2021; Spake et al. 2023).

3.1 Model structures

Let Y be any ecological response variable, such as species richness, diversity index, occupancy, or abundance, F a fragmentation metric, A habitat amount, and ε error. In what follows, A is expressed as deviation from its sample mean, so that A = 0 denotes typical habitat amount in the sampled region.

Here, F refers to fragmentation per se, that is, the effect of habitat configuration considered independently of habitat amount (Fahrig 2003; Riva et al. 2024). Higher values of F indicate greater fragmentation per se. Analysis using aggregation-based metrics where higher values indicate less fragmentation should invert the metric before fitting for full alignment with Table 4.1 and Fig. 4.2. Without inversion, the predicted signs of β1 and β3 are reversed relative to those displays, and the qualitative descriptions in Table 4.1 and the curve directions in Fig. 4.2 apply to the low end of the metric rather than the high end.

The inferential target is not a single fragmentation coefficient but the slope of biodiversity with respect to fragmentation, evaluated as a function of habitat amount. In response-surface notation, this is $\partial Y/\partial F = f(A)$. The HAH is the limiting case in which this function is approximately zero, once habitat amount is controlled. Other frameworks predict non-zero functions whose sign or magnitude changes along the gradient.

The terms “marginal” and “conditional” are used here in an estimand sense. The marginal effect is the local slope $\partial Y/\partial F$ evaluated at a value of A, and the effect is conditional because that slope is evaluated along the habitat-amount gradient, f(A). This usage is distinct from the marginal–conditional distinction in mixed-model prediction, where the contrast concerns whether predictions are averaged over or conditioned on random effects.

The dominant empirical model is the additive specification.

Model 1 (additive). $Y = \beta0 + \beta1 \cdot F + \beta2 \cdot A + \varepsilon$

The minimum linear conditional model adds the fragmentation × habitat amount interaction.

Model 2 (interaction). $Y = \beta0 + \beta1 \cdot F + \beta2 \cdot A + \beta3 \cdot (F \times A) + \varepsilon$

These coefficients correspond, respectively, to the fragmentation effect (β1), the habitat amount effect (β2), and the fragmentation × habitat amount interaction (β3). We retain the

numerical notation used throughout the debate literature (Fahrig 2003; Smith et al. 2009; Koper et al. 2007). The habitat amount effect β2 is included to statistically partial out habitat amount from the fragmentation coefficient and is not interpreted further. β1 denotes the estimated fragmentation effect at mean habitat amount, the reference point used throughout Table 4.1 and Fig. 4.2.

In Model 1, the fragmentation effect is $\partial Y/\partial F = \beta 1$, constant and independent of A. In Model 2, it is $\partial Y/\partial F = \beta 1 + \beta 3 \cdot A$, a function of habitat amount. Model 2 is therefore the linear implementation of the general statement $\partial Y/\partial F = f(A)$, with $f(A) = \beta 1 + \beta 3 \cdot A$. On the linear predictor scale, β3 estimates the linear rate at which the fragmentation slope changes with habitat amount, formally the cross-partial derivative $\partial^2 Y/\partial F \partial A$. If the true f(A) is nonlinear, threshold-shaped, or varies among taxa or systems, β3 is a first-order summary of conditionality rather than a complete description of the response surface.

β1 and β3 have direct ecological readings. β1 is the fragmentation effect at mean habitat amount in the sampled dataset. β3 describes how that effect changes as cover increases or decreases from that baseline. A positive β3 means the fragmentation effect becomes more positive as cover increases, or more negative as cover declines. This direction is consistent with threshold/scarcity when fragmentation is harmful at low cover and with C–C tradeoff when the effect crosses from negative at low cover to neutral or positive at higher cover. A negative β3 means the fragmentation effect becomes more negative as cover increases, as predicted when sensitive taxa remain in higher-cover landscapes and the signal fades after filtering at lower cover. β3 equal to zero means the fragmentation effect is constant across the habitat-amount gradient. This is the invariance component of the HAH; the HAH further predicts that this constant effect is zero, or near zero, once habitat amount is controlled.

These two quantities are most directly communicated through the paired visualization in Fig. 4.2. Predicted biodiversity at representative cover levels on the response scale provides the ecological entry point, showing whether fragmentation is harmful, neutral, or beneficial at low, intermediate, and high cover. The marginal fragmentation effect plotted against habitat amount on the model scale makes $\beta 1 + \beta 3A$ the primary visual object rather than an inference from comparing slopes across panels (Duncan and Kefford 2021). Given the recoverability conditions in Fig. 4.1, the four canonical frameworks generate distinct limiting-case curves. The HAH

predicts a horizontal line at zero. Threshold dynamics predict a curve that is most negative at low cover and weakens toward zero as cover increases. C–C tradeoff predicts a curve that crosses zero at the coexistence crossover and becomes increasingly positive at higher cover. Extinction filtering predicts a curve that intensifies toward more negative values as cover increases. Treating habitat amount as the moderator of the fragmentation effect reflects the causal structure of landscape change, because habitat loss generates configuration. The fragmentation effect conditional on habitat amount is therefore the ecologically interpretable quantity (Ruffell et al. 2016; Spake et al. 2023). In models with non-identity links the two panels are not numerically equivalent. The predicted-biodiversity panel shows response-scale values, and the marginal-effect panel remains on the model scale where $\beta1 + \beta3A$ is defined. For empirical applications using generalized models, back-transformed predictions are appropriate for the conditional plot while the marginal effect stays on the link scale (Spake et al. 2023; Duncan and Kefford 2021).

Model 1 is Model 2 under the restriction $\beta3 = 0$, but it does not test whether this restriction holds. It imposes the restriction structurally. Fitting Model 1 assumes fragmentation effects are invariant across the cover gradient. This is the specific invariance prediction that the HAH makes and that all other frameworks reject. It is not equivalent to adopting the HAH in full. The HAH also predicts $\beta1 = 0$, whereas an additive model with significant $\beta1$ imposes invariance while reporting a non-zero average effect whose value depends on where the sample sits on the gradient. Testing whether $\beta3 = 0$ requires fitting Model 2 and comparing it to Model 1 via likelihood ratio test or information criterion.

Models 1 and 2 are written here in their simplest form to clarify the inferential target. In empirical applications, this structure may be implemented within generalized, hierarchical, multivariate community, spatial, or temporal models as required by the response variable's distribution and the study's sampling design (Pedersen et al. 2019; Tredennick et al. 2021; Ovaskainen and Soininen 2011; Popovic et al. 2024). The conditional fragmentation effect of interest remains $\partial Y/\partial F = f(A)$. The interaction term is its minimal linear parameterization. In all such implementations, the testable prediction is whether the fragmentation slope varies systematically along the habitat-amount gradient.

3.2 Conditions for β3 interpretability

The interaction coefficient β3 is interpretable as a test of the competing theoretical predictions only when several prior conditions are satisfied, operating at different inferential stages. These conditions are empirical and dataset-specific. Their violation cannot be corrected post hoc by model specification alone. Where geometry conditions are not met, the interaction term parameterizes variation along a shared habitat-loss gradient rather than isolating a configuration effect, and may therefore reproduce the same inferential limitations as the additive model in a more complex form.

Design conditions are established before data collection. First, the habitat layer underlying both F and A must be ecologically appropriate to the focal taxon or assemblage. A cover type constitutes habitat only for taxa whose biological requirements it meets, and using generic land-cover classifications as proxies introduces systematic error into both predictors before any model is estimated (Betts et al. 2014; Halstead et al. 2019; Almeida-Gomes et al. 2016; Dennis et al. 2026). For forest-interior taxa, breaking apart of forest at constant forest cover reduces interior habitat amount, while for forest-edge taxa the same configuration change increases edge habitat amount, so fragmentation manipulations holding generic land cover constant do not hold species-specific habitat amount constant (Fahrig 1999). Second, F, A, and Y must be measured at the scale at which the taxon is sensitive to configuration; the scale of effect is frequently outside the range researchers evaluate, producing null results attributable to scale mismatch rather than ecological invariance (Jackson and Fahrig 2015). Third, the fragmentation metric must track the ecological property relevant to the mechanism under investigation, namely aggregation or isolation of correctly defined habitat (Neel et al. 2004; Fletcher et al. 2023).

Metric–mechanism correspondence. Species- and population-level fragmentation effects commonly operate through edge-mediated changes in local conditions or movement-mediated changes in colonization or immigration, which are captured by different metrics and may activate at different positions along the habitat-amount gradient (Fletcher et al. 2016). A connectivity metric can show near-zero F effects where the binding constraint is local habitat sufficiency rather than inter-patch movement, even when the canonical conditional surface predicts strong negative effects at low cover (Villard and Metzger 2014; Fletcher et al. 2016). Connectivity metrics that embed patch area in their formulation (Hanski-type indices) compound this problem because they suppress connectivity values at low A independently of movement biology, partially confounding the metric with the habitat-amount predictor (Fahrig et al. 2021). At the

community level, no single fragmentation metric is likely to capture all operative mechanisms simultaneously, so trait-informed or mechanism-specific complementary metrics may improve inference (Fahrig et al. 2011).

Geometry conditions are evaluable once the dataset is assembled. Habitat amount and configuration must be sufficiently separable in the realized predictor space for the interaction coefficient to identify a conditional configuration effect rather than a position along a shared habitat-loss gradient. When F and A are co-generated along a directed, asymmetric landscape gradient before any landscape is sampled (Didham et al. 2012), additive regression yields statistically valid partial coefficients for the quantities it defines, but those quantities need not correspond to the intended ecological contrasts (Koper et al. 2007; Smith et al. 2009). Low pairwise correlations or acceptable variance-inflation factors do not by themselves establish separability when habitat amount and configuration are generated nonlinearly along the same landscape gradient (Dormann et al. 2013; Martínez-Lanfranco 2026). Additive models recover direct effects rather than total effects. Where habitat loss operates partly through configuration as an indirect pathway, the partial configuration coefficient underestimates the total ecological effect of fragmentation (Ruffell et al. 2016; Püttker et al. 2020).

Configuration variance must also exist across the sampled gradient. Fragmentation metrics exhibit constrained variability near the gradient extremes, where few configurations are geometrically possible. Percolation theory makes this mechanism explicit because the set of possible configurations changes nonlinearly with habitat amount (Gardner et al. 1987; With and Crist 1995; Neel et al. 2004; Fletcher et al. 2023). These geometry constraints correspond to the response-specific recoverability window summarized in Fig. 4.1. Metric choice must also be evaluated against the ecological process tested, because interventions that increase functional connectivity can be scored by simple configuration metrics as increased fragmentation even when they improve habitat reachability (Villard and Metzger 2014).

The specification condition operates at the level of model form. The linear interaction model is interpretable as a first-order summary only when curvature in the amount response, the fragmentation response, and the bivariate A–F surface has been evaluated. If either main-effect relationship is nonlinear, residual curvature can be absorbed by β3, biasing the interaction estimate independently of genuine fragmentation conditionality (Pedersen et al. 2019;

Tredennick et al. 2021). Diagnostic flexible models, including univariate smooths of Y against A and F and a bivariate smooth of Y across the A–F plane, can identify whether the linear interaction adequately summarizes the conditional surface before β3 is interpreted mechanistically (Pedersen et al. 2019; Duncan and Kefford 2021).

3.3 The additive default as rational origin and present insufficiency

Each study that fits an additive fragmentation model is effectively estimating the fragmentation effect only around the part of the gradient it happens to sample. If the fragmentation effect varies with habitat amount, as all frameworks except the HAH predict, that local estimate depends on where on the gradient the study's landscapes are located. Studies positioned at different points on the gradient will therefore measure different tangent slopes and report different fragmentation effects, even when the underlying surface is the same.

When the fragmentation effect varies with habitat amount ($\beta3 \neq 0$), the additive estimate of β1 approximates the average of $\beta1 + \beta3 \cdot A$ over the sample's habitat-amount distribution. This quantity is a property of where the study's landscapes sit on the gradient, not a property of the ecology alone. A study sampling primarily low-cover landscapes will estimate a more negative average β1 than an otherwise identical study sampling high-cover landscapes, even if the underlying surface is identical. Neither result is incorrect given its model. Both are uninformative about the shape of the surface.

This sampling artefact generates a testable prediction. If $\beta3 \neq 0$, studies sampling different regions of the gradient should produce directionally different average fragmentation effects even when the underlying ecology is identical. The interaction model tests that prediction rather than assuming it. The debate's pattern of unresolved directional disagreement is consistent with this prediction. Better controls and larger samples within the additive framework have not resolved it. Recent theory and reanalysis show explicitly that the parameter space of fragmentation is constrained by total habitat amount, such that empirical scenarios often span only restricted portions of the attainable joint space and therefore recover local slopes rather than the full response surface (Zhang et al. 2024).

The additive specification was not adopted carelessly. It arose from legitimate concerns about collinearity. Fragmentation metrics and habitat amount covary in most landscape datasets, and separating their effects has been a central methodological preoccupation of the field

(McGarigal and Cushman 2002; Ewers and Didham 2006; Smith et al. 2009). That preoccupation, however, addresses a different question from the one theory poses. Disentangling the independent effects of F and A is a problem of partial regression. Testing whether the effect of F changes as A changes is a problem of moderation. The field solved for the former and left the latter unaddressed. Adding an interaction between two correlated predictors amplifies collinearity and reduces precision, and for single-study datasets with limited cover range the interaction model may be underpowered (Dormann et al. 2013). These were real constraints that justified the additive default in an earlier data environment, and additive attempts to separate the effects of correlated predictors can also assign shared variance differently across parameterizations, causing their apparent importance to reverse (Koper et al. 2007; Smith et al. 2009).

That practical difficulty does not justify omitting the interaction test across all studies. The appropriate response is to accumulate multi-landscape datasets with adequate gradient coverage and to interpret null interaction results in relation to power and gradient breadth rather than as evidence of invariance by default. The historical rationale for the additive default does not extend to the current data environment, where multi-taxa, multi-landscape compilations spanning the full cover gradient are available (e.g., Zhang et al. 2024).

4. Translating mechanism into testable predictions

The four frameworks generate distinguishable predictions for the sign and shape of β3, the beta-diversity structure of communities, and the gradient position at which effects are expressed and recoverable. Together, these three signals provide the strongest basis for mechanism identification. β3 sign alone does not distinguish threshold from C–C, both of which predict $\beta3 > 0$, but it does distinguish extinction filtering, which uniquely predicts $\beta3 < 0$. Mechanism identification requires evaluating all three signals together. Even then, the results are diagnostic rather than determinative. Cases should remain ambiguous when multiple mechanisms generate overlapping signatures or when gradient coverage is limited. These limiting-case signatures are summarized in Table 4.1 and visualized as marginal-effect geometries in Fig. 4.2.

4.1 β3 sign and gradient shape as primary discriminators

The sign of β3 is the primary empirical clue, but it is not sufficient on its own (Table 4.1). The HAH uniquely predicts $\beta 3 = 0$ (together with $\beta 1 = 0$ at mean cover; Section 2.1). Extinction filtering uniquely predicts $\beta 3 < 0$, and the fragmentation effect therefore intensifies as cover increases, reaching its maximum negative value where sensitive species remain. Threshold and C–C both predict $\beta 3 > 0$ and are distinguished from each other not by β3 sign but by the sign of β1 at mean A, the beta-diversity component signatures, and the gradient position of recoverable effects. The HAH is compatible with an extinction threshold generated by habitat loss, but not with a fragmentation threshold generated by configuration, because HAH requires the fragmentation effect to remain near zero across the habitat-amount gradient.

Where the sample mean sits within the coexistence window, β1 is positive or near zero, the only case among the four frameworks in which fragmentation is ecologically beneficial or neutral at mean cover. Because β1 represents the fragmentation effect at the sample mean of A rather than at a fixed ecological reference point, its interpretation depends on where that mean lies relative to the coexistence window. Cross-study comparisons therefore require explicit reporting of gradient position or re-centering to comparable ecological reference points. β1 is a property of the sample as much as of the system.

This creates a specific inferential collapse under the dominant additive specification. Where a C–C system is sampled with its mean cover near the coexistence crossover, the fitted additive coefficient $\hat{\beta}_1$ approaches zero, which is the same signature HAH produces by prediction. The two frameworks therefore become statistically indistinguishable under the additive model alone, even though they imply different ecological responses away from the sample mean. Null results for β1 at mean cover therefore do not discriminate HAH from C–C without further evidence, including explicit estimation of β3, beta-diversity decomposition, or evaluation of the fragmentation effect away from the sample mean.

Threshold and extinction filtering both yield negative β1 at mean cover but differ in β3 sign. Threshold predicts $\beta 3 > 0$ (effects intensify as cover declines, concentrated below the connectivity threshold), whereas extinction filtering predicts $\beta 3 < 0$ (effects intensify as cover increases, because sensitive species are still present in high-cover landscapes and vulnerable to isolation). Among the four limiting cases considered here, extinction filtering is the only

framework predicting that fragmentation effects grow stronger as habitat becomes more abundant. Neither signal is accessible under the additive model.

At the species level, the extinction-filtering pattern ($\beta3 < 0$) has been documented in dynamic occupancy models of a threatened forest songbird (Hart and Bayne 2025). The C–C tradeoff pattern ($\beta3 > 0$ with sign reversal above approximately 20% habitat cover) has been documented in habitat selection of a farmland passerine across ten spatial scales (Bosco et al. 2021). Multi-taxon implementations recover several Table 4.1 signatures jointly. A Swedish forest analysis fitted 163 forest-associated plant taxa using an aggregation index. Ninety-five taxa showed no significant configuration effect or interaction, 33 showed gradient-invariant negative fragmentation effects, and 4 showed gradient-invariant positive effects. Among 31 taxa with a significant forest area × configuration interaction, 19 recovered $\beta3 > 0$ consistent with the threshold signature, 8 recovered a $\beta3 > 0$ pattern consistent with C–C sign reversal, and 4 showed a positive high-cover effect only (Lehtilä et al. 2020). Species richness recovered the threshold pattern, while extinction filtering was not over-represented. Forest amount fitted jointly with two configuration metrics across 70 eastern Ontario landscapes yielded metric-specific surfaces. Patch-number fragmentation showed little interaction with forest amount, whereas wooded-corridor connectivity showed a sign-reversing interaction, increasing biodiversity at low forest amount and decreasing it above approximately 30–40% (Daly et al. 2025). $\beta3$ therefore expresses different conditional surfaces along both the taxon axis within a community and the metric axis within a pooled response.

The empirical literature on the cover × fragmentation interaction is therefore not convergent on a single result. Studies recovering negative-and-fading $\beta3$ (Hart and Bayne 2025; Betts et al. 2017; Hua et al. 2024), positive-and-sign-reversing $\beta3$ (Bosco et al. 2021; Zhang et al. 2024), threshold-consistent $\beta3$ (Smith et al. 2011), and joint-signature heterogeneity within single datasets (Lehtilä et al. 2020; Daly et al. 2025) document conditional response surfaces. Studies reporting null $\beta3$ (Trzcinski et al. 1999; Watling et al. 2020) document either HAH-consistent invariance, recoverability-window failure under the response–scale–metric–taxon combination tested, or partial-identification artefacts (Sections 3.2 and 3.3). The empirical task the framework specifies is the joint distribution of $\beta3$ sign, gradient shape, $\beta1$ at mean cover, gradient coverage, beta-diversity decomposition, and assemblage filtering history, reported

across response variables and configuration axes within each system and aggregated across systems.

A parallel heterogeneity–diversity literature makes a related point from the opposite direction. Increasing heterogeneity can promote richness through niche diversification, but subdivision can reduce richness when effective area per condition becomes limiting (Allouche et al. 2012; Ben-Hur and Kadmon 2020; Laanisto et al. 2013). The geometric convergence is adjacent rather than mechanistic. The same conditional logic appears in both traditions, with subdivision effects depending on where the system sits along an amount and persistence axis.

Beyond the sign of β3, the shape of the conditional surface provides additional discrimination. Threshold theory predicts a near-flat surface at high cover that bends sharply negative below the threshold, a nonlinear, step-structured change that the linear interaction term approximates poorly and that generalized additive or hierarchical threshold models are better positioned to detect. The C–C framework predicts a more gradual sign reversal across the gradient. Extinction filtering predicts a surface that declines monotonically but with a slope that weakens as cover decreases. For threshold detection specifically, flexible response-surface models, including generalized additive or hierarchical threshold models, are better suited than a single linear interaction term (Pedersen et al. 2019).

4.2 Beta-diversity as a secondary discriminator

Beta diversity, like richness, is an emergent property of assembled multi-species communities, but unlike richness it decomposes into turnover and nestedness components, which are available as a secondary discriminator wherever community-level data are present (Baselga 2010; Legendre 2014; Rybicki et al. 2020; Betts et al. 2019). Under the C–C tradeoff, landscape-scale diversity benefits arise from compositional differentiation among patches, expressed as turnover. Large patches support competitive dominants, small patches support strong colonizers, and the landscape mosaic maintains both. Turnover should therefore be elevated in fragmented landscapes at intermediate-to-high cover, where the coexistence window is open, and should decline at low cover as competitive dominants are globally lost. Nestedness, reflecting species loss rather than replacement, should increase at low cover as extinction progresses.

Under extinction filtering, turnover among patches should be weak across all cover levels, because the community is dominated by matrix-tolerant generalists whose occurrence

patterns reflect dispersal capacity rather than competitive hierarchy (Betts et al. 2019). Nestedness should increase at mid-to-high cover, where filtering is underway but not complete. The sign of β3 for the turnover component provides corroborative contrast between C–C (positive β3 for turnover) and extinction filtering (near-zero or negative β3 for turnover), although the strength of this contrast depends on gradient coverage, sampling completeness, richness gradients, and the decomposition used. These components are not independent of richness gradients, so beta-diversity signatures should be interpreted alongside richness, sampling completeness, and gradient coverage (Baselga 2010; Legendre 2014; Anderson et al. 2011). These signatures should be reported alongside the corresponding richness response, because a turnover increase from coexistence, a nestedness increase from filtering, and a dissimilarity increase driven by local richness loss are not equivalent conservation signals even when each increases some measure of beta diversity (Anderson et al. 2011; Socolar et al. 2016).

For extinction filtering specifically, Baselga's nestedness-resultant component operationalizes the mechanism's prediction more precisely than richness-difference formulations (Baselga 2010; Legendre 2014). Nestedness-resultant quantifies the degree to which compositional differences arise from one community being a proper subset of another, whereas richness-difference formulations detect any count asymmetry between communities without requiring that subset structure.

4.3 The recoverability window

Conditional fragmentation effects are empirically recoverable only where the focal response remains biologically sensitive to configuration and the sampled landscapes provide enough configuration contrast to estimate that sensitivity. These conditions are response-specific and may shift with taxon, spatial scale, matrix permeability, fragmentation metric, and landscape-generation process. Continental analyses confirm that maximally fragmented configurations are structurally achievable only in low-cover landscapes, because high-cover landscapes retain most habitat in a single dominant patch regardless of sampling design (Rueda et al. 2013). For threshold dynamics, the configuration-sensitive window W_R spans the gradient between ET_R and FT_R. Above FT_R configuration effects are weak because habitat remains functionally connected. Below ET_R sensitive taxa may no longer persist. Within W_R, β3 is most likely to be recoverable, but configuration variability peaks at intermediate habitat amounts and

collapses at very low cover (Andrén 1994; Hanski 2015). The C–C sign reversal requires observations spanning both high and low cover to detect the full slope transition (Rybicki et al. 2020; Zhang et al. 2024). Extinction filtering effects are detectable at mid-to-high cover, before filtering is complete, and collapse at low cover where the signal species have been lost (Betts et al. 2019).

Scale dependency of the recoverability window itself is a source of null results. The interaction is recoverable only where F and A are measured at the scale at which the taxon responds to configuration; 44% of species in a global review had their scale of effect at the extreme of the range evaluated, indicating the true scale was not measured (Jackson and Fahrig 2015). The interaction has been detected at territory and home-range scales but becomes non-significant at broader landscape scales in both boreal songbirds and farmland passerines (Hart and Bayne 2025; Bosco et al. 2021). It can also reverse sign across scales, with negative effects at landscape scale and positive effects at local scale in a marbled murrelet occupancy analysis (Valente et al. 2023b). Matrix quality also modulates the recoverability window. Low-quality matrices intensify habitat-loss filtering, while higher-quality matrices relax filtering through mass effects and resource supplementation (de Souza Leite et al. 2022).

For movement-mediated mechanisms specifically, the recoverability window may peak at intermediate rather than lowest habitat amounts, where patches exceed local viability thresholds, but isolation still limits colonization. At very low cover, local habitat sufficiency can become the binding constraint, and the marginal effect of connectivity can collapse regardless of isolation (Villard and Metzger 2014).

A study restricted to a narrow gradient range may produce a null or weakly significant $\beta 3$ not because the interaction is absent, but because the sample does not span the gradient region where it operates. Characterizing the cover distribution of the sample is therefore necessary context for interpreting any null interaction. This includes the range, mean, and representation across the low-cover gradient where threshold and filtering mechanisms operate. A non-significant $\beta 3$ from a study spanning the full gradient, combined with $\beta 1 \approx 0$ at mean cover, narrow confidence intervals, and adequate power, constitutes genuine support for the HAH. A non-significant $\beta 3$ from a study restricted to high-cover landscapes is uninformative about threshold predictions and cannot be presented as resolving the debate. Cover-restricted analyses

can still test local mechanism-specific predictions, but they cannot estimate the full conditional response surface or resolve the debate-level question (Villard and Metzger 2014; Betts et al. 2019). Apparent null fragmentation coefficients can arise from configuration invariance under the HAH, reachability saturation above FT_R, response collapse below ET_R, or cancellation, insufficient configuration contrast, or a flat local surface inside W_R.

Existing syntheses and experiments test different inferential targets. Species-density interaction tests aggregated across studies (Watling et al. 2020) estimate an average interaction coefficient at the alpha-diversity scale, independent of gradient-position effects and of beta-diversity decomposition (Section 4.2). Sign-counting reviews (Fahrig 2017) aggregate across response variables, mechanisms, and gradient positions without estimating any specific β3. Controlled fragmentation experiments with temporal resolution (Haddad et al. 2017a) estimate configuration effects on species–area relationships at specific amount-configuration contrasts and reveal the temporal dynamics of extinction filtering as it unfolds (Section 2.4), a signal that cross-sectional designs cannot directly observe. The full conditional response surface locates each of these inferential targets and clarifies what β3 captures. Their apparent disagreements reflect partial characterization by each design rather than irreconcilable findings about whether configuration matters.

5. Four informative outcomes

The four outcomes below define inferential conditions rather than analytical prescriptions. They identify the minimum evidential standard for contributing to the debate and the circumstances under which the predictions can be distinguished in data.

Outcome 1. β3 significantly ≠ 0 with adequate gradient coverage and power. The data support a conditional fragmentation effect, and the response surface is therefore not flat along the fragmentation axis. Evaluate direction and functional form against Table 4.1, using β1 at mean A and the beta-diversity signatures to identify which framework is most consistent with the data. Where $\beta3 > 0$ and β1 is negative, threshold dynamics are consistent. Extinction filtering is not supported as the dominant mechanism, although mixed filtering-plus-threshold systems remain possible and should be evaluated using beta-diversity signatures and assemblage history. Where $\beta3 < 0$ and β1 is negative, extinction filtering is consistent. Confirmation requires nestedness at mid-to-high cover and absence of turnover elevation. Where $\beta3 > 0$ and β1 is

positive or near zero, C–C tradeoff is consistent, subject to gradient position and turnover signatures. The result is diagnostic rather than causal proof. It identifies patterns consistent with mechanisms but does not assign exclusive causation where multiple processes generate overlapping signatures.

Outcome 2. β1 ≈ 0 at mean cover and β3 ≈ 0 with narrow confidence intervals, adequate gradient coverage, and sufficient power. This constitutes genuine empirical support for the HAH at the spatial scale and taxonomic group analyzed. Because the interaction framework nests all competing hypotheses within the same statistical model, evidentiary standards must be applied symmetrically. Strong support for any framework requires gradient coverage spanning the recoverability windows of all competing frameworks, sufficient statistical power to detect biologically plausible interaction effects, and consistency of the β3 signature across richness and beta-diversity components. For HAH specifically, this means β1 ≈ 0 at mean cover and β3 ≈ 0 at the richness level with narrow confidence intervals under these conditions. Where community composition data are available, the same absence of configuration-driven conditionality should also hold for turnover and nestedness components, distinguishing genuine ecological invariance from a failure to detect conditionality.

A boundary case distinct from Outcome 2 arises when β3 ≈ 0 with narrow confidence intervals but β1 is detectably non-zero. HAH is rejected because it requires both restrictions simultaneously. None of the three conditional frameworks applies, since threshold, C–C tradeoff, and extinction filtering all predict β3 ≠ 0. The result identifies configuration as consequential but gradient-invariant within the sampled range. It should be reported as such and not conflated with Outcome 2. The 33 plant taxa in Lehtilä et al. (2020) with significant negative fragmentation main effects and no significant interaction exemplify this outcome.

Outcome 3. β3 ≈ 0 with wide confidence intervals, restricted gradient coverage, or insufficient power. β3 carries information only when reported with uncertainty and interpreted in relation to gradient coverage and sensitivity to biologically plausible conditional effects. This outcome is inconclusive. The fragmentation effect may be genuinely invariant across the gradient. Opposing conditional effects may cancel at the sample mean, leaving β1 near zero even when β3 is large. The gradient may also be sampled too narrowly to detect an interaction operating elsewhere. These are fundamentally different ecological conclusions, and an additive

model cannot distinguish between them. The result is interpretable only in the context of gradient characterization and sensitivity assessment, and should not be presented as resolution of the debate. A fourth possibility arises when the sample mean lies near the zero-crossing of the conditional fragmentation effect, the C–C coexistence-to-isolation transition. In that case, β1 will be near zero even when β3 is large and ecologically meaningful. This is the sampling artefact of §3.3 at a specific gradient position. Beta-diversity decomposition provides the diagnostic, with turnover revealing the C–C signal even when richness cancels near the mean.

Outcome 4. β3 inconclusive with a significantly negative β1. A detectably negative mean fragmentation effect exists at the sample mean of habitat amount, but the dataset cannot resolve whether that effect intensifies, reverses, or fades along the cover gradient. Where design, scale, and predictor-geometry conditions are adequate for interpreting β1, this outcome rejects the HAH, which requires both $\beta1 \approx 0$ and $\beta3 \approx 0$. It should not be reported as inconclusive in the sense of uninformative. The interaction model has revealed that fragmentation is harmful on average under the realized gradient conditions, while leaving the shape of that effect unresolved. The appropriate conclusion is that conditionality remains unresolved rather than absent, and that the correct remedies are wider gradient coverage, larger sample size, or beta-diversity decomposition rather than abandonment of the conditional framework. Treating this outcome as equivalent to Outcome 3 collapses a distinction the framework is specifically designed to preserve. All four outcomes advance the debate. Only omission of the interaction test, or conflation of Outcome 3 with the β1-negative subcase, preserves stasis.

6. Implications for fragmentation research

Long-term fragmentation datasets contain the raw material for the first test. Re-analyses that estimate β3 with explicit characterization of gradient coverage would establish, without new data, how many published null results are uninformative rather than ecologically invariant. New designs should prioritize landscapes spanning the recoverability window each framework predicts (Section 4.3) rather than those optimized for predictor decorrelation alone. Decorrelating predictors across too narrow a cover range estimates β3 on a sample that cannot distinguish the framework-specific predictions. Community-level analyses should report β3 sign and gradient shape and, where community composition data allow, beta-diversity decomposition. These

quantities allow framework-specific signatures to be recovered jointly rather than conflated under an aggregate effect (Section 4.2).

These design and reporting reforms address a problem that has been institutionalized in empirical practice. Recent design-level critique reaches the same conclusion in the specific context of amount-controlled fragmentation inference, showing that habitat amount control is necessary, but not sufficient, unless predictor separability is demonstrated in the realized dataset (Martínez-Lanfranco 2026). Independent audits across more than two decades confirm that the proportion of fragmentation studies separately testing configuration and amount has remained low, with stricter separability criteria rarely met in earlier reviews and the weaker "control for habitat amount" criterion satisfied by approximately one in five studies in more recent ones (McGarigal and Cushman 2002; Hadley and Betts 2016; Riva et al. 2024). The reanalysis programme above is therefore a remediation of accumulated practice as much as a prospective agenda.

The reframing extends the existing call for theory investigation (Valente et al. 2023a) by identifying the conditional quantity along which the competing theoretical predictions are defined, treating it not as one source of variance among contextual moderators but as the organizing axis of the inferential target. Whether subdivision of remaining habitat is expected to harm, benefit, or leave unaffected a focal assemblage depends on where that landscape sits along the habitat-amount gradient and which response-surface geometry its community reflects. Minimum-patch-size thresholds and connectivity priorities for reserve design and restoration may therefore differ not only across taxa but across positions on the gradient. Additive models cannot recover this dependence, whereas the interaction framework makes it directly estimable (Riva and Fahrig 2022, 2023; Fahrig et al. 2021). To our knowledge, no published multi-taxon implementation has yet demonstrated the full discriminatory potential of this approach by jointly estimating β3 signatures, beta-diversity components, gradient-position diagnostics, and metric–mechanism contrasts across multiple taxa and landscapes. Such a test is the most direct route from conceptual reframing to empirical resolution.

7. Conclusions

The fragmentation debate is about the shape of the biodiversity response surface, not the sign of a single coefficient. Competing theories make predictions about how fragmentation

effects vary across the habitat-amount gradient, not only about whether one average configuration effect exists. Analyses that sample different parts of the gradient or constrain the fragmentation effect to be invariant across it need not converge on the same answer even when the underlying ecology converges.

The field has had the necessary analytical components for decades. The threshold logic was established early on (Andrén 1994). The interaction implication was stated directly soon after (Fahrig 2003). Configuration sensitivity was later placed along the habitat-amount gradient, and recent work has shown habitat-amount-constrained fragmentation domains and slope reversals (Villard and Metzger 2014; Zhang et al. 2024). The persistence of the debate reflects a translation failure. Theoretical predictions were repeatedly formalized, but statistical practice repeatedly omitted their operationalization (Haddad et al. 2017b; Valente et al. 2023a; Didham et al. 2012). In its most compressed form, the problem is that dominant empirical practice has estimated β1, the fragmentation effect at the sample mean of cover, where theory predicts conditionality. In the linear model, β3 summarizes that conditionality as the rate at which the fragmentation effect changes along the habitat gradient.

Null results require a higher evidential standard. A non-significant main effect of fragmentation in an additive model is not evidence that fragmentation has no effect on biodiversity. It is evidence that the average slope of the response surface at mean cover in that sample is not distinguishable from zero, a local and gradient-dependent quantity. Whether this reflects genuine ecological invariance (flat surface), a mispositioned sample on the gradient (non-flat surface sampled away from the active region), extinction filtering (signal species lost), or predictor geometry cannot be determined without estimating β3, characterizing the cover gradient, and assessing assemblage composition. To our knowledge, this is the first synthesis to organize the habitat amount hypothesis, threshold/scarcity dynamics, colonization–competition coexistence, and extinction filtering around a common conditional inferential target. In the linear interaction model, β_3 formalizes that target as the rate at which the fragmentation effect changes along the habitat-amount gradient. This synthesis derives the distinguishable gradient predictions each framework makes and specifies the statistical structure needed to distinguish them empirically. That connection between theoretical prediction and empirical practice is what has been missing from the debate.

References

Almeida-Gomes M, Prevedello JA, Crouzeilles R (2016) The use of native vegetation as a proxy for habitat may overestimate habitat availability in fragmented landscapes. Landscape Ecology 31:711–719

Allouche O, Kalyuzhny M, Moreno-Rueda G, Pizarro M, Kadmon R (2012) Area–heterogeneity tradeoff and the diversity of ecological communities. Proceedings of the National Academy of Sciences 109:17495–17500

Anderson MJ, Crist TO, Chase JM, Vellend M, Inouye BD, Freestone AL, Sanders NJ, Cornell HV, Comita LS, Davies KF, Harrison SP, Kraft NJB, Stegen JC, Swenson NG (2011) Navigating the multiple meanings of β diversity: a roadmap for the practicing ecologist. Ecology Letters 14:19–28

Andrén H (1994) Effects of habitat fragmentation on birds and mammals in landscapes with different proportions of suitable habitat: a review. Oikos 71:355–366

Balmford A (1996) Extinction filters and current resilience: the significance of past selection pressures for conservation biology. Trends in Ecology and Evolution 11:193–196

Baselga A (2010) Partitioning the turnover and nestedness components of beta diversity. Global Ecology and Biogeography 19:134–143

Ben-Hur E, Kadmon R (2020) An experimental test of the area–heterogeneity tradeoff. Proceedings of the National Academy of Sciences 117:4815–4822

Betts MG, Fahrig L, Hadley AS, et al. (2014) A species-centered approach for uncovering generalities in organism responses to habitat loss and fragmentation. Ecography 37:517–527

Betts MG, Wolf C, Pfeifer M, et al. (2019) Extinction filters mediate the global effects of habitat fragmentation on animals. Science 366:1236–1239

Betts MG, Wolf C, Ripple WJ, Phalan B, Millers KA, Duarte A, Butchart SHM, Levi T (2017) Global forest loss disproportionately erodes biodiversity in intact landscapes. Nature 547:441–444

Bosco L, Cushman SA, Wan HY, et al. (2021) Fragmentation effects on woodlark habitat selection depend on habitat amount and spatial scale. Animal Conservation 24:84–94

Daly L, Gabriel J, Hajdasz AC, et al. (2025) Connectivity of forest patches via wooded corridors increases biodiversity at low, but not high, forest amounts. Conservation Letters 18:e13154

De Camargo RX, Boucher-Lalonde V, Currie DJ (2018) At the landscape level, birds respond strongly to habitat amount but weakly to fragmentation. Diversity and Distributions 24:629–639

de Souza Leite M, Boesing AL, Metzger JP, Prado PI (2022) Matrix quality determines the strength of habitat loss filtering on bird communities at the landscape scale. Journal of Applied Ecology 59:2790–2802

Dennis M, Huck J, Holt C, et al. (2026) Beyond the patch: leveraging functional habitat delineation in fragmentation-biodiversity research. Landscape Ecology 41:37

Didham RK, Kapos V, Ewers RM (2012) Rethinking the conceptual foundations of habitat fragmentation research. Oikos 121:161–170

Dormann CF, Elith J, Bacher S, et al. (2013) Collinearity: a review of methods to deal with it and a simulation study evaluating their performance. Ecography 36:27–46

Duncan RP, Kefford BJ (2021) Interactions in statistical models: three things to know. Methods in Ecology and Evolution 12:2287–2297

Ewers RM, Didham RK (2006) Confounding factors in the detection of species responses to habitat fragmentation. Biological Reviews 81:117–142

Fahrig L (1998) When does fragmentation of breeding habitat affect population survival? Ecological Modelling 105:273–292

Fahrig L (1999) Forest loss and fragmentation: which has the greater effect on persistence of forest-dwelling animals? In: Rochelle JA, Lehmann LA, Wisniewski J (eds) Forest fragmentation: wildlife and management implications. Brill, Leiden, pp 87–95

Fahrig L (2002) Effect of habitat fragmentation on the extinction threshold: a synthesis. Ecological Applications 12:346–353

Fahrig L (2003) Effects of habitat fragmentation on biodiversity. Annual Review of Ecology, Evolution and Systematics 34:487–515

Fahrig L, Baudry J, Brotons L, Burel FG, Crist TO, Fuller RJ, Sirami C, Siriwardena GM, Martin JL (2011) Functional landscape heterogeneity and animal biodiversity in agricultural landscapes. Ecology Letters 14:101–112

Fahrig L (2013) Rethinking patch size and isolation effects: the habitat amount hypothesis. Journal of Biogeography 40:1649–1663

Fahrig L (2015) Just a hypothesis: a reply to Hanski. Journal of Biogeography 42:993–994

Fahrig L (2017) Ecological responses to habitat fragmentation per se. Annual Review of Ecology, Evolution and Systematics 48:1–23

Fahrig L, Arroyo-Rodríguez V, Bennett JR, et al. (2019) Is habitat fragmentation bad for biodiversity? Biological Conservation 230:179–186

Fahrig L (2021) What the habitat amount hypothesis does and does not predict: a reply to Saura. Journal of Biogeography 48:1530–1535

Fahrig L, Arroyo-Rodríguez V, Cazetta E, Ford A, Lancaster J, Ranius T (2021) Landscape connectivity. In: Perry G, Francis R, Millington J, Minor E (eds) The Routledge Handbook of Landscape Ecology. Taylor and Francis, pp 67–88.

Fahrig L, Galán-Acedo C, Edwards BPM, Habrich AK, Martin AE, Albert CH, Arroyo-Rodríguez V, Binley AD, Cole JR, Lockhart JK, Nielsen SE, Riva F, Schmoeller MPR, Tscharntke T (2026) Why controlling for habitat amount is critical for resolving the fragmentation debate. Conservation Biology e70245

Fletcher RJ Jr, Bullock JM, Graham LJ, Kortessis N, Watts K (2026) Resurrecting habitat fragmentation as a process over time. Trends in Ecology & Evolution [in press]. https://doi.org/10.1016/j.tree.2026.03.009

Fletcher RJ Jr, Burrell N, Reichert BE, Vasudev D (2016) Divergent perspectives on landscape connectivity reveal consistent effects from genes to communities. Current Landscape Ecology Reports 1:67–79

Fletcher RJ Jr, Didham RK, Banks-Leite C, et al. (2018) Is habitat fragmentation good for biodiversity? Biological Conservation 226:9–15

Fletcher RJ Jr, Betts MG, Damschen EI, et al. (2023) Addressing the problem of scale that emerges with habitat fragmentation. Global Ecology and Biogeography 32:828–841

Gardner RH, Milne BT, Turner MG, O'Neill RV (1987) Neutral models for the analysis of broad-scale landscape pattern. Landscape Ecology 1:19–28

Haddad NM, Gonzalez A, Brudvig LA, Burt MA, Levey DJ, Damschen EI (2017a) Experimental evidence does not support the Habitat Amount Hypothesis. Ecography 40:48–55

Haddad NM, Holt RD, Fletcher RJ Jr, Loreau M, Clobert J (2017b) Connecting models, data, and concepts to understand fragmentation's ecosystem-wide effects. Ecography 40:1–8

Gonçalves-Souza T, Chase JM, Haddad NM, et al. (2025) Species turnover does not rescue biodiversity in fragmented landscapes. Nature 640:702–706

Hadley AS, Betts MG (2016) Refocusing habitat fragmentation research using lessons from the last decade. Current Landscape Ecology Reports 1:55–66

Hua F, Wang W, Nakagawa S, Liu S, Miao X, Yu L, Du Z, Abrahamczyk S, Arias-Sosa LA, Buda K, Budka M (2024) Ecological filtering shapes the impacts of agricultural deforestation on biodiversity. Nature Ecology & Evolution 8:251–266

Halstead KE, Alexander JD, Hadley AS, et al. (2019) Using a species-centered approach to predict bird community responses to habitat fragmentation. Landscape Ecology 34:1919–1935

Hanski I (2015) Habitat fragmentation and species richness. Journal of Biogeography 42:989–993

Hart T, Bayne EM (2025) Habitat amount–fragmentation interactions drive Canada Warbler dynamics across spatial scales. Avian Conservation and Ecology 20(2):15

Henle K, Davies KF, Kleyer M, Margules C, Settele J (2004) Predictors of species sensitivity to fragmentation. Biodiversity and Conservation 13:207–251

Jackson HB, Fahrig L (2015) Are ecologists conducting research at the optimal scale? Global Ecology and Biogeography 24:52–63

Koper, N., Schmiegelow, F.K. and Merrill, E.H., 2007. Residuals cannot distinguish between ecological effects of habitat amount and fragmentation: implications for the debate. Landscape Ecology 22(6):811-820.

Kuussaari M, Bommarco R, Heikkinen RK, et al. (2009) Extinction debt: a challenge for biodiversity conservation. Trends in Ecology and Evolution 24:564–571

Laanisto L, Tamme R, Hiiesalu I, Szava-Kovats R, Gazol A, Pärtel M (2013) Microfragmentation concept explains non-positive environmental heterogeneity–diversity relationships. Oecologia 171:217–226

Legendre P (2014) Interpreting the replacement and richness difference components of beta diversity. Global Ecology and Biogeography 23:1324–1334

Lehtilä K, Vinter T, Dinnetz P (2020) Plant response to habitat amount and configuration in Swedish forests. Diversity and Distributions 26:329–339

Lundberg I, Johnson R, Stewart BM (2021) What is your estimand? Defining the target quantity connects statistical evidence to theory. American Sociological Review 86:532–565

Martínez-Lanfranco JA (2026) Habitat amount control is necessary but not sufficient to resolve the fragmentation debate. Conservation Biology. Accepted. EcoEvoRxiv. DOI 10.32942/X2KT05

McGarigal K, Cushman SA (2002) Comparative evaluation of experimental approaches to the study of habitat fragmentation effects. Ecological Applications 12:335–345

Miller-Rushing AJ, Primack RB, Devictor V, Corlett RT, Cumming GS, Loyola R, Maas B, Pejchar L (2019) How does habitat fragmentation affect biodiversity? A controversial question at the core of conservation biology. Biological Conservation 232:271–273

Neel MC, McGarigal K, Cushman SA (2004) Behavior of class-level landscape metrics across gradients of class aggregation and area. Landscape Ecology 19:435–455

Ovaskainen O, Soininen J (2011) Making more out of sparse data: hierarchical modeling of species communities. Ecology 92:289–295

Pedersen EJ, Miller DL, Simpson GL, Ross N (2019) Hierarchical generalized additive models in ecology: an introduction with *mgcv*. PeerJ 7:e6876

Popovic G, Mason TJ, Drobniak SM, et al. (2024) Four principles for improved statistical ecology. Methods in Ecology and Evolution 15:266–281

Püttker T, Crouzeilles R, Almeida-Gomes M, et al. (2020) Indirect effects of habitat loss via habitat fragmentation: a cross-taxa analysis of forest-dependent species. Biological Conservation 241:108368

Riva F, Fahrig L (2022) The disproportionately high value of small patches for biodiversity conservation. Conservation Letters 15:e12881

Riva F, Fahrig L (2023) Obstruction of biodiversity conservation by minimum patch size criteria. Conservation Biology 37:e14092

Riva, F., Koper, N., & Fahrig, L. (2024). Overcoming confusion and stigma in habitat fragmentation research. Biological Reviews 99:1411–1424.

Rueda M, Hawkins BA, Morales-Castilla I, Vidanes RM, Ferrero M, Rodríguez MÁ (2013) Does fragmentation increase extinction thresholds? A European-wide test with seven forest birds. Global Ecology and Biogeography 22:1282–1292

Ruffell J, Banks-Leite C, Didham RK (2016) Accounting for the causal basis of collinearity when measuring the effects of habitat loss versus habitat fragmentation. Oikos 125:117–125

Rybicki J, Abrego N, Ovaskainen O (2020) Habitat fragmentation and species diversity in competitive communities. Ecology Letters 23:506–517

Saura S (2021) The habitat amount hypothesis implies negative effects of habitat fragmentation on species richness. Journal of Biogeography 48:11–22

Spake R, Bowler DE, Callaghan CT, Blowes SA, Doncaster CP, Antao LH, Nakagawa S, McElreath R, Chase JM (2023) Understanding ‘it depends’ in ecology: a guide to

hypothesising, visualising and interpreting statistical interactions. Biological Reviews 98:983–1002

Smith AC, Koper N, Francis CM, Fahrig L (2009) Confronting collinearity: comparing methods for disentangling the effects of habitat loss and fragmentation. Landscape Ecology 24:1271–1285

Smith AC, Fahrig L, Francis CM (2011) Landscape size affects the relative importance of habitat amount, habitat fragmentation, and matrix quality on forest birds. Ecography 34:103–113

Socolar JB, Gilroy JJ, Kunin WE, Edwards DP (2016) How should beta-diversity inform biodiversity conservation? Trends in Ecology & Evolution 31:67–80

Tredennick AT, Hooker G, Ellner SP, Adler PB (2021) A practical guide to selecting models for exploration, inference, and prediction in ecology. Ecology 102:e03336

Trzcinski MK, Fahrig L, Merriam G (1999) Independent effects of forest cover and fragmentation on the distribution of forest breeding birds. Ecological Applications 9:586–593

Valente JJ, Gannon DG, Hightower J, et al. (2023a) Toward conciliation in the habitat fragmentation and biodiversity debate. Landscape Ecology 38:2717–2730

Valente JJ, Rivers JW, Yang Z, et al. (2023b) Fragmentation effects on an endangered species across a gradient from the interior to edge of its range. Conservation Biology 37:e14091

Villard M-A, Metzger JP (2014) Beyond the fragmentation debate: a conceptual model to predict when habitat configuration really matters. Journal of Applied Ecology 51:309–318

Watling JI, Arroyo-Rodríguez V, Pfeifer M, et al. (2020) Support for the habitat amount hypothesis from a global synthesis of species density studies. Ecology Letters 23:674–681

Weeks TL, Betts MG, Pfeifer M, et al. (2023) Climate-driven variation in dispersal ability predicts responses to forest fragmentation in birds. Nature Ecology & Evolution 7:1079–1091

With KA, Crist TO (1995) Critical thresholds in species' responses to landscape structure. Ecology 76:2446–2459

Zhang H, Chase JM, Liao J (2024) Habitat amount modulates biodiversity responses to fragmentation. Nature Ecology and Evolution 8:1437–1447

Table 4.1. Limiting-case predictions of four fragmentation frameworks for the conditional fragmentation effect along the habitat-amount gradient. β1 is the fragmentation effect at the sample mean of habitat amount. Its interpretation depends on where that mean lies relative to the response-specific threshold, coexistence crossover, or filtering window. β3 is the fragmentation × habitat amount interaction and describes how the fragmentation effect changes as habitat amount increases. A positive β3 means the effect becomes less negative, or reverses to positive, as cover increases. A negative β3 means the fragmentation effect becomes increasingly negative as cover increases. Recoverability windows indicate where each mechanism is both biologically expressed and empirically recoverable. ET_R, FT_R, and W_R denote the response-specific extinction threshold, fragmentation threshold, and configuration-sensitive window, respectively. In systems where multiple mechanisms co-occur, the observed β3 sign depends on which gradient region dominates sample variance, and ambiguity in mechanism identification should be reported. C–C denotes colonization–competition coexistence. F is coded so that higher values indicate greater fragmentation per se. Analysts using aggregation-based metrics should invert them before fitting or read predicted β signs and curve directions as reversed. Beta-diversity predictions apply to community-level implementations. Population-level applications should interpret the corresponding mechanisms through occupancy, abundance, or demographic responses.

Framework	**β1 (F at mean A)**	**β3 (F × A interaction)**	**Recoverability window**	**Beta-diversity prediction**
Habitat amount hypothesis (HAH)	= 0 no effect at mean cover	= 0 invariant across gradient	Full sampled gradient. Support requires β1 ≈ 0 and β3 ≈ 0 with adequate power, gradient coverage, and configuration contrast across the relevant response-specific range	No configuration-driven turnover or nestedness shift once habitat amount is controlled.

Threshold / scarcity	< 0 negative effect at mean cover	> 0 effect weakens as cover increases	Between ET_R and FT_R. Often low-to-intermediate cover, but position depends on landscape-generation process, matrix permeability, species traits, spatial scale, and response variable	Nestedness may increase within the lower portion of WR as persistence becomes limited. Turnover is not expected to be elevated at high cover.
Metacommunity C–C tradeoff sign reversal	> 0 or ≈0 neutral to positive at mean cover	> 0 effect becomes more positive as cover increases, producing sign reversal	Requires sampling across the response-specific coexistence crossover within W_R. The positive high-cover limb occurs inside the configuration-sensitive domain, not above FT_R	Turnover elevated at intermediate-to-high cover when the coexistence window is open. C–C is identified jointly by $\beta1 \approx 0$ and elevated turnover at high cover
Extinction filtering	< 0 negative effect at mean cover	< 0 effect intensifies as cover increases, reaching its most negative at high cover	Mid-to-high cover or any region where sensitive taxa remain. The signal fades after species-specific ET_R crossings accumulate and filtering is complete	Nestedness increases at mid-to-high cover, while turnover remains weak or absent. $\beta3 < 0$ is the unique signature distinguishing extinction filtering from threshold and C–C, both of which predict $\beta3 > 0$

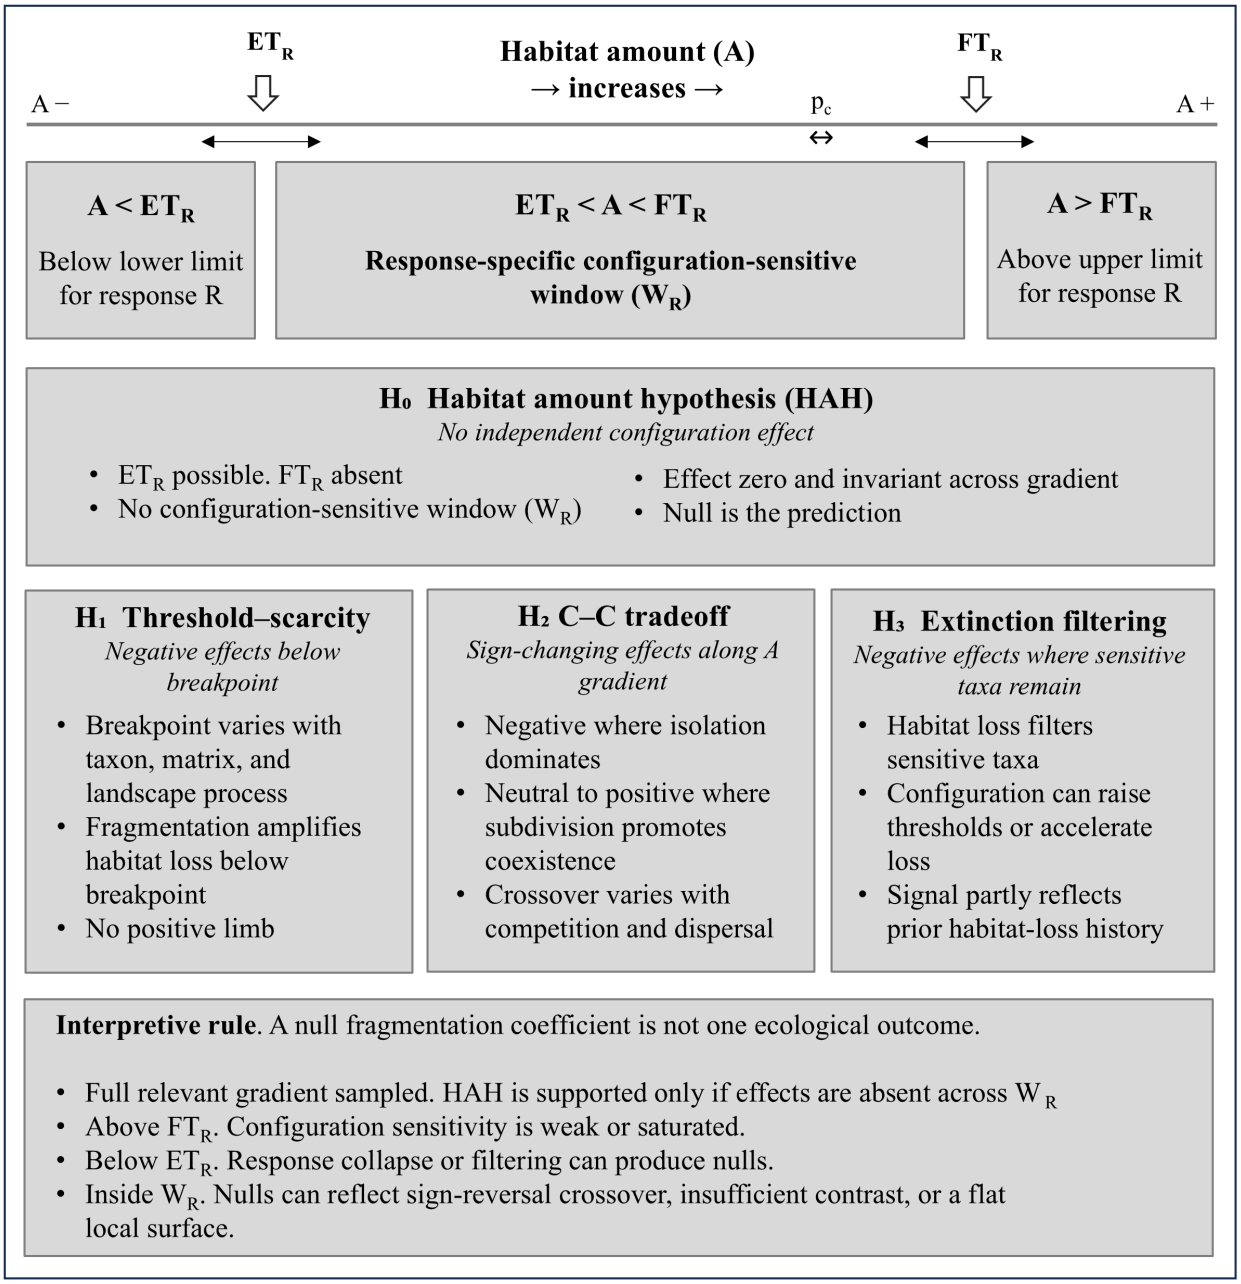


Figure 4.1. Conceptual response-surface framework for four competing fragmentation theories. The top panel defines regions along the habitat-amount axis. ETR is the lower persistence boundary, FTR is the upper configuration-saturation boundary, and WR is the response-specific configuration-sensitive window between them. pc marks a structural percolation threshold that can help determine FTR but is not equivalent to it. H_0 denotes the habitat amount hypothesis, H_1 threshold/scarcity, H_2 colonization–competition coexistence, and H_3 extinction filtering. H_0 predicts no configuration-sensitive window because configuration effects are invariant once

habitat amount is controlled. H_1 predicts negative effects in the lower portion of WR, strongest where the response persists. H_2 predicts sign-changing effects within WR. H_3 predicts negative effects strongest where sensitive taxa remain, fading as assemblage composition simplifies. Boundaries are schematic and vary with response variable, taxon, spatial scale, matrix, configuration metric, and landscape-generation process. The bottom panel shows when null additive fragmentation coefficients can reflect gradient position rather than ecological invariance. The statistical translation of these canonical cases appears in Fig. 4.2.

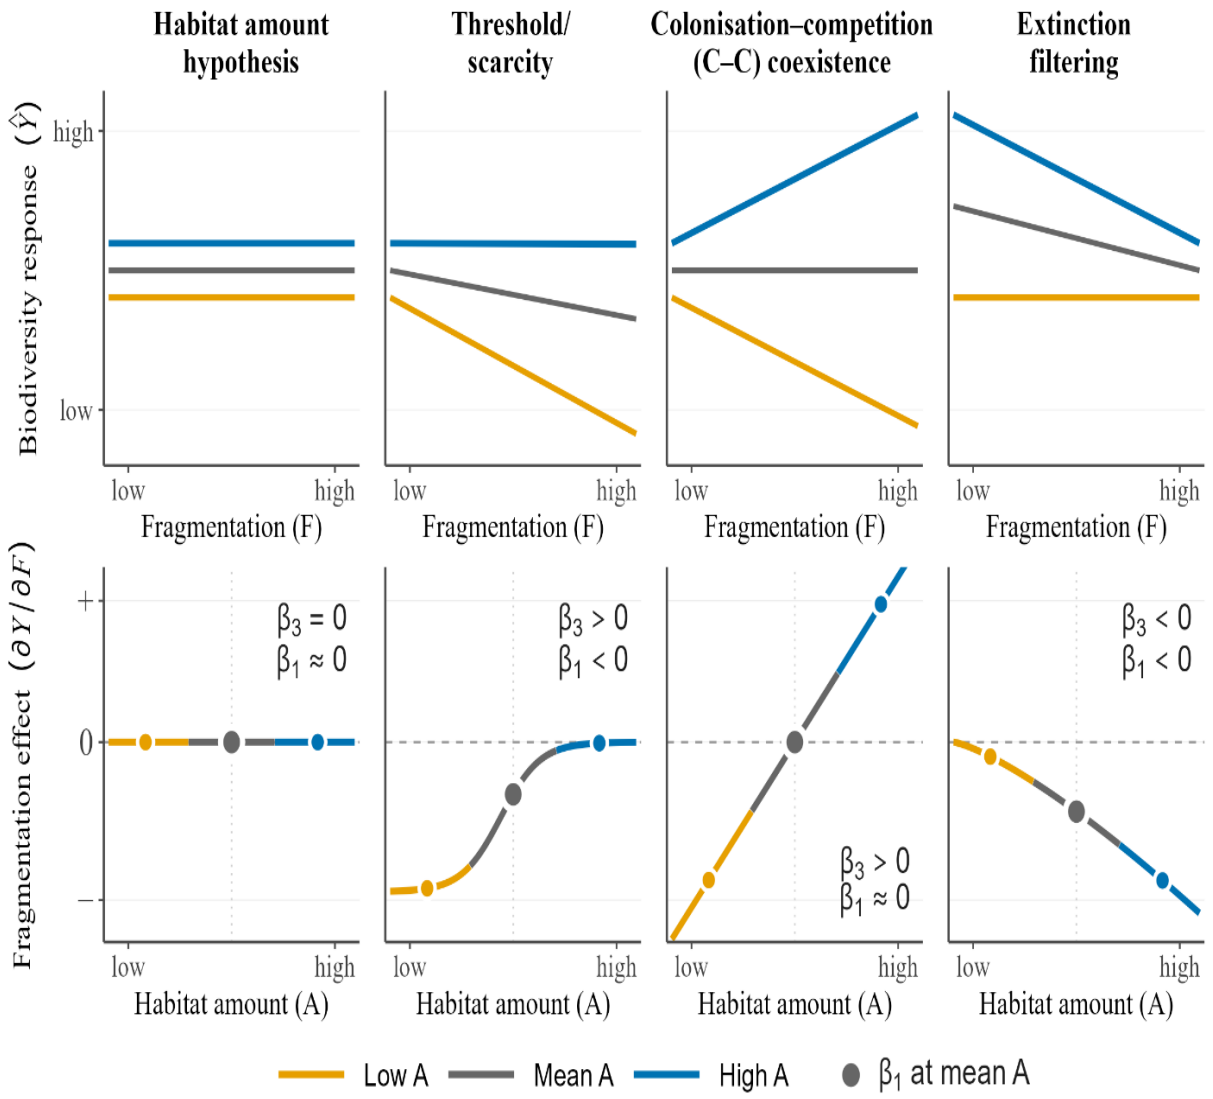


Figure 4.2. Statistical translation of the canonical alternatives in Fig. 4.1. The top row shows schematic response-scale biodiversity predictions ($\hat{Y}$) against fragmentation at low, mean, and high habitat amount. The bottom row shows the corresponding marginal fragmentation-effect functions, $\partial Y/\partial F = \beta_1 + \beta_3 A$, on the model scale. Open circles mark β_1, the fragmentation effect at mean habitat amount and the quantity estimated by additive models. Curve direction represents β_3. HAH predicts a near-zero effect across the gradient. Threshold/scarcity predicts negative effects that weaken as cover increases. C–C tradeoff predicts a sign reversal. β_1 can approach zero when mean A lies near the coexistence crossover, making C–C and the HAH indistinguishable in additive models alone unless β_3 is estimated and shown to be positive. Extinction filtering predicts stronger negative effects at higher habitat amounts, where sensitive

taxa remain, and weaker effects at lower amounts as those taxa are progressively lost. Curves are schematic limiting cases, not fitted estimates. In non-identity-link models, the top row uses back-transformed predictions, whereas the bottom row remains on the link scale where $\beta_1 + \beta_3 A$ is defined. F is coded so that higher values indicate greater fragmentation per se. Analysts using aggregation-based metrics should invert before fitting or read predicted β signs and curve directions as reversed.

CHAPTER 5

Where forest remains, configuration matters most: amount-conditioned associations in Atlantic Forest vertebrates

Introduction

Land-use intensification reshapes landscapes through habitat loss, habitat fragmentation, and habitat degradation, with consequences for biodiversity that propagate across spatial scales (Tscharntke et al. 2012; Newbold et al. 2015; Fischer and Lindenmayer 2007). The negative effects of habitat loss on population persistence and species richness are broadly accepted (Fletcher et al. 2018; Riva et al. 2024), but the ecological role of habitat fragmentation *per se* — the effects of spatial configuration once habitat amount is controlled (Fahrig 2003, 2013, 2017) — remains contested. Empirical studies have alternately reported negative, neutral, and positive effects on species richness after amount control (Fletcher et al. 2018; Fahrig et al. 2019; Miller-Rushing et al. 2019), and the disagreement concerns not only magnitude but direction. Species richness sits at the centre of this controversy because it is the most widely used biodiversity response in fragmentation research, allows direct comparison across studies and systems, and is the response for which the habitat amount hypothesis makes its most explicit quantitative predictions (Fahrig 2013; Watling et al. 2020).

We argue that this apparent contradiction reflects, in part, a mismatch between what theory predicts and what analyses estimating only an average fragmentation effect can recover. Multiple theoretical traditions converge on the prediction that fragmentation effects on species richness are conditional on habitat amount rather than independent of it. Early syntheses of bird and mammal responses showed that configuration effects emerge mainly when habitat cover declines below approximately 10–30% (Andrén 1994). Spatially explicit metapopulation models formalize a related expectation, with spatial subdivision amplifying extinction rates beyond what habitat loss alone predicts when total habitat is limited (Rybicki and Hanski 2013; Hanski 2015). Competitive metacommunity models extend this further. At high cover, fragmentation can promote coexistence through heterogeneity-mediated relaxation of competitive dominance, while at low cover isolation-driven extinctions dominate (Rybicki et al. 2020; Zhang et al. 2024).

Extinction filtering predicts a complementary pattern from a different starting point. Long-term habitat loss progressively removes fragmentation-sensitive taxa from assemblages, so the strongest configuration effects remain detectable only where sensitive species still persist (Betts et al. 2019). These conditional alternatives disagree about the direction and shape of fragmentation effects, but they share the expectation that the fragmentation effect varies with habitat amount. Within this theoretical landscape, the habitat amount hypothesis (Fahrig 2013) corresponds to the limiting null case in which that conditional effect is absent — a strong and testable assumption (Saura 2021; Fahrig 2021), not a default. Empirical tests that estimate only an average fragmentation effect after amount control implicitly assume the invariance component of this null, that the fragmentation slope does not vary with habitat amount (Villard and Metzger 2014; Valente et al. 2023). If fragmentation effects do vary with habitat amount, averaging across that gradient can attenuate or obscure strong conditional effects, and studies sampling different portions of the gradient can recover different average estimates from similarly structured response surfaces. The forest cover × fragmentation interaction is the minimum statistical structure on the linear predictor required to evaluate whether that assumption holds.

Beyond this conceptual mismatch, three empirical contingencies further obscure conditional fragmentation effects in observational landscapes. The first concerns habitat delineation. Vegetation classes are routinely treated as habitat, but habitat is species-specific. What constitutes habitat depends on organismal requirements, while mapped land cover only approximates them (Valente et al. 2023; Almeida-Gomes et al. 2016). Strict canopy-cover thresholds emphasize structurally intact forest interior but exclude degraded vegetation that may still maintain functional connectivity, whereas inclusive thresholds inflate apparent habitat amount and confound configuration with habitat quality (Almeida-Gomes et al. 2016; Dennis and Huck 2025). The second concerns metric architecture. Widely used fragmentation metrics covary nonlinearly with habitat amount because total edge length, patch number, and aggregation are geometrically constrained by cover (Neel et al. 2004; Wang et al. 2014; Li and Wu 2004). Linear amount control can be insufficient when the amount–fragmentation relationship is unimodal, peaking at intermediate cover where habitat is subdivided but not yet reduced to remnants; under such conditions, fragmentation metrics share complex variance with habitat amount that additive covariates cannot remove, and residual confounding may persist (Neel et al. 2004; Wang et al. 2014; Martínez-Lanfranco 2026). The third concerns spatial scale, which

operates on two distinct planes. On the phenomenological plane, the predictor extent over which landscape structure is quantified relative to the grain at which richness is measured determines which ecological processes are captured. Effects detectable at broad extents, where matrix permeability and inter-patch connectivity dominate (Hatfield et al. 2020), may be absent at finer grains where within-patch microhabitat processes prevail (Miguet et al. 2016; Jackson and Fahrig 2015). A recent primary-data synthesis found that flux and patch density had scales of effect wider than forest cover while mean nearest-neighbour distance and cohesion did not, indicating that differences in characteristic extent between composition and configuration are themselves metric-dependent (Gengler et al. 2024). On the analytical plane, patch-scale and landscape-scale metrics share geometric structure rather than functioning as independent predictors, so aggregating responses to the landscape grain can attenuate or reverse fine-scale signals through the modifiable areal unit problem — explaining why patch-scale studies routinely detect negative fragmentation effects while landscape-scale designs often report null or positive ones (Fletcher et al. 2023; Frazier and Kedron 2017).

A further biological consideration shapes how fragmentation effects appear at the community level. Aggregate richness combines species with contrasting habitat affiliations and sensitivities to spatial subdivision. Species with large area requirements, limited dispersal, and narrow forest specialization are often more vulnerable to fragmentation, whereas matrix-tolerant species may respond weakly or in different directions (Ewers and Didham 2006; Henle et al. 2004; Vetter et al. 2011). We therefore retain restricted total richness as an aggregate response and partition it into three mutually exclusive guilds — forest specialists, forest facultatives, and non-forest species — so that the aggregate can be compared with its components. This decomposition distinguishes broadly weak assemblage responses from effects concentrated in a sensitive guild and diluted in total richness, and indicates whether a pattern aligns with forest dependency or extends across the wider fauna. These contrasts localize biological specificity but do not directly identify demographic mechanism.

We address these contingencies jointly using a multi-scale, cross-product validation framework designed to detect conditional configuration effects along the habitat-amount gradient. We quantified configuration across multiple spatial grains and predictor extents to evaluate scale dependence, compared two independent forest products under three canopy-cover thresholds to assess robustness to habitat delineation, contrasted normalized configuration

metrics against widely used enumeration-based metrics in both residualized and non-residualized formulations to evaluate whether conclusions were sensitive to nonlinear amount–configuration coupling, and partitioned assemblages by habitat specialization to compare restricted-total richness with three mutually exclusive guild-richness responses (specialist, facultative and non-forest).

The Atlantic Forest of South America provides the gradient these analyses require. Over 84% of its original cover has been lost (Ribeiro et al. 2009), and the remaining remnants span a broad gradient from near-zero to relatively high local forest cover, and from highly subdivided patch mosaics to consolidated blocks (Vancine et al. 2024). This dual gradient provides the predictor variation needed to estimate whether configuration associations change along the forest-amount gradient across a broad empirical range. Most of that loss predates the satellite record. Cover has fallen by a further 2.4 percentage points across the 34 years since 1986, and biome-wide loss slowed markedly after 2005 (Vancine et al. 2024), so contemporary diversity patterns are likely shaped by long-standing landscape structure rather than only by immediate post-deforestation transients — though extinction debts cannot be excluded.

Hypotheses and predictions

The four theoretical frameworks introduced above — the habitat amount hypothesis as null, scarcity and coexistence dynamics, and extinction filtering — differ primarily in the predicted direction in which the fragmentation effect changes along the forest-amount gradient. We therefore use the sign of the forest cover × fragmentation interaction as the primary diagnostic for evaluating which mechanism is most consistent with the data (Table 5.1; Figure 5.2). Throughout, β_{frag} denotes the fragmentation effect at mean standardized forest cover, and $\beta_{frag \times amount}$ denotes the change in that fragmentation effect along the forest-cover gradient.

A configuration-independent habitat-amount account predicts that both the main effect of fragmentation and the fragmentation × forest-amount interaction are indistinguishable from zero, with marginal fragmentation effects near zero across the entire gradient (Fahrig 2013; Saura 2021). This constitutes the baseline against which the conditional alternatives are evaluated.

Scarcity and coexistence mechanisms both predict a positive fragmentation × forest-amount interaction. Under scarcity dynamics, subdivision is detrimental at all cover levels but

most severe where habitat is limited, so fragmentation effects weaken progressively as forest amount increases (Andrén 1994; Fahrig 1997, 1998; Flather and Bevers 2002; Hanski 2015; Rybicki and Hanski 2013; Villard and Metzger 2014). Under coexistence dynamics, heterogeneity-mediated relaxation of competitive dominance at higher cover may further offset, neutralize, or reverse isolation costs, potentially producing positive marginal effects at the upper end of the gradient (Rybicki et al. 2020; Zhang et al. 2024). The two mechanisms are mechanistically distinct in whether heterogeneity benefits can offset isolation costs at high cover, but they generate statistically indistinguishable predictions of a positive interaction coefficient when no sign reversal is detected. Empirical separation requires evidence that marginal fragmentation effects cross zero at higher forest cover — an outcome unlikely to be testable in severely deforested systems where high-cover landscapes are rare.

Extinction filtering predicts the opposite interaction sign. Under this mechanism, long-term habitat loss progressively removes fragmentation-sensitive taxa from assemblages, so fragmentation-driven richness declines remain detectable mainly where sensitive species still persist — at intermediate-to-high forest cover (Balmford 1996; Betts et al. 2019). The fragmentation × forest-amount interaction is therefore expected to be negative, with marginal effects intensifying as forest amount increases and attenuating at low cover due to prior compositional pruning. Under this mechanism, the absence of negative effects at low cover would reflect ecological simplification rather than benign subdivision. The signal would fade because responsive taxa had been lost, not because configuration had become ecologically inconsequential.

Two further predictions follow from trait-based vulnerability. Species reliant on structurally complex interior microhabitats, with low matrix tolerance and limited dispersal capacity, are disproportionately sensitive to fragmentation (Henle et al. 2004; Vetter et al. 2011), and species–area relationships are steeper for habitat specialists than for whole assemblages (Matthews et al. 2014). Forest-dependent richness should therefore show stronger negative marginal effects than total richness when fragmentation costs are expressed, with the magnitude difference reflecting the disproportionate contribution of fragmentation-sensitive taxa to specialist assemblages.

Forest-specialist richness is the response most directly aligned with sensitivity associated with forest dependency. We therefore treat facultative richness as an intermediate-affiliation contrast and non-forest richness as a specificity diagnostic, without imposing a monotonic ordering among guilds. Stronger negative effects in specialists than in the other guilds would support concentration of fragmentation costs among forest-dependent taxa. Comparable responses in specialists and non-forest species would weaken that dependency-graded interpretation, although they would not by themselves demonstrate artefact or exclude extinction filtering.

Finally, forest-definition thresholds are expected to modify the detectability — but not the mechanistic direction — of fragmentation effects. Strict canopy thresholds reduce mapped forest amount and emphasize interior-forest structure, shifting landscapes toward lower positions along the habitat-amount axis and amplifying isolation-driven responses (Neel et al. 2004; Wang et al. 2014). Inclusive thresholds inflate mapped forest amount and incorporate degraded vegetation, increasing apparent connectivity and blending configuration with habitat-quality heterogeneity (Almeida-Gomes et al. 2016; Dennis and Huck 2025). Thresholds should therefore modify the clarity at which conditional effects are detectable but not reverse the interaction sign predicted by each mechanism.

Because these are richness-based, cross-sectional data, the signatures above identify the mechanism most consistent with the observed response structure rather than directly demonstrating the underlying demographic process. Together, these predictions distinguish the competing frameworks by interaction sign, marginal-effect geometry and guild contrast (Table 5.1; Figure 5.2), while forest definition, metric, scale, guild structure and the cross-sectional response constrain how those signatures can be interpreted (Table 5.S3.3).

Methods

Study region

The Atlantic Forest of South America is one of the most severely transformed biomes on Earth, retaining approximately 12–16% of its original cover (Ribeiro et al. 2009). The remaining forest spans a broad gradient in both habitat amount and spatial configuration, from severely deforested landscapes with small fragmented remnants embedded in agricultural matrices to

relatively intact blocks where local cover exceeds 50%, with cover broadly stable since 2005 (Vancine et al. 2024). The dual gradient in cover and configuration, combined with that stability, makes the Atlantic Forest particularly suitable for testing whether fragmentation effects on richness vary along the forest-amount gradient.

Species assemblage data and richness responses

We analysed 1,360 aggregated taxon-by-method modelling records at the finest response grain, derived from 2,653 curated, georeferenced, effort-reported vertebrate assemblages sampled across the Atlantic Forest between 1981 and 2018. Assemblage data were obtained from the ATLANTIC series of open-access data papers published in Ecology: mid-large mammals from Lima et al. (2017) and Souza et al. (2019), small mammals from Bovendorp et al. (2017) and Figueiredo et al. (2017), bats from Muylaert et al. (2017), passerines from Hasui et al. (2018), and anurans from Vancine et al. (2018). The working dataset comprised five taxonomic groups for which taxon-level forest-affiliation classifications could be assigned consistently: mid-large mammals, small mammals, bats, passerines, and anurans. Snakes were excluded because the source data provided richness totals without species-level composition required for habitat-affiliation stratification, and butterflies were excluded because no standardized habitat-affiliation scheme was available for this group at the time of analysis. The dataset construction is summarised here to make the chapter self-contained; the full account is given in the General Introduction and Appendix 0.S1.

We assign each retained taxon to one of three mutually exclusive forest-affiliation guilds using its habitat records in the IUCN Red List, coded under the IUCN Habitats Classification Scheme (IUCN 2012). Forest specialists have habitat codes restricted to forest environments, forest facultatives have both forest and non-forest habitat codes, and non-forest taxa have no forest habitat code. For each assemblage, facultative richness equals forest-associated richness minus specialist richness, and non-forest richness equals restricted total richness minus forest-associated richness. Forest-associated richness is computed only as the intermediate quantity required to derive facultative richness; it is not fitted as a response. The three guild components are exhaustive within the retained five-taxon universe, so restricted total richness equals the sum of specialist, facultative, and non-forest richness in every assemblage. We verified non-negative component values and exact partition identity before model fitting. We retain restricted total

richness to estimate the aggregate response and evaluate dilution, but do not fit forest-associated richness as an additional overlapping response; its directly modelled aggregate estimand is therefore unavailable rather than statistically redundant. The corresponding fitted responses are restricted total richness and the specialist, facultative, and non-forest guilds (total_restr, specialist, facultative, non_forest). The spatial distribution of assemblages is shown in Figure 5.1.

Spatial design: response grains and predictor extents

We aggregated assemblages by taxon and sampling method into hexagonal response units at three nested grains: 10 km, 25 km, and 40 km radius. These grains represent local, intermediate, and broad landscape integration domains across which fragmentation processes operate through different spatial dependencies (Miguet et al. 2016; Jackson and Fahrig 2015). For each response grain, richness was recalculated within each hexagon and sampling effort was represented as the number of distinct assemblages contributing to that cell. The final dataset included 1,360 10-km, 1,020 25-km, and 833 40-km hexagonal response units.

Predictor extents ranged from 10 to 100 km and were paired with response grains under the constraint that predictor extent is greater than or equal to response grain, yielding 18 valid grain-by-extent combinations (7, 6, and 5 valid extents at the 10-, 25-, and 40-km grains, respectively).

Environmental predictors were quantified within concentric buffers of 10, 25, 40, 55, 70, 85, and 100 km radius around the centroid of the assemblages aggregated into each record, constituting a systematic sweep across local-to-regional predictor extents. We used year-aggregated richness to retain adequate data density within response cells. We did not fit a year-retained comparison, so temporal-aggregation robustness is not claimed.

Forest products and thresholds

Forest cover was represented by three raster streams — Hansen Global Forest Change resampled to 500 m and to 1,000 m, and Tuanmu and Jetz at 1,000 m — each binarized at three canopy-cover thresholds (70, 80, and 90%), yielding nine forest-definition combinations. This design supports a product contrast (Hansen 1,000 m versus Tuanmu 1,000 m at matched

resolution), a resolution contrast (Hansen 500 m versus Hansen 1,000 m within Hansen), and a threshold contrast; the nine combinations are forest definitions, not nine products.

The three raster streams derive from two forest products — the Hansen et al. (2013) Global Forest Change (GFC) product (at 500 m and 1,000 m resampled resolutions) and the Tuanmu and Jetz (2014) global 1-km consensus land-cover product (forest classes combined) — each binarized at canopy-cover thresholds of 70%, 80%, and 90%. Strict thresholds (90%) emphasize structurally intact forest interior and exclude degraded vegetation; inclusive thresholds (70%) incorporate moderately degraded or recovering forest. The two products differ in temporal anchor, spatial resolution, and classification framework; the threshold sweep tests robustness to the operational definition of habitat. The temporal mismatch between biodiversity sampling and forest products is treated as a design constraint typical of large-scale landscape studies. We do not select a single "correct" product or threshold; instead, we treat the nine stream × threshold combinations as a robustness dimension, distinguishing product (Hansen versus Tuanmu) from resolution (Hansen 500 m versus 1,000 m).

All nine forest-definition combinations — three raster streams (Hansen 500 m, Hansen 1,000 m, Tuanmu 1,000 m), each at 70, 80 and 90 % canopy thresholds — were retained throughout the complete primary additive and interaction factorial and the targeted interaction-control analysis. The parametric bootstrap (Fieberg et al. 2020) covers forest-specialist richness only; no additive environmental-control comparison is reported.

Fragmentation metrics

Forest configuration was quantified using four class-level landscape metrics calculated from the binary forest maps: DCLUMP (the negated clumpiness index, indexing landscape-level dispersion), NLSI (normalized landscape shape index, indexing edge geometry relative to maximum aggregation), EDGED (edge density, total forest edge length per unit area), and PATCHD (patch density, number of forest patches per unit area). Metric definitions follow McGarigal et al. (2012), and all four were computed with the R package *landscapemetrics* (Hesselbarth et al. 2019). All four metrics were directionally standardized so that higher values indicate greater fragmentation; for DCLUMP this is achieved by negating the clumpiness index, which is defined on a symmetric −1 to +1 scale so that negation preserves its range while reversing its direction.

We retained four configuration metrics — DCLUMP, NLSI, EDGED, and PATCHD — after excluding DISSAG, which the stage-02 predictor diagnostics showed to be highly redundant with NLSI ($r = 0.92–0.995$ across scale combinations). The retained metrics capture complementary aspects of configuration (dispersion, shape, edge, and patch density) and differ in their expected coupling to forest amount; we do not treat them as cleanly orthogonal, and address residual amount-coupling through the residualized formulation and the collinearity diagnostics.

Environmental controls

The primary analysis estimated amount-conditioned fragmentation effects. Each model included sampling effort, forest cover, fragmentation, the selected hierarchical random-effects structure and an observation-level random effect, together with a comparison-specific spatial autocovariate. That autocovariate was derived from the full non-spatial model and shared by the matched interaction and additive members. No explicit environmental covariates entered the primary models. We therefore refer to these as the primary spatially adjusted models, and where the analysis registry and the tables in Appendix 5.S5 record them as control-free the two terms are equivalent; “raw” and “residualized” refer only to the fragmentation formulation and never to the degree of environmental or spatial adjustment. We evaluated environmental robustness on the interaction surfaces meeting the primary support criterion.

Each primary-supported interaction was refitted after adding two environmental controls. Actual evapotranspiration (AET) represents realized water–energy availability, and was derived at 1 km resolution as the product of annual mean potential evapotranspiration (Title and Bemmels 2018) and the Priestley–Taylor coefficient of the Global High-Resolution Soil-Water Balance dataset (Trabucco and Zomer 2019). Topographic ruggedness (TOPOD) represents terrain heterogeneity, and is the terrain ruggedness index derived at 1 km resolution from the Global Multi-resolution Terrain Elevation 2010 dataset (Amatulli et al. 2018). The refits retained the matched response data, fragmentation formulation, polynomial form and random-effects architecture.

This targeted analysis asks whether the interaction surfaces selected for ecological interpretation remain directionally stable after explicit adjustment for water–energy availability and terrain heterogeneity. The primary criterion remains the basis for inference. The controlled ladder comprised the joint AET + TOPOD specification, an AET-only specification and a

TOPOD-only specification, so that dependence on control composition could be assessed. Vegetation texture, represented by the coefficient of variation in the enhanced vegetation index (EVICV) computed from MODIS imagery at 1 km resolution for 2000–2005 (Tuanmu and Jetz 2015), entered only as a separately labelled secondary sensitivity for forest-specialist richness, where it was added to the joint AET + TOPOD specification. The complete eligibility, support, denominator and mechanism-classification rules are set out in Appendix 5.S3, and Appendix 5.S4 summarises the model-building and inference sequence schematically (Figure 5.S4.1).

Predictor geometry and residualized fragmentation metrics

Before fitting richness models, we characterized the geometric relationship between forest amount and each fragmentation metric across products, thresholds, response grains, and predictor extents. This step was necessary because fragmentation metrics can show nonlinear, often unimodal dependence on forest cover, peaking at intermediate cover where habitat is subdivided but not yet reduced to remnants. Under such conditions, linear amount control in additive models can be insufficient to remove all composition–configuration coupling, and residual confounding may persist or even reverse coefficient signs (Neel et al. 2004; Wang et al. 2014). We treated this geometry audit as a prerequisite for interpretation because standard linear diagnostics — Pearson correlations and variance inflation factors — may fail to capture nonlinear amount–configuration coupling that persists after the linear association is removed; empirical documentation of this problem in a multi-taxa fragmentation dataset is provided by Martínez-Lanfranco (2026).

For each forest product × threshold × response grain × predictor extent combination, we fitted a generalized additive model with the fragmentation metric as the response and forest cover as a thin-plate regression spline smooth (Wood 2017), implemented in the *mgcv* R package. Residuals from this GAM define the amount-adjusted (residualized) fragmentation metric: positive residuals identify landscapes more fragmented than expected for their forest amount; negative residuals identify landscapes more aggregated than expected. We denote the residualized metrics resDCLUMP, resNLSI, resEDGED, and resPATCHD.

Residualization removes the nonlinear dependence represented by the fitted amount–configuration model and defines an operational estimand (Lundberg et al. 2021), the component

of fragmentation not predicted by forest amount under that model, not a causal intervention on landscape configuration.

Residualizing fragmentation against forest amount leaves forest amount unchanged and replaces each fragmentation metric with its deviation from the fitted amount-specific expectation. The raw and residualized formulations therefore target non-symmetric estimands (Koper et al. 2007; Smith et al. 2009), and residuals of one predictor on another do not isolate its independent effect (Freckleton 2002). The raw formulation estimates the joint amount-conditioned fragmentation surface, whereas the residualized formulation asks whether a configuration association persists in fragmentation variation not predicted by forest amount under the fitted nonlinear relationship. It is an amount-prioritizing sensitivity, not a cleaner or unbiased estimate of fragmentation per se. Persistence of direction and support after residualization provides strong robustness evidence, whereas attenuation or loss of support can reflect the changed estimand, predictor variance and precision and cannot by itself establish ecological absence. Reciprocal residualization would instead give fragmentation priority in the predictor representation and answer a diagnostic variance-allocation question rather than the conservative amount-adjusted sensitivity intended here, so it was not applied. Standardized raw and residualized coefficient magnitudes were not interpreted as symmetric estimates or as direct measures of the fraction of an effect removed.

Richness models were fitted using both raw and residualized fragmentation metrics. The two formulations are treated as complementary rather than competing: the raw formulation estimates the average fragmentation effect after linear amount control; the residualized formulation estimates the average effect of more-than-expected fragmentation at a given forest amount after nonlinear amount control. Raw–residualized concordance is one matched robustness axis, not the primary support criterion; the primary interaction gate is model support relative to the matched additive model together with structure-coefficient concordance. A structure coefficient is the correlation between a predictor and the fitted linear predictor, and indexes a term's direct association with the modelled response independently of its estimated coefficient, which makes it diagnostic of suppression under collinearity (Ray-Mukherjee et al. 2014). We computed structure coefficients directly from each fitted model rather than through a variance-partitioning package *partR2* (Stoffel et al. 2021), and interpret them as a concordance diagnostic rather than as a decomposition of explained variance.

Model architecture

The analysis comprises three model families, fitted with Poisson generalized linear mixed models (GLMMs) with an observation-level random effect to account for overdispersion (Harrison 2014; O'Hara and Kotze 2010), implemented in glmmTMB (Brooks et al. 2017). All continuous predictors were standardized to zero mean and unit variance prior to model fitting (Schielzeth 2010), so that β_{frag} is interpreted as the fragmentation effect at mean standardized forest cover, matching the notation introduced in the Hypotheses.

The three model families test progressively more complex fragmentation effects. (i) Additive raw: richness as a function of forest amount and fragmentation (primary, spatially adjusted, no explicit environmental covariates), with random intercepts for TAXA and METHOD plus an observation-level random effect. (ii) Additive residualized: as (i) but with the residualized fragmentation metric. (iii) Interaction: adding the forest-amount-by-fragmentation product term. Environmental controls (AET, TOPOD; and EVICV in sensitivity) enter only in the separate control-comparison arms, not in the primary models. The same factorial is fitted separately to each of the four response estimands (restricted total, specialist, facultative, non-forest), which are parallel aggregate and guild-specific estimands, not independent replicates, since restricted total richness is the arithmetic sum of the three guilds. The model universe comprises 1,296 design cells per response stratum and 5,184 cells in total — nine forest definitions × four fragmentation metrics × eighteen grain × extent combinations × two formulations. Each interaction model carries the random-slope structure selected for its cell at the model-structure stage and the polynomial structure selected for its cell at the functional-form stage. That linkage was audited against the frozen results master on the complete design key — stratum, metric, formulation, forest definition and scale combination — rather than assumed: all 5,184 cells resolve uniquely to one selected random-slope structure and one selected polynomial structure, the polynomial structures agree in every cell, and the random-slope structures agree in 5,180. The four exceptions are cells in which the selected correlated-slope structure failed to converge and the fitter recorded the no-random-slope structure instead, which the stored record carries as the structure fitted rather than the structure selected; none introduces a structure outside the candidate set, and all four occur in the residualized formulation of the two enumeration metrics at broad predictor extents.

The maximal raw interaction model can therefore be represented as:

resp ~ EFFsb + FCOV + FRAG + FCOV:FRAG + space + RS_TAXA + (1 | METHOD) + (1 | obs)

where EFFsb is standardised site-based sampling effort, FCOV is forest cover, FRAG is the focal configuration metric, space is the residual spatial autocovariate, and RS_TAXA denotes the selected TAXA-level random-slope structure. The residualized formulation replaces FRAG with resFRAG throughout, and the additive families drop the FCOV:FRAG product term. RS_TAXA was selected from four candidates: (1 | TAXA), (1 + FCOV | TAXA), (1 + FRAG | TAXA), and correlated (1 + FCOV + FRAG | TAXA). Appendix 5.S4 gives the complete model templates and selection diagnostics.

The interaction family is specified in two formulations (raw and residualized), yielding four model specifications per metric × product × threshold × response grain × predictor extent combination. The two interaction formulations test the same theoretical structure under different amount-control architectures. Their matched concordance is reported as a robustness property; primary support is evaluated within each formulation from model support relative to the matched additive model together with structure-coefficient concordance.

Random effects and spatial autocovariate

AICc served two distinct roles in this analysis. For component selection — polynomial form and TAXA-level random-slope structure — candidates within $\Delta AICc \leq 2$ of the lowest-AICc model were treated as statistically indistinguishable and the simplest member of that band was retained. For hypothesis testing, by contrast, an interaction model had to improve on its matched additive comparator by more than two AICc units, and satisfy the product-term structure-coefficient concordance criterion, to earn primary support. Keeping these uses separate prevents component selection from being conflated with evidence for the fragmentation × amount interaction itself (Burnham et al. 2011).

The random-slope structure was selected among four candidates — no random slope, a forest-cover slope, a fragmentation slope, and correlated forest-cover and fragmentation slopes — with random intercepts for TAXA and METHOD accounting for baseline differences among animal groups and survey protocols, and the interaction-term random slope fixed out on

estimability grounds. The selected random-slope structures are summarized in Appendix 5.S4 (Table 5.S4.2); the selected structure and its AICc difference from the no-slope candidate were retained for every cell, while AICc values for the two remaining candidates are unavailable and are not reported.

Residual spatial dependence was addressed using a derived spatial autocovariate constructed from the residuals of the corresponding non-spatial model and entered as a fixed effect in the final model (Crase et al. 2012; Fletcher and Fortin 2018). The term is derived from residuals rather than from a coordinate smooth fitted alongside the predictors, because forest amount, fragmentation and the environmental covariates are themselves spatially structured. A jointly estimated smooth would compete with them for the same broad-scale variation, leaving the fragmentation coefficients conditional on a surface that may absorb part of what the analysis is meant to estimate.

The autocovariate adjusts for residual spatial pattern of unspecified origin; it is not interpreted as an ecological predictor, and it is not a substitute for named environmental covariates, since it cannot identify the source of that structure and offers no adjustment for omitted influences that are not spatially patterned. Separate spatial surfaces were constructed for the additive raw, additive residualized, interaction raw and interaction residualized model families, so that each model was adjusted using residual structure generated under its own focal predictor architecture rather than a surface inherited from another family. For the environmental-control sensitivity, the spatial autocovariate of the matched primary model was retained, so that the comparison isolates the addition of AET, TOPOD and, where applicable, EVICV; this differs from the family-specific recomputation used in the primary model sequence and is stated as such.

Marginal effects and inference

The marginal fragmentation effect, denoted $\partial \log \mu / \partial \text{frag}$, is the derivative of fitted log-expected richness with respect to fragmentation, evaluated across the observed forest-cover gradient; for interaction models its sign and magnitude vary along that gradient, and for additive models it is constant. The marginal effect was evaluated over the central observed cover domain of each design cell, defined as the 5th to 95th percentile of that cell's own forest-cover distribution, with no extrapolation beyond it, and summarized at the 5th, 25th, 50th, 75th and

95th percentiles. Because no cell selected a polynomial fragmentation term, the marginal effect is linear in forest cover throughout and no cell was excluded from the derivative. Intervals derive from the parametric bootstrap for forest-specialist richness under the control-free specification, computed as percentile intervals across 250 resamples of the marginal effect itself rather than of its component coefficients; the remaining strata carry point estimates without intervals and are reported as such. For every interaction cell we retain, as separate fields, coefficient direction, interval support where available, model support relative to the matched additive model, structure-coefficient concordance, and the resulting primary-support classification, each with its denominator.

We treat the factorial outputs as a dependent robustness space rather than as independent tests. For interaction models we retain, as separate fields, coefficient direction, interval support (95% CI excluding zero), model support relative to the matched additive model, structure-coefficient concordance, and the resulting primary-support classification (model support and structure-coefficient concordance jointly). Product, resolution, threshold, scale, formulation, and response-stratum comparisons use matched eligible cells and are reported as robustness summaries with their paired denominators, not as an alternative support criterion. The evaluation domain, interval procedure and marginal-effect grid are defined from the fitted models. Those formal reporting, support and denominator rules were consolidated after initial directional summaries had been viewed and before the results master was constructed and the manuscript drafted; they are fixed for the present reporting but are not preregistered or result-blind (Appendix 5.S3).

Cross-stratum comparison architecture. We compare response strata only within matched design cells sharing the same fragmentation metric and formulation, forest stream, canopy threshold, response grain, predictor extent, and model architecture. We first compare restricted total and specialist estimates to evaluate whether aggregation dilutes a response concentrated among forest-dependent species. Because specialist richness is a component of restricted total richness, aggregate-component comparisons are interpreted as linked decompositions of the same assemblages rather than as independent ecological replications. We then evaluate specialist-facultative, specialist-non-forest, and facultative-non-forest contrasts to assess biological specificity. Effect-size and marginal-effect contrasts with paired uncertainty provide the inferential comparisons; sign coherence, support-frequency differences, and concordance are

treated as descriptive summaries of the dependent robustness space. We impose no magnitude-ratio threshold and do not interpret a higher frequency of supported cells as evidence of a larger effect. Paired curve comparisons use the intersection of matched central domains.

Model diagnostics and sensitivity analyses

Sensitivity analyses comprise the nine forest definitions, the raw-versus-residualized comparison, the eighteen grain by extent combinations, a leave-one-control-out comparison among controlled specifications, and a parametric bootstrap. The bootstrap covers forest-specialist richness at three levels — control-free, with actual evapotranspiration added, and with topographic ruggedness added — covering 411, 410 and 410 cells respectively; it is an adequacy check on inference already obtained rather than the inferential basis, since model-based inference rests on the matched parsimony comparison together with structure-coefficient concordance across all eligible cells in all four strata. Bootstrap coverage does not extend to the remaining strata. The secondary analysis adding vegetation texture to the two environmental controls covers forest-specialist richness only. Model adequacy is assessed on three layers. All 5,184 designed interaction models converged, and their matched additive comparison converged in 5,180; the remaining four cells are retained for direction summaries but lack the paired model that support classification requires, and are reported as missing paired members rather than as failed fits.

Estimation stability was assessed for forest-specialist richness by parametric bootstrap at three control levels, where interval-based and likelihood-based support agreed in all 411 primary-supported cells, bootstrap coverage extended to all of them and the median relative root-mean-square error was 0.271, giving no evidence of instability under the resampling design. Simulation-based residual diagnostics — scaled quantile residuals with dispersion, uniformity, outlier and zero-inflation tests — were computed for forest-specialist richness from the model objects retained by stage 06, without refitting, across 144 design cells at the reference forest definition, and are reported in Appendix 5.S4.3. Dispersion and zero inflation are not detected in any audited cell and outlier excess is absent in 95.8 % of them; the Kolmogorov–Smirnov uniformity test rejects in every cell, but the departure is small — central-quartile occupancy of 45.6 % against 50 expected and tail occupancy of 12.3 % against 10 — and is consistent with mild tail-heaviness under unconditional simulation rather than diagnostic of a specific

misspecification. These diagnostics assess distributional adequacy and play no part in model selection.

Analyses were run in R (R Core Team 2024), with glmmTMB for the generalized linear mixed models (Brooks et al. 2017), mgcv for the amount–configuration smooths and the spatial autocovariate (Wood 2017), DHARMa for the simulation-based residual diagnostics (Hartig 2022), marginaleffects for the marginal-effect derivatives (Arel-Bundock et al. 2024), ape for the spatial tests (Paradis and Schliep 2019) and lmtest for the likelihood-ratio tests (Zeileis and Hothorn 2002).

*Abbreviations used in figures and tables***.** AET, actual evapotranspiration; CLUMPY, clumpiness index; DCLUMP, negated clumpiness index, indexing landscape-level dispersion; DISSAG, disaggregation index; EDGED, edge density, total forest edge length per unit area; EVICV, coefficient of variation of EVI; FCOV, forest cover; FRAG, focal configuration metric; NLSI, normalised landscape shape index; PATCHD, patch density, number of forest patches per unit area; resFRAG, amount-residualised fragmentation; SC, structure coefficient; TOPOD, topographic ruggedness. Full definitions and provenance are given in Appendix 5.S2.

Results

Sampled assemblages and landscape coverage

The analysis universe, documented in full in Appendix 0.S1, comprises 1,360 assemblage records across 1,323 sampling locations at the 10 km response grain, 1,020 records at 25 km and 833 at 40 km, aggregated from 2,653 curated assemblages sampled between 1981 and 2018 and drawn from five habitat-classifiable animal groups and fourteen taxon-by-method sampling groups (Appendix 0.S1, Table 5.S1.1): anurans (496, 360 and 292 records across the three grains; three methods), mid- and large-bodied mammals (269, 221, 184; three methods), passerines (276, 192, 154; three methods), small mammals (206, 162, 130; three methods) and bats (113, 85, 73; two methods). Aggregation to coarser grains reduces the sample by 25.0 % from the 10 to the 25 km grain and by a further 18.3 % to 40 km, retaining 61.2 % of the finest-grain records at the coarsest. Sampling protocols and effort units differed substantially among taxonomic groups, as expected from their contrasting survey methods (Appendix 0.S1), and the sampling structure is carried in the random-effect architecture rather than assumed away.

Richness distributions differ markedly among the four response strata, and those differences bear on the precision with which each can be estimated. At the 25 km grain, restricted total richness has a median of 17 and a mean of 23.1 (range 1 to 212), forest-facultative richness a median of 12 and mean of 15.9, forest-specialist richness a median of 2 and mean of 4.6, and non-forest richness a median of 1 and mean of 2.6 (range 0 to 33). Zero prevalence separates the strata more sharply than the means: no restricted-total observation is zero, against 1.4 to 2.6 % of facultative observations, 22.2 to 24.6 % of specialist observations and 30.3 to 35.7 % of non-forest observations across the three grains. The non-forest stratum therefore carries roughly a third zeros and a mean of fewer than three taxa, which is the empirical basis for treating it as a specificity diagnostic with its own precision properties rather than as a negative control, and for evaluating its support frequency against that baseline rather than against the other strata.

Forest-cover and fragmentation geometry

The four configuration metrics relate to forest cover nonlinearly rather than linearly: fitted smooths used a median 7.9 to 8.4 effective degrees of freedom against a basis dimension of ten, and at the metric-summary level the smooths explained approximately 8 to 32 additional percentage points of variation relative to the corresponding median linear r^2. Absolute linear correlation with forest cover ran 0.34 to 0.77, and variance inflation for the raw fragmentation term in the joint AET + TOPOD controlled specification 1.18 to 2.61. Pairwise correlations, curve shapes and interior maxima are reported in Appendix 5.S2.3 and S2.4.

Residualization validation

Forest amount constrained the four fragmentation metrics to sharply different degrees. The fitted amount model explained a median 19.2 % of the variation in DCLUMP, 37.7 % in NLSI, 57.6 % in PATCHD and 77.1 % in edge density, so edge density is the metric most tightly determined by how much forest remains and DCLUMP the least. The coupling was also substantially nonlinear. At the metric-summary level PATCHD showed the largest excess: median deviance explained by the smooth was 0.576 against a median linear r^2 of 0.252, so a linear amount term captures less than half of the fitted relationship. Residualization reduced the fragmentation term's variance inflation in the joint AET + TOPOD controlled specification to between 1.03 and 1.06 for every metric, and the forest-cover term to between 1.91 and 1.94. In

that specification, forest-cover variance inflation above five occurred in 53 raw edge-density cells and reached 10.8. In the EVICV-augmented specialist secondary specification it occurred in 101 of the 648 raw cells — 77 under edge density, 18 under NLSI and 6 under PATCHD — with a maximum of 11.9. No residualized cell exceeded five under either specification, the maxima being 2.1 and 3.4. The residualized predictors carried in the fitted models were verified against refitted smooths across all 648 design cells and reproduced throughout, with a minimum correlation of 0.9998 and no cell below 0.99. The stored residual columns were subsequently standardized to unit variance, with a median scaling factor of 1.371 overall and 2.09 for edge density. One standardized residual unit therefore spans a smaller interval of the original standardized metric than one raw standardized unit, especially for edge density; relative residual quantities are interpreted here and absolute residual variance is not. Validation detail is reported in Appendix 5.S2.4.

Additive effects: raw and residualized formulations

Forest-amount effects were larger than fragmentation effects in specialist, facultative, and restricted-total richness under both formulations. Across raw and residualized models, median forest-cover coefficients ranged from +0.322 to +0.350 for specialists, +0.153 to +0.156 for restricted total richness, and +0.122 to +0.127 for facultative richness, with intervals excluding zero in 86.7 to 100 % of cells. Fragmentation coefficients are not summarized across metrics here, because neither their magnitude nor their interval support is invariant to metric choice; the metric-resolved distributions are shown in Figure 5.3 and Appendix 5.S5.2. Forest-amount effects were much more consistently interval-supported than fragmentation effects in every stratum and formulation. Fragmentation coefficients were negative in 69 to 90 % of cells depending on stratum and formulation, so direction was consistent while per-cell precision was modest; support frequency is not read as effect magnitude.

Non-forest richness inverts both terms. Its median fragmentation coefficient is positive (+0.035 raw, +0.031 residualized) and its median amount coefficient negative (−0.071, −0.069), but the amount interval excludes zero in only 23.5 % of raw and 9.3 % of residualized cells, against 86.7 to 100 % in specialist, facultative, and restricted-total richness. The inversion is therefore directional rather than precisely estimated, and the specificity argument rests on the interaction contrast rather than on this reversal. Descriptively, the median magnitude ratio

differed sharply between strata and was not invariant to the fragmentation metric. Non-forest richness responded to configuration about as strongly as to amount, with a median amount-to-fragmentation magnitude ratio between 1.18 and 2.18 across metrics and formulations. In specialist, facultative and restricted-total richness the amount effect exceeded the fragmentation effect under every metric, with ratios ranging from 2.42 to 8.50, both extremes occurring under NLSI in the raw and residualized formulations respectively; the metric-specific values are shown in Figure 5.3.

Residualization sharpened direction, but its effect on interval support differed among strata. For specialist richness, interval support increased from 40.4 to 53.2 %. Facultative richness moved from 69 to 84 % negative while interval support fell from 34 to 21 %, and restricted total from 71 to 83 % while support fell from 40 to 27 %. Residualization therefore increased directional consistency across the three strata, while interval support increased for specialist richness and declined for facultative and restricted-total richness; the cause of that divergence is not established here.

Interaction effects: raw and residualized formulations

Each stratum comprised 1,296 cells, divided equally between the raw and residualized formulations. All 648 raw cells were eligible in every stratum. In the residualized formulation, 648 cells were eligible for specialist and non-forest richness, 647 for restricted total richness, and 645 for facultative richness; the four ineligible cells arose where the matched additive comparison model failed to converge although the interaction model succeeded, so no parsimony baseline exists. A separate set of four cells, identified by a different criterion, is the one in which the selected correlated random-slope structure was replaced by the no-random-slope structure. Interaction coefficients were negative in 88.0 % of raw and 78.5 % of residualized forest-specialist cells, 71.3 and 77.6 % of restricted total, 68.1 and 71.5 % of facultative, and 35.8 and 36.1 % of non-forest cells. These are descriptive summaries of a dependent robustness space rather than independent tests.

Primary support, requiring both a parsimony gain over the matched additive model and structure-coefficient concordance, was reached in 214 of 648 raw and 197 of 648 residualized forest-specialist cells, 188 of 648 and 174 of 647 restricted total cells, 184 of 648 and 90 of 645 facultative cells, and 122 of 648 and 79 of 648 non-forest cells. Primary support occupied a

minority of the robustness space, so support frequency was interpreted jointly with all-cell direction, supported effect magnitudes, marginal-effect geometry, and the specialist bootstrap rather than as a stand-alone measure of evidence. Facultative richness lost half its support under residualization, the largest formulation effect in the design.

Direction showed moderate-to-high concordance across forest definitions, scales, and formulations, whereas magnitude was less stable. Matched sign concordance was 68 to 88 % between forest products at matched resolution, 73 to 83 % between resolutions within Hansen, and 61 to 84 % between the 70 and 90 % canopy thresholds. Formulation profiles are reported in Appendix 5.S5.4. Hansen at 1,000 metres yielded interaction coefficients approximately 1.8 times as large in absolute magnitude in the raw formulation, and approximately 3.2 times as large in the residualized formulation, relative to Tuanmu at matched resolution for forest specialists. Complete product, resolution, threshold, scale and formulation profiles are reported in Appendix 5.S5.4.

Treating the robustness space as a set of design factors rather than replicates sharpens this. The median interaction coefficient for forest-specialist richness is negative at all thirty marginal factor levels, ranging from −0.109 to −0.005, and negative in sixty-three of the seventy-two joint stream × threshold × metric × formulation profiles, each marginalized over the eighteen scale combinations. The nine departures form three complete threshold triplets rather than scattering: raw NLSI under Hansen 500 m, and residualized EDGED and PATCHD under Tuanmu 1,000 m, each positive at all three canopy thresholds and none exceeding +0.024. Restricted-total and facultative richness are negative in sixty and fifty-four of those profiles, departing at the marginal level only at the three smallest scale combinations, which are also those with least support. Computed within formulation, fragmentation metric shows the largest median absolute deviation of factor-level medians in all eight stratum × formulation profiles (0.014 to 0.056, against 0.009 to 0.039 for scale), and it governs guild specificity. Pooling both formulations, on matched cells the specialist-minus-non-forest paired difference falls from −0.124 under NLSI and −0.120 under DCLUMP to −0.059 under PATCHD and only −0.023 under edge density, where sign agreement between the two strata rises from 23.5 to 62.3 %. The formulation-separated values are shown in Figure 5.6. Within the scale matrix the gradient is carried by predictor extent rather than response grain: coefficients strengthen with extent to roughly 70 to 85 km before turning back at 100 km, while medians differ little across grains once

extent is held common. Canopy threshold raises primary support in five of the six specialist, facultative and restricted-total formulation profiles and lowers it for non-forest richness in both. Predictor extent carried the ordered scale pattern, whereas response-grain trajectories were comparatively similar within metrics (Figure 5.4).

Environmental adjustment did not alter the direction of any interaction carried into the controls analysis. Across the 1,248 primary-supported cells — 362 restricted-total, 411 specialist, 274 facultative and 201 non-forest — adding AET and TOPOD produced zero sign reversals in all four response strata. Median changes in the interaction coefficient were −4.11 % for restricted-total richness, −1.24 % for specialist richness, +1.33 % for facultative richness and +0.75 % for non-forest richness, with the specialist interquartile range running −12.8 to +8.0 %. The leave-one-control-out specifications likewise retained the interaction direction in every converged cell, indicating that directional stability did not depend on either AET or TOPOD alone. EVICV was evaluated as a secondary sensitivity for forest-specialist richness only: adding it to AET and TOPOD retained the interaction sign in all 411 specialist cells with a median coefficient change of approximately +2 %, so vegetation texture had little influence on the direction of the detected specialist surfaces. The primary spatially adjusted models remain the inferential basis; the environmental-control models evaluate the robustness of surfaces already selected under the primary criterion.

Marginal-effect curves and cross-stratum contrasts

The marginal fragmentation effect became more negative toward higher forest cover in 88.0 % of raw and 78.5 % of residualized specialist cells, 71.3 and 77.6 % of restricted-total cells, and 68.1 and 71.5 % of facultative cells. These are the interaction-direction percentages of the preceding subsection, since the sign of the interaction coefficient defines the direction of change. Non-forest richness moved in the opposite direction, becoming more positive toward higher cover in all but 35.8 and 36.1 % of cells; its median marginal effect rose from +0.010 at the 5th cover percentile to +0.096 at the 95th. Median specialist marginal effects across all eligible cells ran from +0.011 at the 5th percentile to −0.173 at the 95th in the raw formulation, and from +0.006 to −0.178 in the residualized formulation, with strong metric- and formulation-specific differences in trajectory (Figure 5.5).

Bootstrap percentile intervals were available for the 214 raw and 197 residualized specialist cells carrying primary support. The two sets are within-formulation descriptive summaries rather than a direct formulation contrast: the raw and residualized predictors were standardized separately, and the support-selected populations differ in metric composition, edge-density cells making up 32 % of the raw and 9 % of the residualized set against 6 and 24 % for NLSI. Metric-resolved comparisons are reported in Figure 5.5 and Appendix 5.S5.2. Median marginal effects were +0.120 at the 5th cover percentile, −0.140 at the 75th and −0.275 at the 95th in the raw formulation, and +0.043, −0.180 and −0.290 in the residualized formulation. Intervals excluded zero in 67.8 % of raw and 91.9 % of residualized cells at the 75th percentile, and in 88.3 and 95.9 %, respectively, at the 95th. At the 75th cover percentile these estimates correspond to approximately 13 % lower specialist richness per standard deviation of fragmentation in the raw formulation and 16.4 % lower richness in the residualized formulation, increasing to approximately 24 and 25 % at the 95th. Values at the 25th and 50th percentiles are reported in Appendix 5.S5. Headline additive and conditional estimates for every response stratum are summarized in Table 5.2.

The direction of interval support at the lower endpoint was heterogeneous. Among the 214 raw cells, 100 intervals lay entirely above zero at the 5th cover percentile, 3 lay entirely below it and 111 overlapped zero. The corresponding residualized counts were 40, 8 and 149 among 197 cells. In 79 of 214 raw cells and 38 of 197 residualized cells, the interval was positive at the 5th percentile and negative at the 95th. These are pointwise percentile intervals evaluated separately at each endpoint rather than a simultaneous confidence procedure for the location of a zero crossing, and no operational crossing rule was implemented (Appendix 5.S3.1).

Opposite-signed endpoint intervals were concentrated in specific stream–metric–formulation combinations rather than distributed evenly across the design. They occurred in 51 of 69 primary-supported raw Hansen EDGED cells but in none of 10 residualized Hansen EDGED cells. In the residualized formulation, 23 of 40 Hansen PATCHD cells carried the pattern. By contrast, no such endpoint pattern was observed among 27 raw or 29 residualized Tuanmu DCLUMP cells. Response grain showed no consistent profile in either formulation.

When the fragmentation main effect and interaction were both negative, the lower-end point estimate became positive whenever the absolute standardized cover value at the 5th

percentile exceeded the absolute main-effect-to-interaction ratio. This condition was met in at least half and usually most primary-supported cells across metrics, so it did not explain why opposite-signed endpoint intervals concentrated in particular metric profiles. Pointwise support at both endpoints additionally required marginal effects of sufficient absolute magnitude relative to their bootstrap uncertainty, and it was this condition that differed sharply among metrics and formulations. The opposite-signed endpoint pattern was therefore metric- and formulation-sensitive rather than a general feature of the specialist response. Complete stream × metric × formulation, threshold, grain and extent profiles are reported in Appendix 5.S5 as post hoc descriptive summaries.

Across all 648 specialist cells in each formulation, point-estimate zero crossings occurred within the observed cover domain in 51.2 % of raw and 58.5 % of residualized cells. These crossings were descriptive properties of the fitted marginal-effect lines and were not treated as interval-supported mechanism classifications.

Seventeen of the 411 primary-supported specialist interactions were positive, two in the raw and fifteen in the residualized formulation, and every one occurred at the 10 km response grain. Six arose under residualized PATCHD at predictor extents of 10 to 40 km, eight under residualized edge density at Tuanmu 1,000 m at extents of 40 to 100 km, and three under DCLUMP or NLSI where the predictor extent equalled the response grain. Within the residualized-PATCHD cells the other configuration metrics ran predominantly neutral or negative, with DCLUMP and NLSI interval-supported negative in two cases, while none of the corresponding raw PATCHD interactions was supported (Appendix 5.S5). Cross-stratum contrasts were computed on matched design cells. Specialist and non-forest interaction coefficients shared sign in only 40.6 % of matched raw cells and 35.2 % of residualized cells, with median paired differences of −0.107 and −0.087. The contrast is metric-dependent: non-forest interactions were predominantly positive under DCLUMP, NLSI and PATCHD (8 to 31 % negative) but predominantly negative under EDGED (76.5 % in both formulations), so the mirror is carried by three of the four configuration metrics rather than all four. Specialist coefficients were more negative than restricted-total coefficients on matched cells, consistent with dilution of a response concentrated among forest-dependent species when guilds are aggregated. Because restricted total richness is the arithmetic sum of the three guilds, these comparisons are linked decompositions of the same assemblages rather than independent replications, and no magnitude-

ratio threshold was applied. Formulation-separated coefficient levels and matched specialist–non-forest differences show this metric dependence directly (Figure 5.6).

Discussion

Conditional fragmentation effects under amount control

Negative β_{frag} would indicate a fragmentation effect at mean standardized forest cover, whereas negative $\beta_{frag \times amount}$ would indicate that this effect changes along the forest-cover gradient. Both are observed. Additive fragmentation coefficients, which estimate the average effect after linear amount control, were negative across most cells in specialist, facultative and restricted-total richness but modest in magnitude, and the interaction-model coefficient is a distinct estimand that differs from them, most markedly for residualized specialist richness. Forest-amount effects exceeded fragmentation effects under every metric, by roughly two to nine times in absolute magnitude, and were much more consistently interval-supported. The interaction was also negative across most cells in these three response strata, most consistently for forest specialists. Configuration is therefore not irrelevant, but neither is its effect invariant along the amount gradient, which is the component of the habitat-amount hypothesis this design can address (Fahrig 2013, 2021; Saura 2021), and the invariance component of the configuration-independent null is not supported for these responses. We do not read this as an affirmative rejection of that null. No equivalence bound was declared in advance, so a near-zero interaction would have been inconclusive rather than confirmatory, and the asymmetry is stated rather than exploited.

Residualization clarifies amount–configuration geometry

Residualization is a robustness axis in this chapter, not a causal separation of configuration from amount. It removes the component of each metric predicted by a fitted amount–configuration model and defines an operational estimand — the part of configuration not predicted by forest amount under that model — so whether residual dependence is negligible is an empirical property rather than one the method guarantees (Koper et al. 2007; Smith et al. 2009), and the quantity it defines is operational, not causal (Lundberg et al. 2021). Both formulations are carried because the amount–configuration relationship is substantially nonlinear. At the metric-summary level the fitted smooths explained approximately 8 to 32

additional percentage points of variation relative to the corresponding median linear r^2, with the largest gap under PATCHD. That geometry differs among metrics rather than being a shared property of fragmentation. Forest amount accounts for roughly a fifth of the variation in DCLUMP but over three quarters in edge density, so residualization removes a very different share of each metric and rescales each by a different factor. A standardized residual unit therefore spans a smaller interval of the original standardized metric in edge density than in DCLUMP, which is why raw and residualized coefficients may be compared in direction, support and fitted geometry but not as a fraction of an effect removed. It also gives a measurement reason, alongside the ecological ones, for the distinctive behaviour of edge density. Among the four metrics, its raw variation contains the largest amount-predictable component. The two are similar in pooled central tendency while disagreeing on sign in a minority of individual cells, but that pooled similarity should not be read as agreement. The supported populations differ in metric composition between formulations, and metric-specific changes are asymmetric, with interaction magnitude increasing under DCLUMP and NLSI and decreasing under PATCHD and edge density. No headline conclusion rests solely on either formulation; where they diverge, most visibly in the threshold response of forest-specialist richness, the divergence is reported rather than resolved in favour of either.

Guild partitioning localizes fragmentation responses

Guild partitioning indicated that the weaker aggregate response reflected dilution rather than uniformly weak responses across the assemblage. Specialist interactions were more negative than restricted-total interactions on matched cells and more negative than facultative interactions in both raw and residualized formulations. Because restricted total richness is the arithmetic sum of the three guild components, these comparisons localize which component carries the response rather than providing independent ecological replication. Concentration among habitat specialists and dilution in aggregate richness are both expected from trait-based sensitivity (Henle et al. 2004; Vetter et al. 2011; Matthews et al. 2014; Valente and Betts 2019).

The specialist–non-forest contrast was sharper still. The two guilds shared interaction sign in only a minority of matched cells, so opposite directions predominated. This pattern weighs against artefacts expected to act uniformly through forest-cover–fragmentation geometry, effort–richness scaling or residualization. It does not exclude composition-dependent sampling, because

guild differences in response distribution, composition and detectability can translate the same predictor geometry into different fitted coefficients (Section IX). The specialist–non-forest sign mirror was pronounced under DCLUMP, NLSI and PATCHD but disappeared under edge density because the non-forest interaction was itself negative. The credibility of this localization rests on choices made before the guild comparison. Predictor geometry was characterized rather than assumed, support did not rest on coefficient sign alone, and metric dependence was retained rather than averaged away. Forest definition, scale and formulation define the domain within which the result holds, and none of that converts an observational association into a causal one (Fletcher et al. 2023). The formulation-specific contrasts also differed in direction under edge density, so the guild-specific interpretation is not metric-invariant. Structural configuration does not map uniquely onto ecological function. Positive edge associations may occur in relatively permeable matrices, whereas connected interior habitat can become more important where matrix hostility and edge effects are severe, and these relationships can change with analysis scale (Dennis et al. 2024).

Interaction geometry is most consistent with extinction filtering

Specialist interaction coefficients were predominantly negative, and their marginal-effect geometry was most consistent with extinction filtering rather than the positive interaction predicted by scarcity and coexistence dynamics, under which fragmentation should be most damaging at low cover and weaken, potentially reversing, as cover increases (Table 5.1). Specialist effects became increasingly negative across the upper half of the domain, while low-cover effects were often weak or unresolved but positive with interval support in a substantial subset. At the 75th cover percentile, a one-standard-deviation increase in fragmentation corresponded to approximately 13 % lower specialist richness in the raw formulation and 16.4 % in the residualized formulation, increasing to 24 to 25 % at the 95th. Interval support was more frequent at the 95th percentile than at the 75th in both formulations. This direction and high-cover intensification recurred broadly across forest definitions, metrics, scales and formulations, although support was not universal and magnitude remained design-dependent.

Residualization preserved the broad negative high-cover pattern but altered magnitude and support asymmetrically among metrics. Although the interaction coefficient itself was somewhat less consistently negative after residualization, the fragmentation effect at mean standardized

cover was negative in a larger share of residualized than raw specialist cells. Interval support at the 75th percentile was also higher after residualization. The median high-cover marginal effect was similar between formulations in pooled terms, but the two supported populations differ in metric composition — edge-density cells make up 32 % of the raw and 9 % of the residualized set, NLSI cells 6 and 24 % — so that similarity is not a like-for-like comparison and the metric-resolved values are given in Figure 5.5 and Appendix 5.S5.2. Because residualization restricts fragmentation to variation not predicted by forest amount, the persistence of negative direction and interval support across several metric profiles indicates that the negative high-cover configuration association was not carried solely by the amount-coupled component of the raw metrics. It remains evidence consistent with, rather than proof of, a fragmentation-per-se association because residualization cannot create the equal-amount configuration contrasts required for causal identification.

The dominant cross-metric response remained negative and concentrated among forest specialists, supporting extinction filtering as the mechanism family most consistent with the overall interaction structure. The seventeen positive specialist interactions were not evidence for a general benefit of fragmentation per se, and they were heterogeneous rather than a single metric-specific process. Facultative coefficients shared the positive sign in all seventeen matched design cells and restricted-total coefficients in sixteen, so the positive specialist interaction was not specialist-exclusive in direction. That parallel is what the scarcity and coexistence frameworks predict for guild direction (Table 5.1), so it neither identifies a guild-specific process nor rebuts those frameworks. That concentration is a property of the design space rather than of these seventeen cells. The residualized edge-density cluster under Tuanmu 1,000 m stands apart from the short-support cells, with non-forest coefficients negative in seven of eight, and it remains an open metric–mechanism alignment issue (Gengler et al. 2024; Dennis et al. 2024). The predominant high-cover intensification is consistent with fragmentation-sensitive taxa having already been filtered from many low-cover assemblages, leaving fewer responsive species, while a larger sensitive complement persists at higher cover (Balmford 1996; Betts et al. 2019). If we interpret this geometry through extinction filtering, prior loss of fragmentation-sensitive taxa may also attenuate the configuration signal that remains detectable in contemporary low-cover assemblages. We would therefore expect the fitted configuration associations to under-represent cumulative fragmentation costs or latent susceptibility after

filtering has occurred, although our cross-sectional design cannot quantify that difference. Coexistence and scarcity arguments predict the opposite conditionality, configuration mattering most where habitat is scarcest (Rybicki et al. 2020; Zhang et al. 2024). The distinction rests on the sign of the change in the fragmentation effect along the cover gradient, not on a negative average coefficient, which both frameworks admit. That interpretation classifies consistency rather than demonstrating demographic loss. Cross-sectional richness cannot observe the process, and similar high-cover intensification could arise if low-cover assemblages were dominated by broadly tolerant species for reasons not specific to fragmentation. Richness also cannot distinguish ordered species loss from replacement. A turnover–nestedness decomposition would expose that distinction, although compositional pattern alone would not identify the mechanism uniquely (Socolar et al. 2016). Across all specialist cells, point estimates crossed zero within the observed cover domain in about half of both formulations. These crossings ran from positive effects at low cover toward increasingly negative effects at high cover, the reverse of the negative-to-positive high-cover reversal that would distinguish coexistence from scarcity. They therefore provide no support for coexistence.

The pointwise-supported positive-to-negative endpoint pattern nevertheless departs from the prospective filtering prediction, under which negative effects strengthen toward high cover without sign reversal (Table 5.1). It occurred chiefly in raw Hansen EDGED and residualized Hansen PATCHD cells and was not observed among the primary-supported Tuanmu DCLUMP cells in either formulation. Point-estimate sign changes were much more widespread across metrics, so the concentration of pointwise-supported endpoint patterns reflected differences in absolute marginal-effect magnitude and bootstrap uncertainty rather than differences in whether the fitted line changed sign. Residualization changed the estimand and the standardized interaction magnitude differently among metrics, so disappearance of EDGED interval support after residualization cannot be treated as an independent refutation of the raw EDGED pattern. The limited convergence of positive low-cover support across stream–metric combinations nonetheless weighs against interpreting it as a general benefit of fragmentation for forest specialists at low cover, whereas increasingly negative high-cover effects recurred broadly. Whether the positive endpoint estimates reflect a genuine low-cover response, metric geometry and residual amount coupling, or the composition-dependent sampling considered in Section IX cannot be resolved here.

Forest-definition choice alters effect magnitude and detectability

Stricter canopy thresholds reclassify degraded and open-canopy vegetation from forest to matrix, altering the landscape boundary the predictors represent without changing the observed assemblages (Almeida-Gomes et al. 2016; Dennis and Huck 2025). In these data they increased detectability more consistently than they increased interaction magnitude. In the raw formulation, raising the threshold from 70 to 90 % made interactions progressively more negative for forest-specialist, facultative and restricted-total richness and attenuated the positive non-forest interaction toward zero. After residualization the same magnitude pattern persisted for facultative and restricted-total richness but reversed for specialists, whose coefficient became less negative at the strictest threshold. Primary support (an operational measure combining the parsimony gain with structure-coefficient concordance, not a general index of statistical power) generally increased with stricter thresholds across specialist, facultative and restricted-total richness, but not uniformly. Raw specialist support peaked at the intermediate threshold, whereas residualized specialist support and non-forest support in both formulations declined. That response- and formulation-dependence makes the threshold sweep a habitat-delineation dimension rather than a uniform amplifier of specialist coefficients (Table 5.S3.3). The residualized-specialist exception was concordant across coefficient magnitude and support frequency, although those summaries derive from the same fitted cells (the coefficient sign enters the support decision) and are not independent evidence against sampling variation. One possible explanation is that stricter thresholds intensify class-area constraints on configuration metrics (Neel et al. 2004), so that residualization removes a different share of configurational variance where forest class area is smallest; the threshold-specific geometry diagnostics of Appendix 5.S2 are required to evaluate that interpretation, which is not established here. Restricted-total richness tracked the facultative pattern across the threshold sweep, indicating that aggregate design-factor patterns largely reflected the forest-associated components.

Forest-product choice displaced factor-level medians more than canopy-threshold choice did for specialist, facultative and restricted-total richness, though not for non-forest richness, where the two were comparable. Comparing median coefficients, Hansen at 1,000 m yielded specialist interactions about 1.8 times as large in absolute magnitude as Tuanmu at matched resolution in the raw formulation and about 3.2 times as large in the residualized formulation, and support was markedly sparser under Tuanmu. Because predictors were standardized on

definition-specific distributions, that ratio is descriptive rather than a common-unit biological comparison. Since the two products differ in temporal anchor, classification framework and native resolution simultaneously, this contrast cannot attribute the difference to any one of them. It does establish that magnitude, and to a lesser extent detectability, are contingent on the operational definition of forest in a way that a single-product analysis would not reveal.

Resolution acted consistently within the Hansen product. The coarser 1,000 m stream returned more negative interaction coefficients than the 500 m stream across all four response strata, on matched cells marginalized over metric, canopy threshold and scale combination. Resolution and threshold are therefore not interchangeable sensitivity axes, and reporting nine definitions as nine products would conflate three distinct contrasts.

Threshold change also alters guild contrasts asymmetrically. The predominantly positive non-forest interaction attenuated toward zero at the strictest threshold, the largest absolute threshold response in the raw formulation; after residualization the largest shift was that of specialist richness, in the opposite direction. Reclassification changes the forest amount and configuration attributed to each landscape without changing the observed species or their guild assignments. The stronger response of non-forest coefficients therefore reflects how those assemblages sit relative to a redefined forest–matrix boundary, not a change in what was sampled.

Predictor extent carries the ordered scale dependence

Predictor extent, rather than response grain, carried the ordered scale dependence of the interaction. Extent–coefficient correlations were negative in twenty-two of twenty-four grain × stratum × formulation combinations. By contrast, specialist medians differed little among response grains at predictor extents shared by all three grains. Conditional fragmentation effects generally strengthened toward broader predictor extents before levelling or turning back at the broadest extent. Predictor extent therefore formed the strongest ordered gradient, although fragmentation metric produced the largest overall displacement of factor-level medians in every response stratum and formulation.

The scale profile is compatible with both ecological scale dependence and increasing predictor stability, and the data do not separate them (Fletcher et al. 2023). The metric-resolved profiles are consistent with recent evidence that some configuration metrics operate over wider

characteristic extents than forest cover while others do not (Gengler et al. 2024). We retain coefficient trajectories across the factorial grain × extent design rather than selecting an information-criterion-optimal buffer, so these profiles identify a metric-dependent scale of association rather than formal scales of effect. Configuration may genuinely operate at broad extents where dispersal and landscape-scale population connectivity are more influential than local edge processes (Miguet et al. 2016; Jackson and Fahrig 2015). Alternatively, broader extents may average over more landscape and yield more stable predictor values, improving estimation without implying a broader ecological scale of effect. The levelling and slight turnback at the broadest extents do not distinguish these explanations because both ecological responses and gains in predictor stability can saturate with increasing extent. The scale combinations form a dependent sensitivity profile, not independent replications. A study using only the smallest extent could therefore have understated the amount-conditioned fragmentation signal. The narrowest combinations are also those in which the predictor window and the response unit share a footprint, and interaction coefficients differ there. Across all fitted cells, positive interaction coefficients were concentrated toward the narrowest predictor extents in every response stratum, though the decline was not monotonic, and every primary-supported positive specialist interaction occurred at the 10 km response grain. Same-support geometry is therefore one feature of that positive tail rather than a general explanation for it, since the tail also contains a residualized edge-density cluster at broad extents. This remains an open interpretation requiring a dedicated matched-support analysis and is not carried further here.

Conditionality helps explain divergent fragmentation estimates

The within-biome pattern parallels the global finding that fragmentation effects on animals were most negative in regions where historical habitat loss had been least severe (Betts et al. 2019). That analysis inferred filtering from cross-biome variation in disturbance history, whereas the present design recovers a comparable gradient within one biome by estimating forest amount and configuration jointly across contemporary landscapes. The conditional architecture recovered here is not specific to forest amount as the conditioning variable. Thresholds in matrix hostility likewise determine which landscape attribute drives mammalian richness, with structural, functional and compositional attributes each dominating in different matrix contexts and each capable of positive or negative effects (Dennis et al. 2024). The two analyses provide convergent observational evidence rather than independent confirmation of extinction filtering,

because both infer the mechanism from spatial response patterns rather than observing demographic loss directly. The within-biome design avoids the specific between-region contrast on which the global analysis depends, but it remains cross-sectional and observational.

The relationship to the habitat amount hypothesis (Fahrig 2013) and the global species-density synthesis supporting it (Watling et al. 2020) is one of scope rather than contradiction. Those analyses estimate an average configuration effect after amount control, whereas the present design estimates how that effect varies with amount. An average coefficient can be attenuated toward zero when sampled landscapes are concentrated where marginal effects are weak or unresolved, even if effects become strongly negative at higher cover. This does not require a supported positive response at low cover, although such effects did receive interval support in a subset concentrated in particular combinations of metric, forest product and formulation. The interaction is the minimum structure needed to distinguish gradient invariance from averaging over a conditional surface.

The coefficients estimated here are conditional associations of configuration at fixed forest amount, not total effects of forest loss transmitted through changes in configuration. Estimating the latter requires a path framework that separates direct, indirect and total effects. An Atlantic Forest analysis using that framework reported direct configuration effects alongside indirect pathways through habitat loss (Püttker et al. 2020). The present design does not identify those pathways and therefore cannot determine how its conditional configuration coefficient compares with a total-effect estimand. Forest-amount effects were substantially larger than average fragmentation effects, while the interaction shows that the conditional configuration slope is not constant along the amount gradient. Sensitivity to forest definition is expected when mapped vegetation imperfectly represents usable habitat and edge-related responses depend on that boundary (Almeida-Gomes et al. 2016; Willmer et al. 2022). Both detectability and magnitude varied among forest definitions, although the product contrast jointly changes temporal anchor, classification framework and native resolution and cannot identify which feature drives the difference.

Cross-study additive estimates may therefore differ partly because an additive coefficient averages over the sampled cover distribution. Studies sampling different portions of a conditional gradient can recover different average effects even when the response surfaces are

similarly structured. Scale choice provides a further reason, because small predictor extents can understate an amount-conditioned signal whose ecological and analytical components remain inseparable (Fletcher et al. 2023).

Limitations and future tests

Fragmentation per se remains an operational estimand. Residualizing each configuration metric against forest cover removes the modelled nonlinear dependence described in Section II without creating matched landscapes that differ only in configuration, and the raw and residualized formulations remain non-symmetric views whose standardized coefficients are not common-unit magnitudes. Pairing both formulations across metrics, forest products, canopy thresholds and spatial scales keeps the conclusion from resting on one adjusted coefficient. The paired formulations support an amount-dependent configuration association that persists after modelled amount dependence is removed. They do not estimate the effect of rearranging a fixed amount of forest while other landscape properties are held constant. DCLUMP, NLSI, PATCHD and edge density quantify different structural dimensions, and the residualized-PATCHD departure reported in Section IV marks the limit of cross-metric generality. Forest-cover percentiles are specific to each design cell, so high cover does not define one universal canopy-cover threshold across products, definitions and scales. Coefficients are also standardized on definition-specific predictor distributions, which makes cross-definition magnitude ratios descriptive rather than common-unit comparisons. The nested grain–extent matrix establishes scale coherence and extent sensitivity without separating ecological scale from predictor smoothing, buffer overlap and spatial dependence.

The interaction geometry is consistent with extinction filtering without observing demographic loss. Cross-sectional richness cannot distinguish historical filtering from ongoing extinction debt, delayed recolonization, compositional replacement or other pathways that generate a similar spatial pattern. Composition-dependent sampling remains the strongest live threat to the extinction-filtering interpretation of the specialist–non-forest sign mirror. Finite survey effort or edge- and road-biased placement could change which guild is detected without equivalent demographic change. The disappearance of the mirror under edge density limits one narrow placement-artefact account, and survey-point proximity to edges and roads was not measured and is not equivalent to landscape-level edge density. Survey duration is a further

unmodelled dimension of sampling intensity: the effort covariate represents reported effort and the number of consolidated source assemblages, so assemblages differing in the number of years over which surveys accumulated are treated alike. The guild pools also differ in taxonomic composition, which cannot be separated from forest affiliation in this contrast. Extinction filtering therefore remains the mechanism family most consistent with the observed geometry, and the amount-dependent specialist concentration remains the empirical finding.

Landscape structure enters as contemporary context, not as a reconstructed exposure history. The forest products differ from one another and from the biodiversity surveys in temporal anchor, and stability in regional forest extent does not establish that the configuration of individual landscapes has been stable over the period relevant to assemblage response. Configuration can be rearranged at near-constant cover, so forest amount and configuration need not share a change history, and an interaction estimated on contemporary landscapes may combine present conditioning with the legacy of earlier rearrangement. For the high-cover interpretation, this ambiguity matters because substantial rearrangement can occur without a corresponding loss of forest amount. Historical landscape reconstruction paired with the same assemblages would test whether the high-cover intensification tracks present structure or accumulated change.

The targeted environmental-control arm tests whether the interpreted surfaces retain their direction under measured environmental adjustment. It does not establish how those controls would alter support classification across the eligible factorial. Absolute cover distributions or narrower amount bands, historical landscape reconstruction, and within-taxon and within-method assessment of the specialist–non-forest contrast would clarify the remaining uncertainty. Within the present design, configuration associations vary with forest amount, and the strongest negative responses concentrate in forest-specialist richness.

Conservation implications

Across specialist, facultative and restricted-total richness, we found larger and more consistently supported contemporary associations with forest amount than with configuration. Standardized amount coefficients exceeded fragmentation coefficients under every metric by roughly two to nine times, but these ratios compare fitted associations rather than the underlying

ecological importance of habitat loss and configuration. The much more consistent support for forest amount provides the firmer evidential basis for management priority. Where management priorities force a choice, we therefore support forest retention and restoration as the first priority, while our results provide no basis for trading cover for configuration.

The interaction changes where configuration management is most consequential. Negative specialist effects intensified toward higher cover, so further subdivision is most concerning in well-forested landscapes that still retain sensitive taxa. Low-cover effects were often weak or unresolved, although low-cover positive support was metric- and product-specific (Section VII), whereas negative effects were substantially more consistent at high cover; the pattern therefore supports protecting the internal configuration of high-cover landscapes without treating configuration in already-degraded landscapes as ecologically irrelevant.

Aggregate-richness targets can understate configuration costs to forest-dependent taxa because specialist interactions were substantially stronger than restricted-total interactions on matched cells. High cover denotes the upper part of each design cell's forest-cover distribution, not one absolute canopy-cover value. Applying the high-cover recommendation to a particular landscape therefore requires its absolute forest-cover distribution. The quoted magnitudes are per standard deviation of a landscape metric, not per management unit, so applying them also requires the distribution of the configuration metric and cannot be inferred from the coefficient alone.

Conclusions

Fragmentation effects on Atlantic Forest vertebrate richness were conditional on forest amount rather than uniform along the cover gradient. Standardized forest-amount coefficients exceeded fragmentation coefficients under every fragmentation metric for specialist, facultative and restricted-total richness, by roughly two to nine times in absolute magnitude, but configuration costs intensified toward higher cover, especially for forest specialists. This gradient was weaker in restricted-total richness, consistent with dilution, and opposed by non-forest richness under three of the four metrics, although edge density removed the guild sign reversal. For forest specialists, the negative interaction direction recurred broadly across the robustness space, while primary support occupied a minority of cells and magnitude remained

contingent on metric, forest definition and scale. The pattern is most consistent with extinction filtering, but cross-sectional richness cannot demonstrate demographic loss, and composition-dependent sampling remains a plausible alternative. In this system, forest retention and restoration remain the first priority, while configuration management is most consequential in well-forested landscapes that still retain sensitive taxa.

References

Almeida-Gomes M, Prevedello JA, Crouzeilles R (2016) The use of native vegetation as a proxy for habitat may overestimate habitat availability in fragmented landscapes. Landscape Ecology 31:711-719.

Amatulli G, Domisch S, Tuanmu M-N, Parmentier B, Ranipeta A, Malczyk J, Jetz W (2018) A suite of global, cross-scale topographic variables for environmental and biodiversity modeling. Scientific Data 5:180040.

Andrén H (1994) Effects of habitat fragmentation on birds and mammals in landscapes with different proportions of suitable habitat: a review. Oikos 71:355-366.

Arel-Bundock V, Greifer N, Heiss A (2024) How to interpret statistical models using marginaleffects for R and Python. Journal of Statistical Software 111:1-32.

Balmford A (1996) Extinction filters and current resilience: the significance of past selection pressures for conservation biology. Trends in Ecology and Evolution 11:193-196.

Betts MG, Wolf C, Pfeifer M, et al. (2019) Extinction filters mediate the global effects of habitat fragmentation on animals. Science 366:1236-1239.

Bovendorp RS, Villar N, de Abreu-Junior EF, et al. (2017) Atlantic small-mammal: a dataset of communities of rodents and marsupials of the Atlantic forests of South America. Ecology 98:2226.

Brooks ME, Kristensen K, van Benthem KJ, et al. (2017) glmmTMB balances speed and flexibility among packages for zero-inflated data. The R Journal 9:378-400.

Burnham KP, Anderson DR, Huyvaert KP (2011) AIC model selection and multimodel inference in behavioral ecology: some background, observations, and comparisons. Behavioral Ecology and Sociobiology 65:23-35.

Crase B, Liedloff AC, Wintle BA (2012) A new method for dealing with residual spatial autocorrelation in species distribution models. Ecography 35:879-888.

Dennis M, Huck JJ (2025) The continuity-contiguity problem in fragmentation-biodiversity research. Journal of Biogeography 52:e15077.

Dennis M, Huck JJ, Holt CD, da Conceição Bispo P, McHenry E, Speak A, James P (2024) Land-cover gradients determine alternate drivers of mammalian species richness in fragmented landscapes. Landscape Ecology 39:146.

Dormann CF, Elith J, Bacher S, et al. (2013) Collinearity: a review of methods to deal with it and a simulation study evaluating their performance. Ecography 36:27-46.

Ewers RM, Didham RK (2006) Confounding factors in the detection of species responses to habitat fragmentation. Biological Reviews 81:117-142.

Fahrig L (2017) Ecological responses to habitat fragmentation per se. Annual Review of Ecology, Evolution and Systematics 48:1-23.

Fahrig L (2003) Effects of habitat fragmentation on biodiversity. Annual Review of Ecology, Evolution and Systematics 34:487-515.

Fahrig L (1997) Relative effects of habitat loss and fragmentation on population extinction. Journal of Wildlife Management 61: 603-610.

Fahrig L (2013) Rethinking patch size and isolation effects: the habitat amount hypothesis. Journal of Biogeography 40:1649-1663.

Fahrig L (2021) What the habitat amount hypothesis does and does not predict: a reply to Saura. Journal of Biogeography 48:1530-1535.

Fahrig L (1998) When does fragmentation of breeding habitat affect population survival? Ecological Modelling 105:273-292.

Fahrig L, Arroyo-Rodríguez V, Bennett JR, et al. (2019) Is habitat fragmentation bad for biodiversity? Biological Conservation 230:179-186.

Fieberg JR, Vitense K, Johnson DH (2020) Resampling-based methods for biologists. PeerJ 8:e9089.

Figueiredo MSL, Barros CS, Delciellos AC, et al. (2017) Abundance of small mammals in the Atlantic Forest (ASMAF): a data set for analyzing tropical community patterns. Ecology 98:2981.

Fischer J, Lindenmayer DB (2007) Landscape modification and habitat fragmentation: a synthesis. Global Ecology and Biogeography 16:265-280.

Flather CH, Bevers M (2002) Patchy reaction-diffusion and population abundance: the relative importance of habitat amount and arrangement. The American Naturalist 159:40-56.

Fletcher RJ Jr, Betts MG, Damschen EI, et al. (2023) Addressing the problem of scale that emerges with habitat fragmentation. Global Ecology and Biogeography 32:828-841.

Fletcher RJ Jr, Didham RK, Banks-Leite C, et al. (2018) Is habitat fragmentation good for biodiversity? Biological Conservation 226:9-15.

Fletcher RJ Jr, Fortin M-J (2018) Spatial Ecology and Conservation Modeling. Springer, Cham.

Frazier AE, Kedron P (2017) Landscape metrics: past progress and future directions. Current Landscape Ecology Reports 2:63-72.

Freckleton RP (2002) On the misuse of residuals in ecology: regression of residuals vs. multiple regression. Journal of Animal Ecology 71: 542-545.

Gengler NW, Acevedo MA, Branch LC (2024) Habitat configuration influences mammal populations at a wider spatial extent than habitat composition: a meta-analysis of forest mammal datasets. Landscape Ecology 39:2.

Hansen MC, Potapov PV, Moore R, et al. (2013) High-resolution global maps of 21st-century forest cover change. Science 342:850-853.

Hanski I (2015) Habitat fragmentation and species richness. Journal of Biogeography 42:989-993.

Harrison XA (2014) Using observation-level random effects to model overdispersion in count data in ecology and evolution. PeerJ 2:e616.

Hartig F (2022) DHARMa: Residual diagnostics for hierarchical (multi-level/mixed) regression models. R package version 0.4.6.

Hasui E, Metzger JP, Pimentel RG, et al. (2018) ATLANTIC BIRDS: a data set of bird species from the Brazilian Atlantic Forest. Ecology 99:497.

Hatfield JH, Barlow J, Joly CA, et al. (2020) Mediation of area and edge effects in forest fragments by adjacent land use. Conservation Biology 34:395-404.

Henle K, Davies KF, Kleyer M, Margules C, Settele J (2004) Predictors of species sensitivity to fragmentation. Biodiversity and Conservation 13:207-251.

Hesselbarth MHK, Sciaini M, With KA, Wiegand K, Nowosad J (2019) landscapemetrics: an open-source R tool to calculate landscape metrics. Ecography 42:1648-1657.

IUCN (2012) IUCN Habitats Classification Scheme (Version 3.1). IUCN, Gland, Switzerland and Cambridge, UK.

Jackson HB, Fahrig L (2015) Are ecologists conducting research at the optimal scale? Global Ecology and Biogeography 24:52-63.

Koper, N., Schmiegelow, F.K. and Merrill, E.H., 2007. Residuals cannot distinguish between ecological effects of habitat amount and fragmentation: implications for the debate. Landscape Ecology 22(6):811-820.

Li H, Wu J (2004) Use and misuse of landscape indices. Landscape Ecology 19:389-399.

Lima F, Beca G, Muylaert RL, et al. (2017) ATLANTIC-CAMTRAPS: a dataset of medium and large terrestrial mammal communities in the Atlantic Forest of South America. Ecology 98:2979.

Lundberg I, Johnson R, Stewart BM (2021) What is your estimand? Defining the target quantity connects statistical evidence to theory. American Sociological Review 86:532-565.

Martínez-Lanfranco JA (2026) Habitat-loss-driven predictor coupling limits inference about the independent effects of configuration in additive habitat-amount models: implications for the fragmentation debate. bioRxiv preprint. https://doi.org/10.64898/2026.04.12.718042

Matthews TJ, Cottee-Jones HE, Whittaker RJ (2014) Habitat fragmentation and the species-area relationship: a focus on total species richness obscures the impact of habitat loss on habitat specialists. Diversity and Distributions 20: 1136-1146.

McGarigal K, Cushman SA, Ene E (2012) FRAGSTATS v4: Spatial Pattern Analysis Program for Categorical and Continuous Maps. University of Massachusetts, Amherst.

Miguet P, Jackson HB, Jackson ND, Martin AE, Fahrig L (2016) What determines the spatial extent of landscape effects on species? Landscape Ecology 31: 1177-1194.

Miller-Rushing AJ, Primack RB, Devictor V, Corlett RT, Cumming GS, Loyola R, Maas B, Pejchar L (2019) How does habitat fragmentation affect biodiversity? A controversial question at the core of conservation biology. Biological Conservation 232:271-273.

Muylaert RDL, Stevens RD, Esbérard CEL, et al. (2017) ATLANTIC BATS: a data set of bat communities from the Atlantic Forests of South America. Ecology 98:3227.

Neel MC, McGarigal K, Cushman SA (2004) Behavior of class-level landscape metrics across gradients of class aggregation and area. Landscape Ecology 19:435-455.

Newbold T, Hudson LN, Hill SLL, et al. (2015) Global effects of land use on local terrestrial biodiversity. Nature 520:45-50.

O'Hara RB, Kotze DJ (2010) Do not log-transform count data. Methods in Ecology and Evolution 1:118-122.

Paradis E, Schliep K (2019) ape 5.0: an environment for modern phylogenetics and evolutionary analyses in R. Bioinformatics 35:526-528.

Püttker T, Crouzeilles R, Almeida-Gomes M, et al. (2020) Indirect effects of habitat loss via habitat fragmentation: a cross-taxa analysis of forest-dependent species. Biological Conservation 241:108368.

R Core Team (2024) R: A Language and Environment for Statistical Computing. R Foundation for Statistical Computing, Vienna, Austria. https://www.R-project.org/.

Ray-Mukherjee J, Nimon K, Mukherjee S, Morris DW, Slotow R, Hamer M (2014) Using commonality analysis in multiple regressions: a tool to decompose regression effects in the face of multicollinearity. Methods in Ecology and Evolution 5:320-328.

Ribeiro MC, Metzger JP, Martensen AC, et al. (2009) The Brazilian Atlantic Forest: How much is left, and how is the remaining forest distributed? Implications for conservation. Biological Conservation 142:1141-1153.

Riva, F., Koper, N., & Fahrig, L. (2024). Overcoming confusion and stigma in habitat fragmentation research. Biological Reviews, 99, 1411–1424.

Rybicki J, Abrego N, Ovaskainen O (2020) Habitat fragmentation and species diversity in competitive communities. Ecology Letters 23:506-517.

Rybicki J, Hanski I (2013) Species-area relationships and extinctions caused by habitat loss and fragmentation. Ecology Letters 16:27-38.

Saura S (2021) The habitat amount hypothesis implies negative effects of habitat fragmentation on species richness. Journal of Biogeography 48:11-22.

Schielzeth H (2010) Simple means to improve the interpretability of regression coefficients. Methods in Ecology and Evolution 1:103-113.

Smith AC, Koper N, Francis CM, Fahrig L (2009) Confronting collinearity: comparing methods for disentangling the effects of habitat loss and fragmentation. Landscape Ecology 24:1271-1285.

Socolar JB, Gilroy JJ, Kunin WE, Edwards DP (2016) How should beta-diversity inform biodiversity conservation? Trends in Ecology and Evolution 31:67-80.

Souza Y, Gonçalves F, Lautenschlager L, et al. (2019) ATLANTIC MAMMALS: a data set of assemblages of medium- and large-sized mammals of the Atlantic Forest of South America. Ecology 100:e02785.

Stoffel MA, Nakagawa S, Schielzeth H (2021) partR2: partitioning R2 in generalized linear mixed models. PeerJ 9:e11414.

Title PO, Bemmels JB (2018) ENVIREM: an expanded set of bioclimatic and topographic variables increases flexibility and improves performance of ecological niche modeling. Ecography 41:291-307.

Trabucco A, Zomer R (2019) Global High-Resolution Soil-Water Balance. figshare Dataset. https://doi.org/10.6084/m9.figshare.7707605.v3

Tscharntke T, Tylianakis JM, Rand TA, et al. (2012) Landscape moderation of biodiversity patterns and processes - eight hypotheses. Biological Reviews 87:661-685.

Tuanmu M-N, Jetz W (2014) A global 1-km consensus land-cover product for biodiversity and ecosystem modelling. Global Ecology and Biogeography 23:1031-1045.

Tuanmu M-N, Jetz W (2015) A global, remote sensing-based characterization of terrestrial habitat heterogeneity for biodiversity and ecosystem modelling. Global Ecology and Biogeography 24:1329-1339.

Valente JJ, Betts MG (2019) Response to fragmentation by avian communities is mediated by species traits. Diversity and Distributions 25:48-60.

Valente JJ, Gannon DG, Hightower J, et al. (2023) Toward conciliation in the habitat fragmentation and biodiversity debate. Landscape Ecology 38:2717-2730.

Vancine MH, Duarte K da S, de Souza YS, et al. (2018) ATLANTIC AMPHIBIANS: a data set of amphibian communities from the Atlantic Forests of South America. Ecology 99:1692.

Vancine MH, Muylaert RL, Niebuhr BB, et al. (2024) The Atlantic Forest of South America: spatiotemporal dynamics of the vegetation and implications for conservation. Biological Conservation 291:110499.

Vetter D, Hansbauer MM, Végvári Z, Storch I (2011) Predictors of forest fragmentation sensitivity in Neotropical vertebrates: a quantitative review. Ecography 34: 1-8.

Villard M-A, Metzger JP (2014) Beyond the fragmentation debate: a conceptual model to predict when habitat configuration really matters. Journal of Applied Ecology 51:309-318.

Wang X, Blanchet FG, Koper N (2014) Measuring habitat fragmentation: an evaluation of landscape pattern metrics. Methods in Ecology and Evolution 5: 634-646.

Watling JI, Arroyo-Rodríguez V, Pfeifer M, et al. (2020) Support for the habitat amount hypothesis from a global synthesis of species density studies. Ecology Letters 23:674-681.

Willmer JNG, Püttker T, Prevedello JA (2022) Global impacts of edge effects on species richness. Biological Conservation 272:109654.

Wood SN (2017) Generalized Additive Models: An Introduction with R, 2nd edn. CRC Press, Boca Raton FL.

Zeileis A, Hothorn T (2002) Diagnostic checking in regression relationships. R News 2(3):7-10.

Zhang H, Chase JM, Liao J (2024) Habitat amount modulates biodiversity responses to fragmentation. Nature Ecology and Evolution 8:1437-1447.

Table 5.1. Predicted fragmentation–amount interaction signatures under the configuration-independent habitat-amount null, coexistence, scarcity and extinction filtering. HAH is a limiting null rather than a fourth mechanism: configuration irrelevance would require both the average fragmentation effect and the interaction to be near zero across the evaluation domain. Coexistence and scarcity share the same positive interaction prediction and are separable only by a sign reversal supported within the observed cover domain.

Framework	**Interaction prediction**	**Marginal-effect geometry and sign reversal**	**Guild prediction**	**Diagnostic interpretation**
HAH — configuration-independent limiting null (not an ecological mechanism).	≈ 0 (no interaction)	Flat marginal effect near zero across the cover gradient; no sign reversal.	Parallel across strata; no amount dependence	Interaction absent. Cannot be affirmed without a declared equivalence region; a non-supported, near-zero or imprecise estimate is inconclusive, not evidence for the null.

Framework	Interaction prediction	Marginal-effect geometry and sign reversal	Guild prediction	Diagnostic interpretation
H_1 — coexistence	$\beta > 0$ (negative at low cover → weak, neutral or positive at high cover).	Marginal effect weakens with cover and may cross zero at high cover†; sign reversal possible.	Same direction; specialists more negative but parallel.	Separable from scarcity only by an interval-supported zero crossing within the observed cover domain; no such rule is operationalized here (Appendix 5.S3.1).
H_1' — scarcity	$\beta > 0$ (strongest negative at low cover, weakening with increasing amount).	Monotonic negative marginal slope, steepest at low cover; no sign reversal.	Same direction; specialists more negative but parallel.	Shares the positive interaction prediction with coexistence; reported jointly as the positive-interaction family, directionally consistent with scarcity and coexistence.

Framework	**Interaction prediction**	**Marginal-effect geometry and sign reversal**	**Guild prediction**	**Diagnostic interpretation**
H_2 — extinction filtering	$\beta < 0$ (negative effects strongest at high cover).	Negative marginal effect strengthens toward high cover; no sign reversal.	Same direction; specialists substantially more negative at intermediate–high cover.	$\beta < 0$ together with stronger negative marginal effects in specialists at mid–high cover; filtering-direction-consistent, not causal proof. The full signature also requires the absence of a sign reversal, which the cell-level rule does not test (Appendix 5.S3.1).

† In severely deforested systems such as the Atlantic Forest, where most landscapes occur at low to intermediate forest cover, sign reversal may be statistically undetectable regardless of which mechanism operates. Absence of sign reversal should therefore not be interpreted as decisive support for H_1' over H_1, nor as evidence against the broader $\beta > 0$ family.

Table 5.2. Formulation-separated summary of additive and conditional fragmentation results across richness strata. Coefficient medians are marginal over the four fragmentation metrics and are not metric-invariant estimates; the metric-resolved values are shown in Figure 5.3 and Table 5.S5.3. The amount-to-fragmentation ratio is not reported here, because it is not invariant to metric choice and ranges from 2.42 to 8.50 across metrics and formulations in specialist, facultative and restricted-total richness, and from 1.18 to 2.18 in non-forest richness. Direction and support columns are robustness-space accounting across the design and are not effect estimates. Interaction direction is the percentage of negative interaction coefficients among all 648 fitted cells in each stratum and formulation. Primary support is reported as the number of supported cells over the eligible denominator, which is smaller where the matched additive model was unavailable. Marginal effects are shown at the 75th and 95th forest-cover percentiles as raw/residualized values. Forest-specialist marginal effects are medians among the 214 raw and 197 residualized primary-supported cells with bootstrap coverage. Values for restricted-total, forest-facultative and non-forest richness are median point estimates across all eligible cells and do not carry matched bootstrap uncertainty. Raw and residualized coefficients are expressed per one standard deviation of their respective formulation-specific predictor and are therefore comparable in direction and support, but not as common-unit magnitudes or fractions of one another.

Response stratum	**Amount effect βFCOV (median, % CI excl. 0)**	**Fragmentation effect βFRAG (median, % CI excl. 0).**	**Interaction βFCOV×FRAG (% negative, all cells).**	**Primary support (raw / residualized)**	**Marginal effect at p75 / p95 cover**
Forest specialist	+0.322 / +0.350 (99.7 / 100)	−0.062 / −0.060 (40.4 / 53.2)	88.0 / 78.5	214/648 / 197/648	p75 −0.140 / −0.180; p95 −0.275 / −0.290

Response stratum	Amount effect βFCOV (median, % CI excl. 0)	Fragmentation effect βFRAG (median, % CI excl. 0).	Interaction βFCOV×FRAG (% negative, all cells).	Primary support (raw / residualized)	Marginal effect at p75 / p95 cover
Forest facultative	+0.127 / +0.122 (86.7 / 100)	−0.020 / −0.022 (34.3 / 20.7)	68.1 / 71.5	184/648 / 90/645	p75 −0.025 / −0.038; p95 −0.046 / −0.046
Non-forest	−0.071 / −0.069 (23.5 / 9.3)	+0.035 / +0.031 (32.3 / 22.5)	35.8 / 36.1	122/648 / 79/648	p75 +0.070 / +0.065; p95 +0.096 / +0.077
Restricted total	+0.156 / +0.153 (100 / 100)	−0.021 / −0.023 (40.1 / 26.9)	71.3 / 77.6	188/648 / 174/647	p75 −0.035 / −0.042; p95 −0.055 / −0.059

Restricted total richness is the arithmetic sum of the three guilds, so its row is a linked decomposition of the other three rather than an independent response. Non-forest richness is a specificity diagnostic, not a negative control: its amount effect is directionally negative but interval-supported in under a quarter of cells, so the guild contrast is read as evidence about response specificity rather than about the absence of an amount effect, and its support frequency is evaluated against its own precision baseline rather than against the other strata.

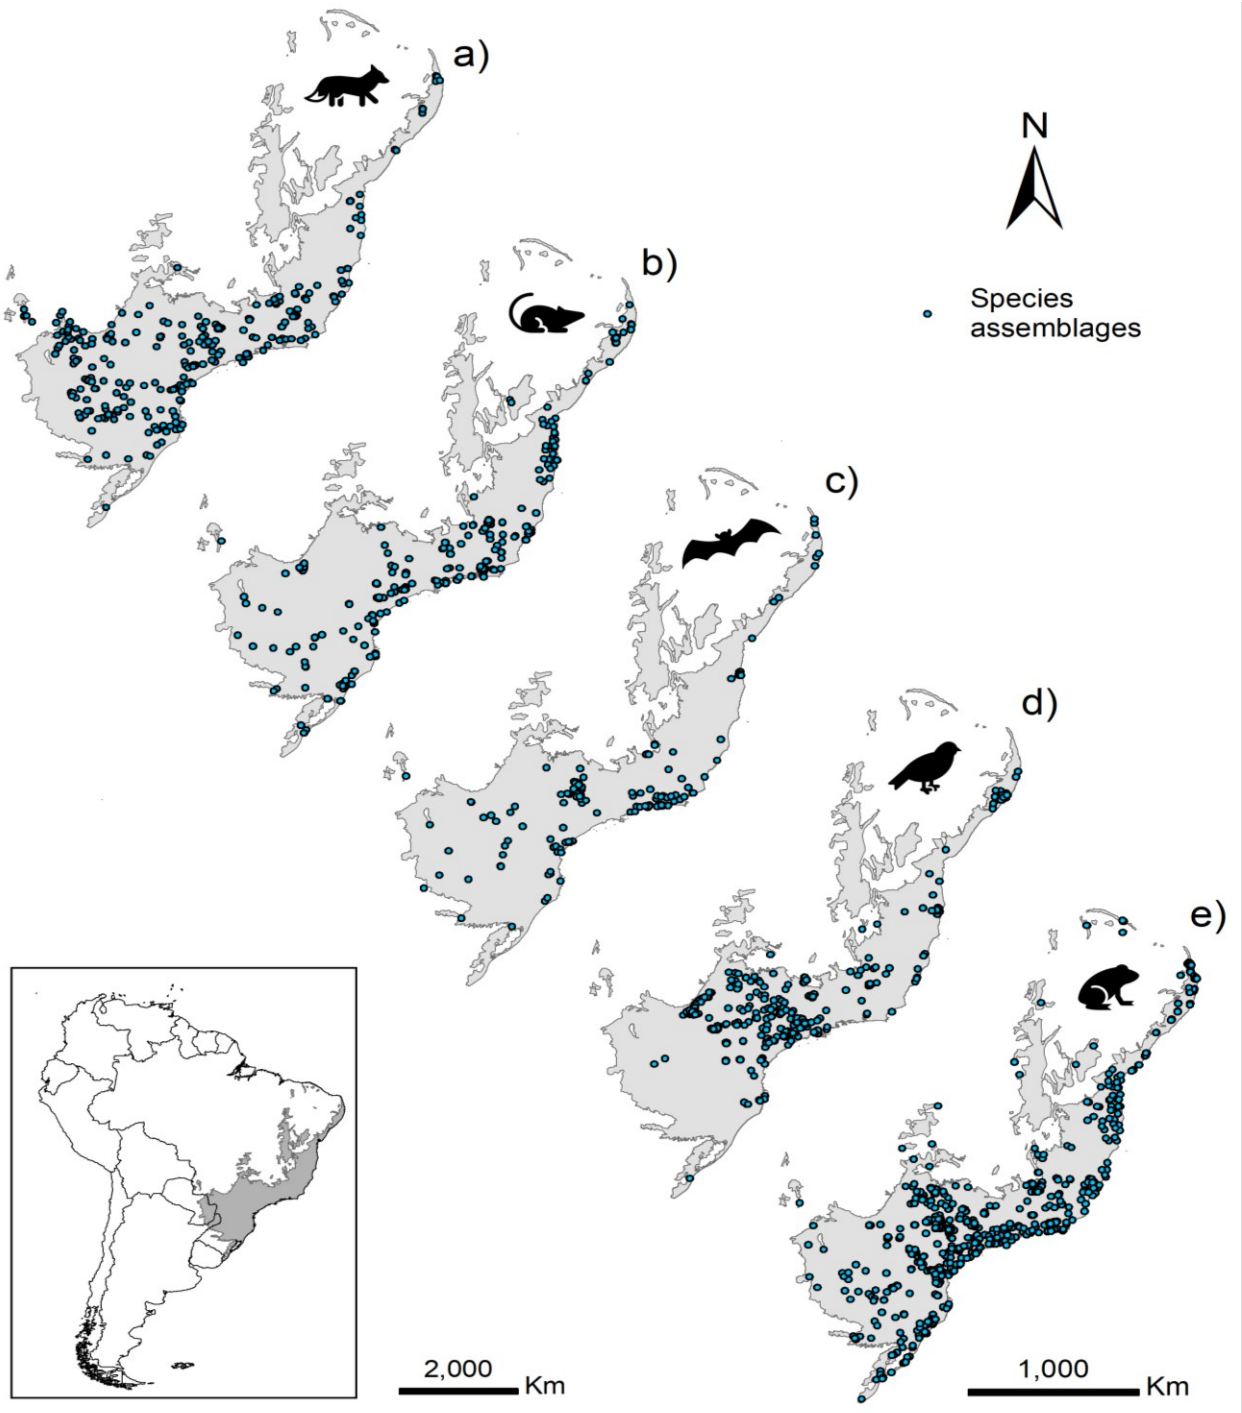


Figure 5.1. Study area and assemblages. .Spatial distribution of curated species assemblages, prior to aggregation, across the Atlantic Forest in South America (shaded) by animal group: mid-large mammals (a), small mammals (b), bats (c), songbirds (d) and anurans (e). Points are source assemblages before consolidation; the models are fitted to aggregated taxon × method × spatial-unit records, so plotted point density is not analytical sampling density.

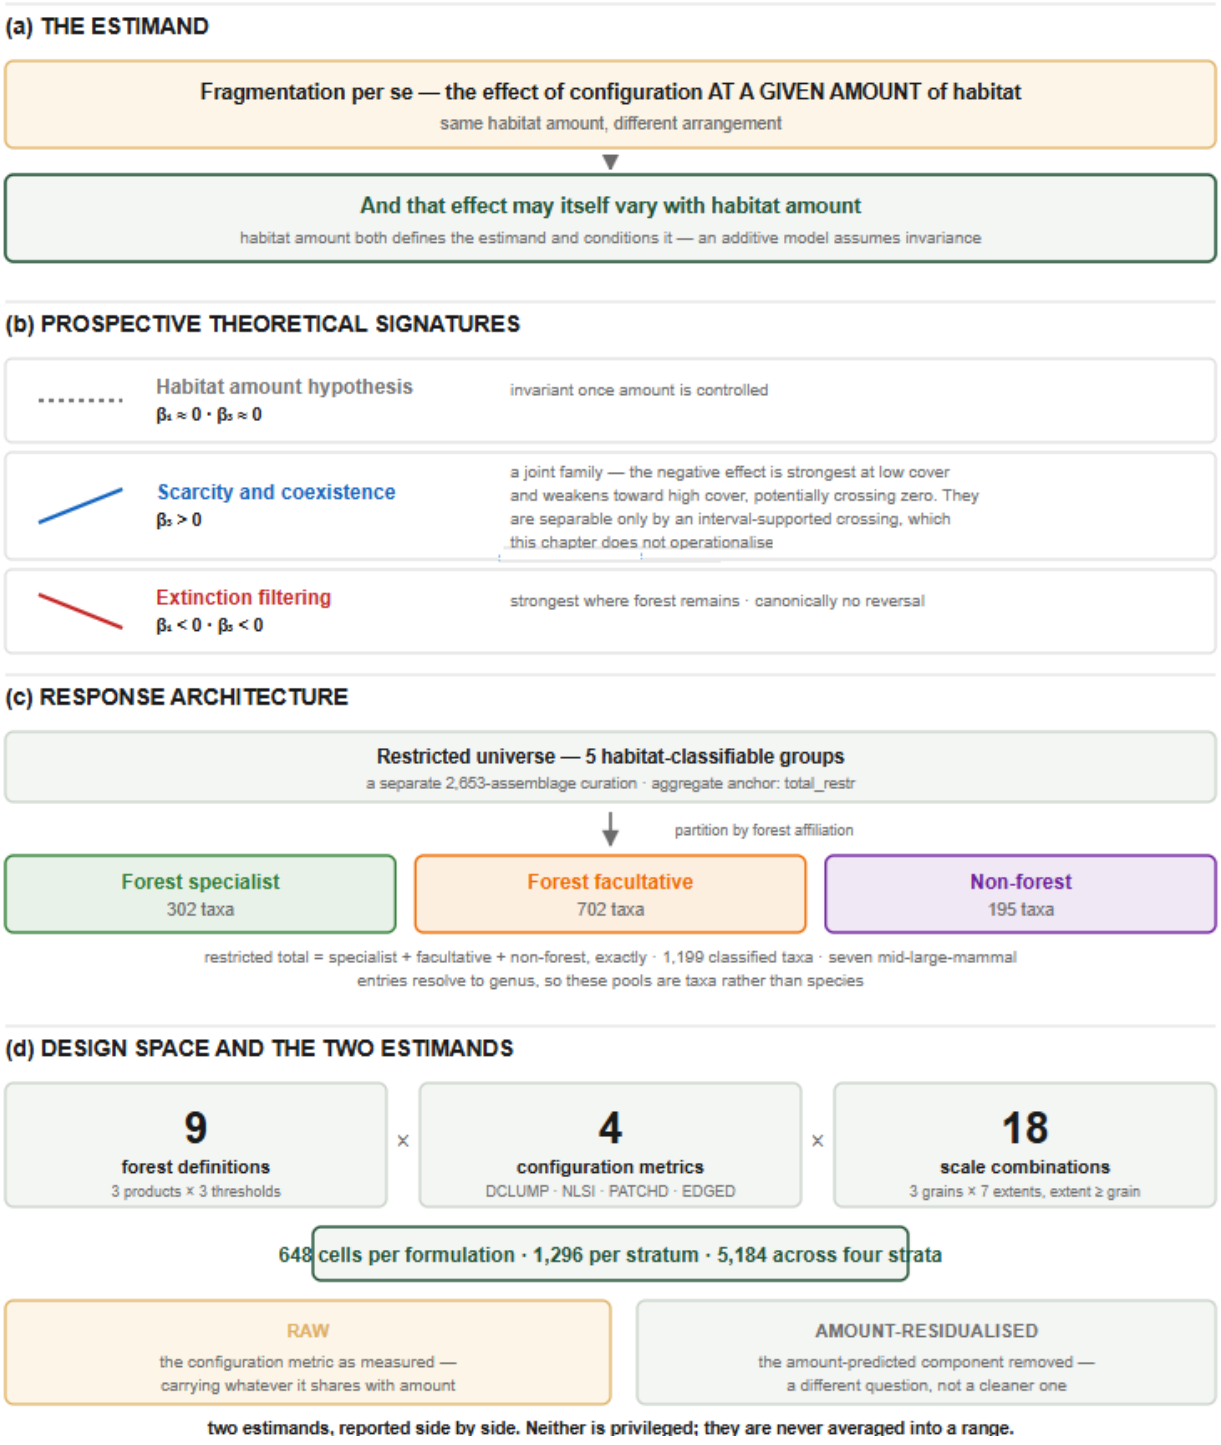


Figure 5.2. The estimand, prospective theoretical signatures, response strata and design space.(a) The estimand. Fragmentation per se is the effect of configuration at a given amount of habitat — the same habitat amount, differently arranged. That effect may itself vary with habitat amount, so habitat amount both defines the estimand and conditions it; an additive amount-plus-

configuration model assumes that the amount-conditioned association is invariant across the gradient. (b) Prospective theoretical signatures. Predicted marginal fragmentation effects across the standardized forest-cover gradient, shown as schematic predictions rather than fitted values and distinguished by line type as well as colour. The flat dotted line represents both the configuration-independent habitat-amount null and the zero reference. Scarcity and coexistence share a positive interaction: the negative fragmentation effect is strongest at low cover and weakens toward high cover, potentially crossing zero. They are separable only by an interval-supported crossing, which this chapter does not operationalize, so they are reported as a joint family. Extinction filtering predicts a negative interaction, under which the negative effect instead strengthens as cover increases. β_1 is the fragmentation effect at the sample mean of standardized forest cover and is therefore a property of the sample as much as of the mechanism. (c) Response strata. The 1,199 classified taxa partitioned by IUCN forest affiliation, drawn from the 2,653-assemblage subset comprising the five habitat-classifiable groups. Seven mid-large-mammal entries are resolved to genus, so the pools are taxa rather than species. Restricted total richness is the arithmetic sum of the three guilds and is a linked decomposition, not an independent replicate. (d) Design space and the two estimands. Nine forest definitions arise from three product-and-resolution streams crossed with three canopy thresholds; four fragmentation metrics; eighteen scale combinations from three response grains and seven predictor extents subject to extent ≥ grain. Scale combinations are nested, overlapping windows on the same assemblages, so repeated recovery across them is scale coherence rather than independent replication. Two formulations, raw and amount-residualized, target non-symmetric estimands and are reported side by side rather than averaged. The design comprises 648 interaction cells per formulation, 1,296 per response stratum and 5,184 across the four strata. Metric and formulation are resolved throughout and never pooled; forest definition and scale are summarized as distributions with stated denominators and resolved cell by cell in Figures 5.S5.3 to 5.S5.5. The AET and TOPOD environmental controls are not part of this factorial; they were applied as a targeted robustness arm to the primary-supported cells, with vegetation texture as a further specialist-only sensitivity.

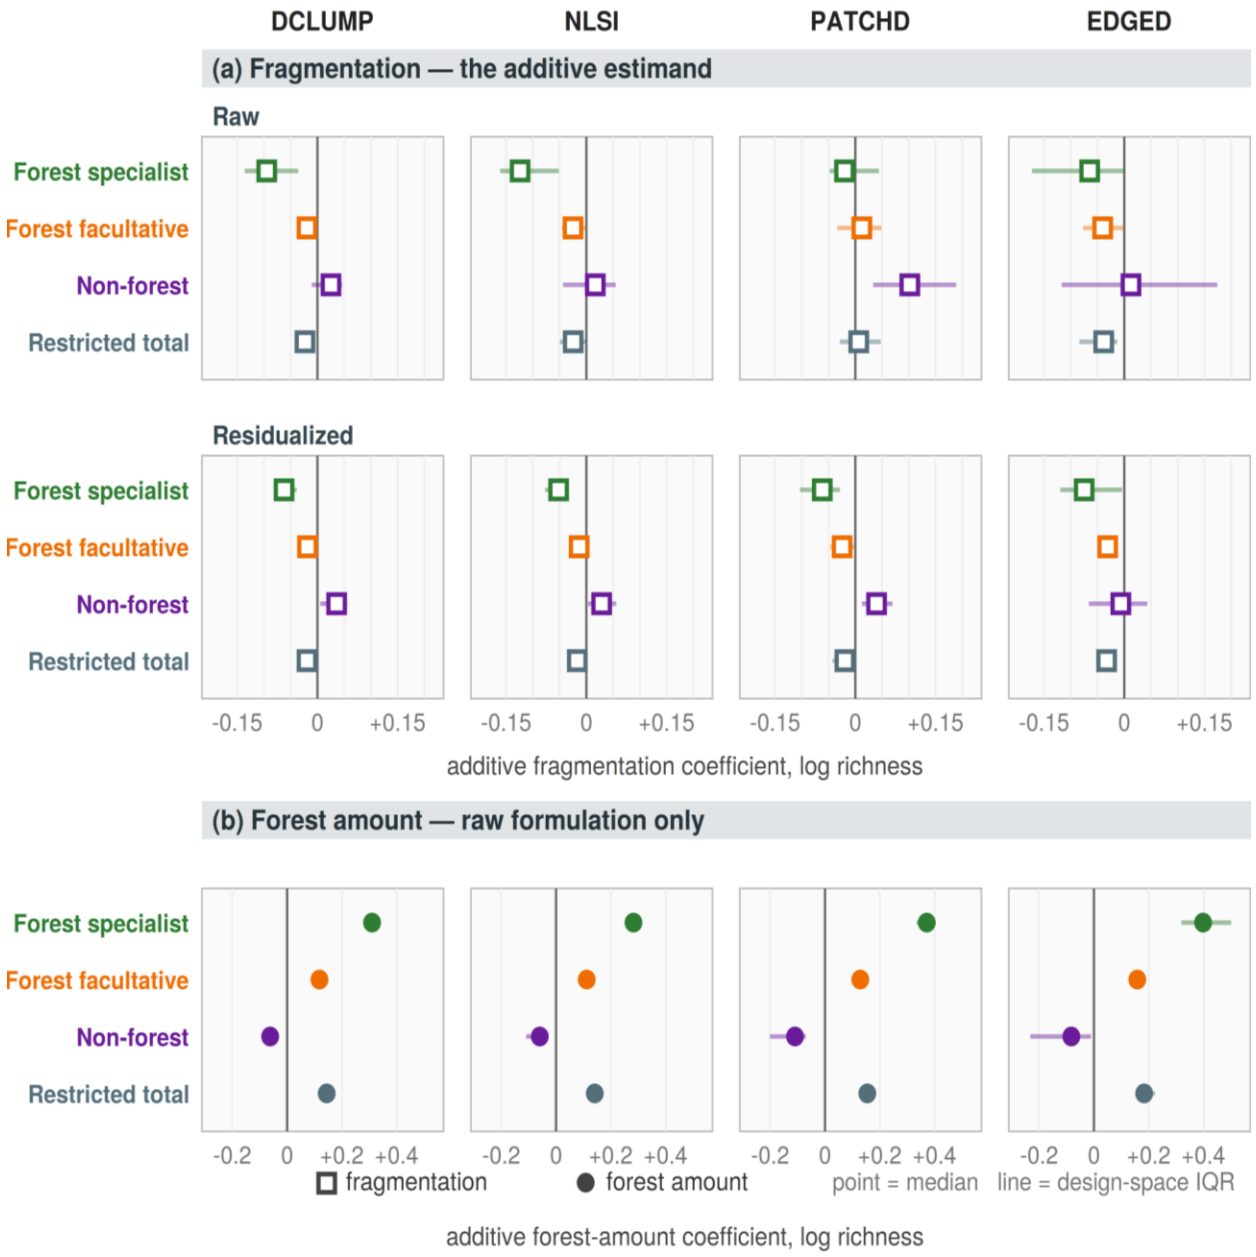


Figure 5.3. Additive fragmentation and forest-amount coefficients by metric and response stratum.

(a) The fragmentation coefficient from the additive family, in which forest amount and fragmentation enter without an interaction term, so the quantity shown is fragmentation per se after amount control — the estimand at issue in the fragmentation debate. Columns are fragmentation metrics and rows are formulations, never pooled, because the coefficient is not invariant to how configuration is measured: for forest specialists in the raw formulation it runs

from −0.020 under PATCHD to −0.124 under NLSI. (b) The forest-amount coefficient, plotted on an independent scale so that the fragmentation estimand is not compressed by the wider amount range. Amount is shown for the raw formulation only, because forest cover enters unresidualized in both formulations and only the metric is residualized; across all sixteen stratum × metric combinations the two formulations differ by a median of 0.012 and at most 0.074, which follows from the changed fragmentation term rather than being a separate result. Points are medians across the 162 cells contributing to each metric, stratum and formulation, obtained by marginalizing over nine forest definitions and eighteen scale combinations; horizontal lines are the interquartile range across those cells and describe design variation, not uncertainty. Unlike the interaction, the additive family has stored standard errors for every cell in all four strata, so interval support is available throughout and is reported in Table 5.2 and Appendix 5.S5.2. Forest amount is positive and fragmentation negative in specialist, facultative and restricted-total richness, and both invert in non-forest richness; in those three strata the amount effect exceeds the fragmentation effect under every metric, by ratios from 2.42 to 8.50, whereas the non-forest ratio runs from 1.18 to 2.18.

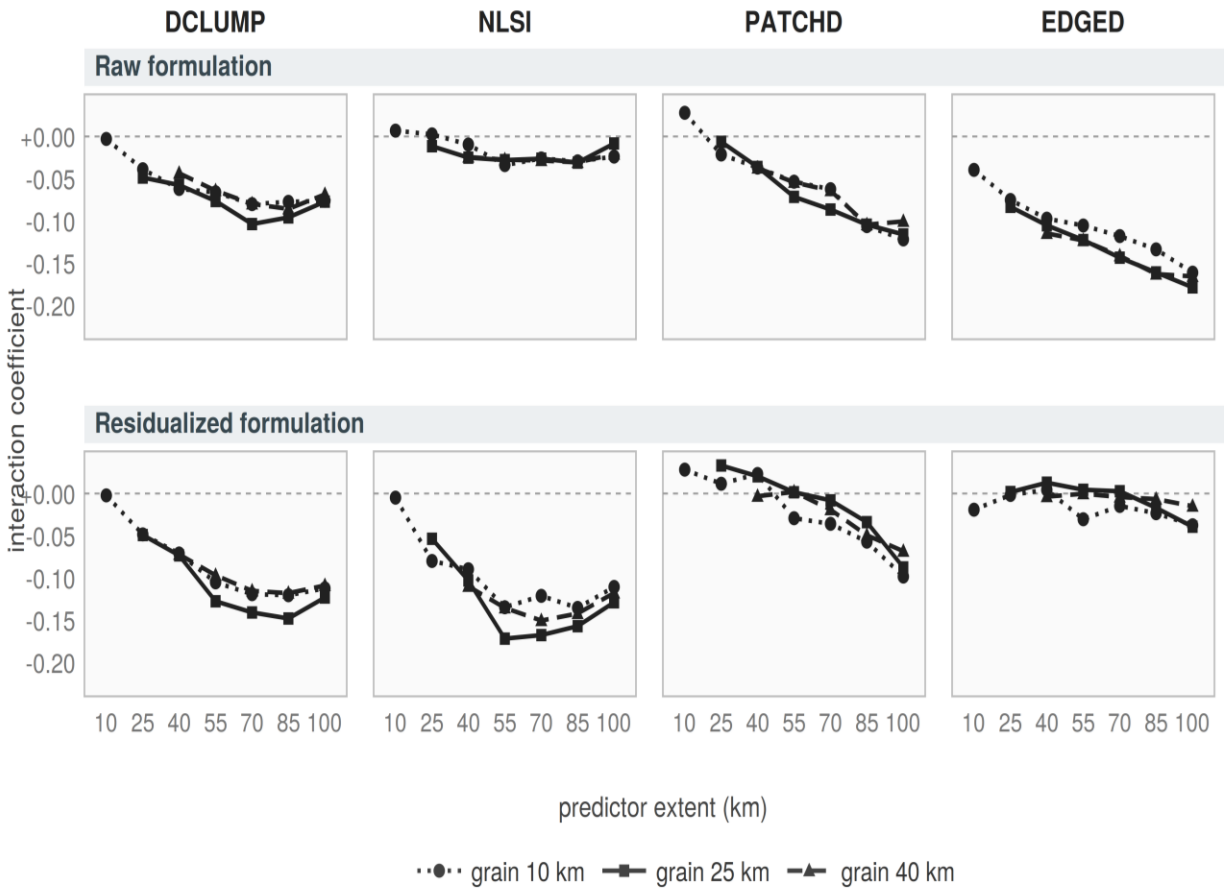


Figure 5.4. Specialist interaction coefficient across predictor extent. Median forest-cover × fragmentation interaction coefficient for forest-specialist richness across predictor extent, with rows for formulation, columns for fragmentation metric, and response grain distinguished by both line type and symbol shape. Each point is a median across the nine forest definitions at that grain, extent, metric and formulation; combinations in which extent is smaller than grain do not exist and are absent. Grain lines nearly coincide within panels while the extent gradient is pronounced, which is the basis for reporting predictor extent rather than response grain as the ordered scale dimension; the extent trajectory itself differs among metrics. The eighteen combinations are nested, overlapping windows on the same assemblages, so the panels are a profile rather than eighteen independent tests. Primary-support frequencies and cell denominators across the same scale grid are given in Figure 5.S5.2.

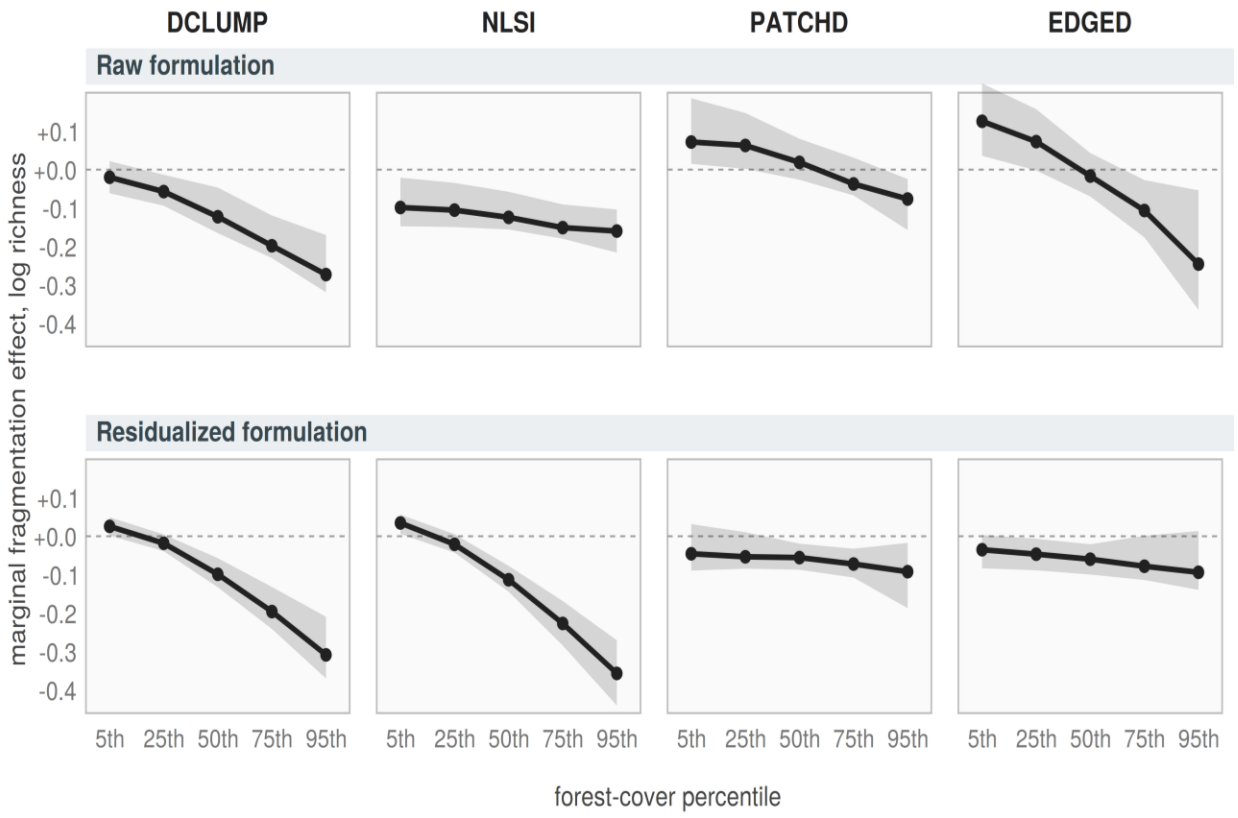


Figure 5.5. Conditional fragmentation effect by metric and formulation. Marginal effect of fragmentation on log richness for forest-specialist richness across five percentiles of the observed forest-cover domain, with rows for formulation and columns for fragmentation metric. Lines are medians across the 162 cells contributing to each metric and formulation; bands are the interquartile range across those cells and are a consistency envelope over a dependent robustness space, not a confidence interval. Metrics are shown separately because the geometry differed jointly by metric and formulation: DCLUMP declined consistently across both formulations, raw NLSI was comparatively flat while residualized NLSI steepened, and PATCHD and edge density showed the strongest formulation-dependent differences at low cover. A single pooled curve would describe none of them. These curves summarise all eligible cells, whereas the percentile estimates reported in the Results are medians among primary-supported cells with bootstrap coverage; the two populations are legitimately different.



Figure 5.6. Interaction coefficient by response stratum and fragmentation metric. Interaction coefficients for forest-specialist, forest-facultative, non-forest and restricted-total richness, plotted as estimated rather than as a derived contrast, with rows for formulation. Strata are distinguished by symbol shape as well as colour, and restricted total richness is shown as a hollow diamond because it is the arithmetic sum of the three guilds rather than a fourth guild. Points are medians across 162 cells; vertical lines are the interquartile range across cells and describe design variation. No model-based confidence interval is shown, because standard errors for the interaction term are unavailable outside forest-specialist richness. The shaded strip beneath each panel gives the median of the per-cell matched differences between specialist and non-forest interactions, which is not the same quantity as the difference between the two plotted medians. Under DCLUMP, NLSI and PATCHD the non-forest interaction is positive while the specialist interaction is negative, giving the guild sign mirror. Under edge density the non-forest interaction is itself negative, so the sign mirror disappears; the raw matched difference nevertheless remains comparable to the other metrics, while the residualized matched difference reverses in sign, and the two partly offset in the metric-level marginal summary.

CHAPTER 6

Putting it all together: coupled ecological and inferential conditionality in biodiversity responses to land-use change

Land-use effects across the empirical results of the thesis varied in sign, magnitude and detectability with thermal exposure, heterogeneity origin, forest amount and configuration, assemblage composition, spatial scale and model estimand. This variation is the dissertation's main ecological and methodological result rather than noise around a single average response. Across the dissertation, biodiversity responses to land-use change are conditional in two linked senses. Ecological conditionality arises because land-use effects vary with climate, heterogeneity origin, productivity, habitat amount and forest configuration. Inferential conditionality arises because coefficients support ecological interpretation only when the model estimates the contrast assigned to them. Assemblage composition and spatial scale cut across both axes because they determine which signals remain in the response and which processes enter the analysed grain and extent. The two axes expose a reciprocal pattern–process problem. Ecological processes generate the patterns available for analysis, while the representation and structure of those patterns determine which ecological contrasts an analysis can recover. Comparable land-use change can therefore generate different biodiversity responses across ecological contexts, and similar fitted patterns can represent different ecological contrasts when response definition, predictor geometry, scale or domain differ.

Ecological conditionality in Atlantic Forest assemblages

In the Atlantic Forest, land-use effects intensified with daily and annual thermal variability and concentrated in forest specialists. Greenness-magnitude variability amplified those effects, whereas configurational greenness diversity did not, persistent geophysical heterogeneity supported baseline richness without moderating conversion, and water–energy availability conditioned part of the vegetational surface. The diagnostic and perspective chapters then showed why fragmentation coefficients must be interpreted through realised predictor geometry and conditional response surfaces. Applying that framework in Chapter 5 revealed negative amount-adjusted configuration associations in forest-specialist richness that became more negative towards higher forest amount. The dissertation therefore shifts the question from whether land-use change reduces richness on average to where effects intensify, which

assemblage components carry them, which scales reveal them and which estimands support ecological interpretation.

Thermal variability and forest dependence

Thermal variability amplified the negative effects of agriculture and broader non-natural cover on richness, most strongly in forest specialists. Daily and annual temperature range together formed a coherent thermal-variability axis. Maximum temperature of the warmest month did not reproduce that pattern, and the dry-season precipitation association did not survive the water-energy and terrain control set.

The interpretation favoured here is that forest canopies reduce thermal variability, most directly at the daily scale to which local canopy buffering applies beneath them, and that land-use conversion weakens or removes this buffering capacity (De Frenne et al. 2019; Ewers and Banks-Leite 2013). Where daily temperature range is already high, conversion can expose organisms to a wider and more acute range of thermal conditions than in climatically buffered landscapes. This mechanism is consistent with the broader expectation that canopy loss, degradation, and fragmentation increase exposure to local climatic variability. It is also consistent with physiological arguments that many tropical species may have limited capacity to tolerate additional short-term thermal stress (Sunday et al. 2014; Khaliq et al. 2014; Senior et al. 2017).

The result should still be interpreted with restraint. Chapter 1 did not directly measure canopy-level microclimate. Daily temperature range was used as an atmospheric moderator, not as a direct measurement of thermal conditions experienced by organisms inside forest fragments, edges, or matrix habitats (Kemppinen et al. 2024). The buffering mechanism is therefore inferred from the interaction between land use and daily temperature range, rather than measured directly. Direct microclimate data from forest interiors, edges, and matrix environments would be needed to test whether canopy-buffering collapse is the operative mechanism rather than a correlated climatic or geographic gradient.

Even with this limitation, the result changes the interpretation of land-use change. Habitat conversion does not only reduce or modify habitat. It also changes the climatic conditions under which remaining organisms persist. In the Atlantic Forest, the biodiversity cost of land-use change is therefore partly a function of thermal exposure. Conservation strategies that consider

land-use intensity or forest cover without considering daily climatic variability may miss landscapes where additional conversion carries especially high ecological costs.

Those climate-conditioned associations were estimated from compounded land-use surrogates that conflate habitat amount, spatial configuration, edge exposure and matrix context (Martínez-Lanfranco 2026a, 2026b). This made the result ecologically important but mechanistically unresolved. Chapter 5 decomposed the land-use side of that compounded signal by asking whether richness losses could be interpreted through forest amount, spatial subdivision, and the amount–configuration response surface.

Land-use effects depended on thermal variability, heterogeneity origin, productivity, habitat amount and forest configuration, but these moderators represent different ecological pathways. Daily and annual temperature ranges index climatic exposure. Greenness-magnitude variability can index degradation, whereas configurational greenness diversity does not reproduce the same surface. Persistent geophysical heterogeneity supports baseline richness without moderating conversion, water-energy availability conditions the greenness-magnitude surface without a demonstrated abundance mediator, and configuration can impose costs on forest specialists. Land-use change therefore acts through multiple pathways whose effects vary with ecological context.

Heterogeneity origin, energy, and response aggregation

Environmental heterogeneity did not behave as a single ecological quantity. Variability in greenness magnitude amplified the negative effects of land use while configurational greenness diversity derived from the same imagery did not reproduce that surface. Persistent geophysical heterogeneity supported baseline richness as an additive template rather than moderating conversion, and water-energy availability conditioned the vegetational surface at fine to mid scales. This contrast is one of the clearest synthesis points of the dissertation.

Classical heterogeneity-diversity logic predicts that spatial variation should increase richness by expanding niche space and allowing species to partition resources or conditions (Hutchinson 1957; MacArthur 1972; Tews et al. 2004; Stein et al. 2014). That expectation is most plausible when heterogeneity reflects persistent environmental structure, such as topographic, geomorphic, or climatic variation. In the Atlantic Forest, persistent geophysical heterogeneity supported baseline richness as an additive template. This pattern is consistent with

the idea that geophysical variation can provide environmental opportunity, refugial conditions, and local climates that allow species to persist under broader land-use pressure (Hjort et al. 2015; Tukiainen et al. 2017; Stein et al. 2014).

Vegetational heterogeneity behaved differently. In the transformed Atlantic Forest landscapes analysed here, vegetational heterogeneity was not simply a proxy for beneficial habitat complexity. It was more consistent with land-cover replacement, edge proliferation, structural degradation, and mosaics of forest, secondary vegetation, agriculture, pasture, and matrix environments (Laurance et al. 2012; Tuanmu and Jetz 2015). Under those conditions, high vegetational heterogeneity can indicate spatial disruption rather than ecological opportunity. Where vegetational heterogeneity is land-use-coupled, it can compound richness loss rather than buffer it.

The productivity and water-energy results add a third pathway, and the three must be kept apart. Variability in greenness magnitude compounded land-use effects in forest specialists. Persistent geophysical heterogeneity supported baseline richness as an additive template and produced no corresponding moderation surface, so it is a background condition rather than a buffer of conversion. Water-energy availability conditioned part of the greenness-magnitude surface and produced guild-specific direct land-use interactions. More productive landscapes can support larger local populations in principle, but the design carries no abundance mediator, so this is energy-conditioned moderation rather than a demonstrated demographic mechanism (Currie 1991; Hawkins et al. 2003; Evans et al. 2005; Hurlbert and Stegen 2014). Richness alone cannot separate resource availability from species-area effects, habitat quality, or covariance with forest amount.

These results refine the landscape-moderation idea by showing that the ecological meaning of heterogeneity depends on its origin. Heterogeneity can promote richness, but only after asking what generated the heterogeneity and what ecological conditions it represents. In a transformed tropical forest biome, vegetational heterogeneity can signal degradation, while abiotic heterogeneity can signal opportunity or refuge. The same abstract term can therefore describe ecological contexts with opposite implications for biodiversity.

Species richness was the empirical throughline of the dissertation because it can be harmonised across taxa, sampling designs, response grains, and predictor extents. Several

conditional signatures nevertheless remained pronounced in this aggregate response. That is notable because aggregate community measures such as species richness are generally considered less sensitive to environmental change than compositional measures, partly because contrasting species responses can compensate one another (Spake et al. 2022). Chapter 2 demonstrates the corresponding aggregation effect through opposing guild-level moderation. Richness still aggregates species with different ecological requirements and retains no information on species identity, abundance or persistence (Fletcher et al. 2025). A landscape can retain similar total richness while losing forest-dependent species and gaining disturbance-tolerant or matrix-associated species (McKinney and Lockwood 1999; Newbold et al. 2018). Total richness can therefore be stable even when assemblage composition shifts towards a more degraded state.

The dissertation addressed this problem through guild partitioning. Chapter 5 separated restricted total richness into forest-specialist, forest-facultative and non-forest components. Forest-specialist interactions were stronger than restricted-total interactions on matched cells, localising the amount-conditioned configuration response to forest specialists. Ecological specialisation and matrix use can mediate species responses to habitat loss and fragmentation, providing a biological basis for examining forest specialists separately from less specialised taxa (Henle et al. 2004; Valente and Betts 2019). Because restricted total richness is the arithmetic sum of those components, the comparison identifies where the aggregate signal originates rather than supplying an independent confirming dataset.

Richness aggregation obscured the compositional consequences of land-use change when guild-level moderation was opposed. Aggregate richness retained shared direct land-use damage but lost much of the moderation signal when guild responses opposed. Chapter 5 reveals a second form of aggregation loss, in which an association concentrated in forest specialists is diluted in restricted total richness. Together, the chapters show that aggregation can obscure response structure through cancellation of opposed effects or dilution of a response concentrated in one guild. The same risk is consistent with habitat-loss analyses in which total richness understates impacts concentrated in habitat specialists (Matthews et al. 2014). Aggregate richness is therefore selectively rather than universally inadequate. It remains informative for assemblage-wide responses but becomes lossy when ecological responses are opposed or strongly concentrated.

Assemblage composition cuts across both axes. The response variable determines what can be detected. Total richness can dilute specialist declines when matrix-tolerant species persist or colonise. Guild-resolved responses can separate a configuration association carried by forest specialists from the aggregate in which it may be diluted. Beta-diversity decomposition reveals whether compositional change reflects ordered loss or replacement. Across the thesis, finer assemblage resolution generally made land-use and fragmentation effects more interpretable.

Inferential conditionality in fragmentation research

The first two Atlantic Forest chapters establish that land-use effects on richness depend on the ecological context in which conversion occurs, and that land-use change is a compound restructuring of habitat, climatic exposure, vegetation texture, resource context, and spatial subdivision rather than a single process. Both chapters, however, work through compounded surrogates that conflate habitat amount, spatial configuration, edge exposure, and matrix transformation into one gradient. They establish where land-use effects intensify without determining which structural component carries them. The dissertation therefore turns to a diagnostic question, whether the quantities estimated by common fragmentation models correspond to the ecological contrasts they are used to represent. That diagnostic work and the response-surface framework it motivates form the inferential core of the thesis, and the Atlantic Forest arc resumes once those standards are specified.

The Atlantic Forest finding of a negative amount-adjusted configuration association in forest-specialist richness appeared difficult to reconcile with global analyses reporting weak or null fragmentation coefficients after additive habitat-amount control. Chapter 3 treated this not as a dispute over which dataset was more representative, but as an inferential problem. The key question was whether additive habitat-amount models estimate the ecological contrast they are used to interpret when habitat amount and configuration are co-generated by habitat loss. Chapter 3 answers it through four conditions of predictor-space adequacy, response-estimand alignment, metric-mechanism correspondence, and domain stability, treating statistical detectability as an intervening diagnostic tier rather than a fifth condition.

Inferential conditionality concerns whether a coefficient estimates the ecological contrast assigned to it. A coefficient estimates a particular quantity under a particular model and predictor distribution rather than automatically representing an ecological effect. Co-generated habitat

amount and configuration can prevent additive coefficients from identifying independent configuration effects, while pooled ecological heterogeneity can make a global average non-diagnostic for any specific mechanism. The dissertation's central methodological contribution is therefore to show that biodiversity inference depends on whether the model's estimand matches the theoretical quantity of interest.

The global diagnostic reanalysis showed that habitat amount and configuration can remain non-separable in realised predictor space. This matters because fragmentation per se is defined as the effect of configuration at a given amount of habitat. If landscapes with the same habitat amount but different configuration are rare or unevenly distributed, then the model may not contain the contrast required for that coefficient to identify an independent configuration effect.

Standard linear diagnostics can miss this problem (Dormann et al. 2013; Koper et al. 2007; Martínez-Lanfranco 2026a). Correlations and variance inflation factors may appear acceptable even when amount and configuration remain coupled nonlinearly or asymmetrically. Configuration can vary widely at intermediate habitat amounts but collapse towards limited states at very low or very high amounts. Under that geometry, the additive partial coefficient for configuration can carry a suppressor signature rather than an independent ecological estimate (Prunier et al. 2015; Smith et al. 2009; Ray-Mukherjee et al. 2014; Martínez-Lanfranco 2026b).

In the global reanalysis, near-zero configuration coefficients shifted towards negative values when residual coupling was reduced through alternative parameterisations, subset analyses, or residualisation sequences. This shift is consistent with geometry-conditioned attenuation of a negative configuration signal rather than a stable ecological null. Under this suppressor attenuation, co-generation of habitat amount and configuration by habitat loss prevents the partial coefficient from estimating an independent configuration effect. A companion analysis of Atlantic Forest tree beta diversity revealed a sign-symmetric failure on the response side. A negative configuration coefficient in total beta diversity did not establish fragmentation per se when the signal did not survive alpha correction and metric-mechanism comparison. The cases are sign-symmetric because either outcome can be non-diagnostic, not because both arise from the same geometric mechanism. The predictor-side case attenuates a negative configuration signal under constrained geometry. The response-side case can amplify a

coefficient whose metric and estimand do not support the inferred mechanism. Both converge on the principle that coefficient sign alone cannot identify fragmentation per se. The chapter therefore identifies the additive null as non-diagnostic unless predictor separability and response-estimand alignment are demonstrated, without estimating a universal negative fragmentation effect.

A partial coefficient is informative only when the realised predictor space supports the contrast being interpreted (Koper et al. 2007; Ruffell et al. 2016; Lundberg et al. 2021; Martínez-Lanfranco 2026a). In habitat-loss-generated landscapes, the biological meaning of an additive configuration coefficient must be demonstrated rather than assumed. Without that demonstration, a near-zero coefficient can reflect the geometry of the sampled landscapes rather than the ecological absence of fragmentation effects.

The domain-stability condition adds a biological limitation to the geometric one. Global multi-taxa datasets pool systems that differ in forest dependence, matrix permeability, land-use history, disturbance chronology, habitat quality, species pools, and the habitat-amount ranges over which configuration effects are expressed (Ewers and Didham 2006; Martínez-Lanfranco 2026b; Henle et al. 2004; Bevan et al. 2025). These differences are ecological, not merely statistical. If they are not measured or modelled, they enter the residual structure and can attenuate the average configuration coefficient. Sensitivity can also vary within species across their geographic ranges, so pooled coefficients can average populations occupying different ecological positions (Orme et al. 2019; Hasui et al. 2024).

This problem is distinct from predictor geometry. Geometry concerns whether the model can estimate the intended contrast from the sampled predictor space. Ecological heterogeneity concerns whether the pooled response is combining systems that should not be forced into a single average coefficient. The remedies are also different. Geometry requires predictor diagnostics, model structures that match the estimand, and designs that sample habitat amount by configuration space deliberately. Ecological heterogeneity requires finer biological information, including forest dependence, matrix quality, landscape history, habitat quality, taxonomic sensitivity, and system-specific gradient positions.

Predictor-geometry failure and biological pooling can operate simultaneously, so a global additive null can be weak for two reasons at once. It may fail to identify the independent

configuration contrast, and it may average over systems whose expected response surfaces differ. The two are distinct problems rather than two instances of one suppressor, and they require different remedies. The null then becomes a model outcome under constrained predictor geometry and unmodelled ecological heterogeneity, not strong evidence that fragmentation is ecologically irrelevant.

The apparent conflict between the Atlantic Forest fragmentation result and the global additive null is not a direct ecological contradiction. It arises because the two analyses assign ecological meaning to coefficients estimated under different predictor geometries. Chapter 5 tested explicit configuration effects within the Atlantic Forest under a design built to reduce amount-configuration confounding and to resolve responses across forest-affiliation guilds. Chapter 3 evaluated whether a broader additive global model could identify the same ecological contrast and found that the sampled predictor geometry did not demonstrate that it could.

This distinction changes the status of additive nulls in two ways. First, a near-zero configuration coefficient may reflect true habitat-amount invariance, but it may also arise from constrained predictor geometry, suppressor attenuation or pooled ecological heterogeneity. Second, even when the intended contrast is recoverable, an additive coefficient averages a configuration association that may vary along the forest-amount gradient. Studies sampling different parts of that gradient can therefore recover different average coefficients from similarly structured conditional surfaces (Duncan and Kefford 2021; Spake et al. 2023). A null result that cannot separate identification failure from gradient-weighted averaging should not be treated as settled evidence that configuration is ecologically irrelevant.

This inferential point has conservation consequences. If model coefficients fail to identify the intended ecological contrast, then using them to deprioritise forest configuration may underprotect sensitive species in systems where configuration effects are real. The detection problem is therefore not only statistical. It affects how evidence is translated into conservation priorities.

The predictor-geometry problem is also broader than fragmentation. It is likely to recur in observational landscape ecology whenever several predictors are co-generated by a shared upstream process. Habitat loss co-generates amount, edge, patch size, isolation, and matrix contrast, and comparable structures are plausible wherever a single driver reshapes multiple

landscape properties at once (Didham et al. 2012; Martínez-Lanfranco 2026a). In such cases, additive coefficients should not be read as independent ecological effects unless the realised predictor space supports the contrast being estimated. Fragmentation is one instance of a broader inferential problem, in which predictor separability must be demonstrated rather than assumed.

From additive effects to conditional response surfaces

The perspective chapter generalised the diagnostic result into a response-surface framework for the fragmentation debate. The central argument is that the debate has often targeted the wrong quantity. Empirical practice has frequently asked whether fragmentation has a single average effect after habitat amount is controlled, whereas the relevant ecological theories predict that configuration effects vary along the habitat-amount gradient.

The habitat amount hypothesis predicts that configuration effects are absent or negligible once habitat amount is accounted for (Fahrig 2013; Saura 2021). Threshold and scarcity dynamics predict that configuration effects intensify after habitat amount falls below critical levels (Andrén 1994; Villard and Metzger 2014). Colonisation–competition coexistence predicts that fragmentation can be beneficial, neutral, or harmful depending on where a system lies along the habitat-amount gradient and which species are being considered (Rybicki et al. 2020; Zhang et al. 2024). Extinction filtering predicts that configuration effects can be strongest where sensitive taxa still persist and fade after those taxa have been lost from heavily transformed landscapes (Betts et al. 2019; Balmford 1996).

These frameworks differ not only in the sign of a mean fragmentation effect but in the geometry of the conditional response surface. The fragmentation by habitat-amount interaction is therefore not a modelling embellishment. It is the minimum linear structure needed to ask whether the fragmentation effect changes along the habitat-amount gradient. An additive model imposes invariance of the configuration effect across habitat amount (Saura 2021; Spake et al. 2023; Fahrig 2003). That restriction is exactly what most competing mechanisms reject.

Additional discrimination comes from response disaggregation and metric-mechanism correspondence. Nestedness-dominated beta diversity supports an ordered-loss interpretation consistent with extinction filtering (Baselga 2010; Legendre 2014). Turnover-dominated beta diversity is more consistent with species replacement, colonisation–competition dynamics, or

other forms of spatial sorting (Baselga 2010; Socolar et al. 2016). Configuration metrics also differ mechanistically. Edge-related metrics speak more directly to exposure and effective area (Fletcher et al. 2016), whereas isolation or connectivity metrics speak more directly to dispersal and recolonisation (Taylor et al. 1993; Bélisle 2005; Fletcher et al. 2016). A model can be statistically valid and still be ecologically non-diagnostic if the metric does not correspond to the mechanism being invoked. A linear interaction is interpretable as a first-order summary only after curvature in the amount response, the fragmentation response, and the bivariate surface has been evaluated, because residual curvature can be absorbed by the interaction term (Pedersen et al. 2019; Tredennick et al. 2021). Gradient coverage sets a further limit, since a conditional effect can be recovered only across the part of the habitat-amount range the sampled landscapes occupy (Duncan and Kefford 2021; Spake et al. 2023). Beta-diversity components are not independent of richness gradients and should be read alongside richness, sampling completeness, and gradient coverage (Legendre 2014).

The response-surface framework makes the fragmentation debate more tractable without declaring one mechanism universally correct. It reframes fragmentation inference around conditional response surfaces, gradient position, response-variable choice and metric-mechanism correspondence, shifting the question from whether fragmentation is positive, negative or null on average to when, for whom and through which mechanism configuration effects are recoverable.

Moving from additive coefficients to conditional response surfaces removes the invariance imposed by an additive model, but it does not supply independent evidence for the ecological meaning of the fitted surface. In a linear interaction model, the marginal configuration curve is a deterministic transformation of the fitted configuration main effect and amount-by-configuration interaction across selected forest-amount values. Pointwise intervals can show where that linear combination is distinguishable from zero, but the curve cannot provide an independent line of evidence for the coefficients or ecological mechanism that generate it. A sign change along the curve may represent a genuine ecological reversal, but it gains evidential weight only when the underlying coefficient pattern survives checks of predictor geometry, metric-mechanism correspondence, formulation, gradient coverage, response specificity and uncertainty. Conditional response surfaces are the appropriate inferential target, but they do not identify the ecological mechanism on their own.

Joint application in the Atlantic Forest

The Atlantic Forest application used a vertebrate fragmentation analysis distinct from the external datasets used in the diagnostic case studies. It is the dissertation's first Atlantic Forest test of explicit forest configuration conducted after the dissertation had specified what fragmentation-per-se inference requires, namely amount control, nonlinear geometry diagnostics, metric architecture validation, habitat-definition sensitivity, spatial-scale robustness and response stratification across forest-affiliation guilds. After the modelled dependence on forest amount was removed, the fitted configuration association in forest-specialist richness was negative, and the negative effect strengthened towards higher forest amount. The fitted association is neither the outcome of rearranging a fixed amount of forest (Koper et al. 2007; Lundberg et al. 2021; Martínez-Lanfranco 2026a) nor the total effect of forest loss transmitted through changes in configuration (Ruffell et al. 2016; Püttker et al. 2020; Martínez-Lanfranco 2026b). The pattern was evaluated across forest products, canopy thresholds, metrics, formulations, spatial scales and response strata rather than inferred from one coefficient. Within the completed targeted control arm, supported interactions retained their direction after water-energy and terrain adjustment, showing directional robustness to those measured alternatives without establishing independence from unmeasured environmental structure.

The fitted marginal associations exposed this limitation empirically. Point-estimate marginal associations crossed zero within the observed forest-cover domain in about half of the evaluated specialist cells under each formulation, with crossings distributed across metrics. Interval-supported positive low-cover associations formed a much narrower subset, concentrated chiefly in raw Hansen edge-density and residualised Hansen patch-density cells and absent among the primary-supported Tuanmu clumpiness cells. This difference reflects variation in absolute marginal-effect magnitude and bootstrap uncertainty rather than merely whether the fitted line crossed zero. Residualisation also changed the estimand and standardised magnitude differently among metrics, so redistribution of interval support across formulations does not independently refute the raw pattern. What weakens their interpretation as a general ecological benefit is the limited convergence of interval-supported positive associations across forest-product, metric and formulation combinations. The dominant conclusion remains the increasingly negative specialist association towards higher forest amount.

The amount-conditioned specialist surface is most consistent with extinction filtering because configuration associations remain strongest where sufficient habitat persists to retain sensitive taxa (Püttker et al. 2020; Betts et al. 2019). The chapter does not observe that historical filtering directly. The specialist–non-forest sign mirror adds guild specificity but is metric-dependent, appearing for three of the four configuration metrics and not for edge density. Composition-dependent sampling remains the strongest live threat to interpreting that mirror as ecological filtering, because finite effort, edge- or road-biased placement, and protocol differences could change which species are recorded and therefore which guild components enter the observed richness response, without equivalent demographic change (Sobral-Souza et al. 2021a, 2024; Rodrigues et al. 2020). Lower specialist richness towards the low-cover end of the gradient may also reduce the recoverability of configuration associations there. The pointwise-supported positive-to-negative endpoint pattern departs from the prospective filtering prediction, which expects negative effects to strengthen towards higher cover without sign reversal, further limiting a uniquely filtering interpretation. Extinction filtering is therefore the best-supported mechanism family, but it is not uniquely identified.

The connection between Chapters 3 and 5 also suggests that contemporary fragmentation effects may be conservative in some assemblages. If heavily transformed landscapes have already lost the species most sensitive to configuration, then present-day richness gradients may underestimate historical fragmentation costs. The stronger conclusion is that contemporary assemblages need not be neutral records of fragmentation effects. If prior filtering has removed configuration-sensitive species, contemporary assemblages would provide a survivor-conditioned response surface. Assemblage history can therefore constrain the response surface available to a cross-sectional analysis even when the contemporary model is correctly specified. This links assemblage conditionality directly to inferential conditionality. If prior loss has removed the most configuration-sensitive species from heavily transformed landscapes, a contemporary richness model estimates the configuration-richness relationship among surviving assemblages rather than the full historical response surface. Better model specification can reduce geometric confounding, but it cannot recover species that are no longer present in the response variable. In that sense, assemblage history determines what any cross-sectional model can estimate, and the limitation is biological, not statistical.

What ecological and inferential conditionality imply

The two forms of conditionality meet rather than run in parallel. Ecological context shapes which species persist and therefore which responses remain available to measure, while predictor geometry, response choice, model structure, and spatial scale determine which part of that response an observational analysis can recover. Assemblage composition and scale operate across both. Understanding land-use effects therefore requires measuring the right moderators, disaggregating the right response components, fitting models that target the intended contrast, and evaluating relationships at ecologically meaningful scales.

Spatial scale likewise cuts across both axes. In this dissertation, scale was both a design principle and, where effects varied across response grain or predictor extent, an empirical signal. The grain at which richness was aggregated and the extent at which predictors were quantified helped determine which ecological processes could enter the analysis and which contrasts were recoverable. Fine grains may capture local habitat structure, canopy buffering, and edge exposure. Broader extents may capture dispersal limitation, source-sink dynamics, productivity gradients, and regional environmental constraints. A weak effect at one grain or extent does not necessarily indicate ecological absence. It may indicate that the analysis is being conducted at a scale that does not match the process.

Shared inferential boundaries

This dissertation establishes that land-use effects on species richness in the Atlantic Forest were generally negative across the analyses and context-dependent. Daily and annual temperature ranges amplified those effects, most strongly in forest specialists, which is more consistent with loss of canopy buffering than with a general effect of peak heat. Heterogeneity origin mattered, because variability in greenness magnitude amplified land-use damage whereas persistent geophysical heterogeneity supported baseline richness as an additive template. Additive habitat-amount models in global fragmentation analyses can produce near-zero configuration coefficients that are non-diagnostic. Amount-adjusted configuration associations were negative in forest-specialist richness and became more negative towards higher forest amount. Conditional response surfaces provide a more informative inferential target than average additive fragmentation coefficients.

The three Atlantic Forest chapters form one coordinated evidential system rather than independent replications, sharing assemblage evidence, forest-affiliation architecture, scale matrices and modelling framework. Across them, land-cover composition is represented by cropland prevalence and non-natural cover in Chapters 1 and 2 and by forest cover in Chapter 5, while land-cover configuration is analysed explicitly in Chapter 5. Recurring specialist responses therefore triangulate across processes rather than confirming one result three times. Nor are the chapters a sequential variable-selection or cumulative-adjustment exercise: each tests a different ecological contrast and carrying a supported moderation relationship forward would require asking whether it also conditions the next focal relationship, thereby defining a new, potentially higher-order question rather than a neutral adjustment. Table 6.1 summarises how the four conditions defined in Chapter 3 play out across them and what inference each retains.

Predictor-space adequacy is met differently in each chapter, and the differences are in kind. In the climate chapter the coupling sits within the land-use surrogate, because cropland prevalence and non-natural cover each aggregate canopy removal, habitat amount, matrix composition, edge exposure and degradation, and no contrast in the data separates them. In the heterogeneity chapter the coupling sits within the realised joint predictor space of the land-use gradient and the moderator, because the vegetation state that indexes moderation is generated in part by that gradient (Martínez-Lanfranco 2026b). In the fragmentation chapter it sits between two predictors that one upstream process generates together, which is why removing the modelled dependence on forest amount defines an operational amount-adjusted component. The shared response is to compare alternative representations, condition on measured structure, and report divergence instead of averaging it. None of that substitutes for the predictor contrast a causal claim would require (Martínez-Lanfranco 2026a, 2026b), so each fitted interaction identifies a conditional association within the predictor space the sampled landscapes actually contain.

Response-estimand alignment turns a measurement limitation into an analytical instrument. Richness is the common currency that makes the three chapters comparable. Chapters 1 and 2 analyse both the full universe and the habitat-classifiable universe, whereas Chapter 5 is restricted to the habitat-classifiable universe because forest affiliation is required throughout. Within that universe, partitioning restricted total richness into forest-specialist, forest-facultative and non-forest components changes the inference rather than adding detail.

What the partition exposes differs by chapter. Aggregation weakens a climate-conditioned signal in one case, cancels oppositely signed guild responses in another, and dilutes the specialist configuration association in the third. Because restricted total richness is the arithmetic sum of its three components, that third comparison localises which component carries the association and does not supply an independent confirming dataset. Abundance, occupancy, turnover and demographic rates lie outside the response, and the guilds rest on range-wide habitat classifications that need not describe habitat use within this biome. Richness is a net state variable, so small net changes can conceal substantial species gains and losses (Hillebrand et al. 2018). In these cross-sectional analyses, a stable or positive guild association cannot be interpreted as evidence of demographic compensation. Cross-taxon coefficients therefore describe average responses across the sampled groups, and fine-grained taxon-specific mechanisms remain a separate target. The three chapters also share one guild construction, so the specialist concentration recurring across them does not independently validate the guild architecture. A deeper problem is biological rather than statistical. Where earlier transformation has already removed the most sensitive species from the most altered landscapes, a contemporary richness model estimates the response surface of surviving assemblages, so part of the effect of interest is absent from the response variable rather than mismeasured within it (Betts et al. 2019; Balmford 1996; Lawes et al. 2007).

Metric-mechanism correspondence does more work through divergence than through agreement. Where two representations of one broad quantity agree, they identify a conclusion that survives a change of product, metric or formulation. Where they disagree, they narrow the construct to which the conclusion applies. Daily and annual temperature range separate from maximum temperature, and greenness-magnitude variability separates from configurational greenness diversity. Chapter 5 shows two pattern-specific metric limits. Residualised patch density departs from the broader cross-metric configuration pattern, whereas edge density does not reproduce the specialist–non-forest sign mirror. Each narrowing has the same form, because a broad label covered quantities that the analysis then showed to behave differently. Climate, heterogeneity and fragmentation must therefore retain the biological and statistical meaning of the variables that operationalise them, and a result stated at the level of the label claims more than any chapter established.

Domain stability is where the analytical and the geographic boundaries meet. The grain and extent matrices are nested and overlapping views of many of the same assemblages, so recurrence supports robustness across analytical scales while variation across scales marks where an association changes without identifying whether ecology, spatial support or predictor averaging produces that change (Miguet et al. 2016; Jackson and Fahrig 2015; Fletcher et al. 2023). Neither counts as independent replication, and no chapter identifies a unique ecological scale of process. In Chapter 3, independent reconstruction from a different forest representation reproduced the direction of the total-beta patch-number association and approximately three quarters of its magnitude, showing that the fitted signal was not a geospatial reconstruction artefact. The association nevertheless remained narrow and scale-sensitive across the 25 to 400 km^2 extents, while non-collinear mechanism-matched metrics remained unsupported at every extent. Representation robustness therefore strengthened confidence in the fitted association without identifying the proposed fragmentation mechanism. Habitat definition works the same way at a finer level, because forest is an operational category throughout, and the sensitivity of magnitude and detectability to product, canopy threshold and resolution shows that conclusions expressed through forest amount and forest specialists carry the definition that produced them (Almeida-Gomes et al. 2016; Dennis and Huck 2025; Betts et al. 2014; Hadley and Betts 2016). Because species can differ in what constitutes habitat, a common mapped forest class can create false equivalence across taxa (Fletcher et al. 2025). Varying forest products and canopy thresholds does not eliminate the mismatch between mapped forest and species-specific habitat. The Atlantic Forest itself defines the outer boundary, since its long transformation history may affect the strength and detectability of these patterns and limits direct transfer to less altered systems.

Temporal alignment cuts across all four conditions in the way that assemblage composition and spatial scale cut across the two axes. The empirical chapters substitute space for time throughout: contemporary spatial contrasts stand in for processes that unfold over time, an approach standard in large observational syntheses but one that cannot reconstruct the histories that produced the landscapes compared, nor the dispersal trajectories, time lags and unpaid extinction debts those histories imply. Biodiversity surveys span several decades, whereas climate, land-use, vegetation and forest products represent different historical windows, so temporal mismatch can alter the magnitude and interpretation of estimated associations (De

Palma et al. 2018; Purvis et al. 2018). Its consequences also differ among chapters. In the climate and heterogeneity chapters the mismatch obscures which historical context corresponds most closely to each assemblage. In the fragmentation chapter it bears on the estimand itself, because configuration can be rearranged at near-constant cover, so stability in regional forest extent does not establish that landscape structure was long-standing. Longitudinal assemblage data and historical landscape reconstruction would help distinguish a contemporary spatial snapshot from processes defined by landscape change through time (Fletcher et al. 2026), including delayed loss, recolonisation and extinction debt (Kuussaari et al. 2009; Daskalova et al. 2020; Semper-Pascual et al. 2021).

A common statistical architecture supports these comparisons and bounds them. Mixed models, observation-level random effects, hierarchical taxon and method terms, residual spatial adjustment, information-theoretic comparisons and structure-coefficient diagnostics make chapter contrasts directly comparable, and model-based estimates carry the primary inference throughout. Multi-source biodiversity data retain residual variation in sampling design, effort and taxonomic coverage which the modelling architecture was designed to accommodate without establishing that all such heterogeneity was removed, and the framework cannot externally validate its own assumptions (Purvis et al. 2018). Independence is limited beyond the dissertation as well, because the assemblage evidence derives from compiled data series that other published analyses of this biome also use, so agreement with a published result drawn from the same compilation reflects a shared data lineage. The fragmentation chapter is likewise an internal application of the framework the perspective chapter developed, testing whether that framework yields discriminating and empirically supported predictions in this system. External validation requires independent systems or datasets, while independent analytical teams would provide an additional reproducibility test.

Future directions

Direct measurements are needed for the two mechanisms that the dissertation infers rather than observes. The canopy-buffering interpretation rests on atmospheric thermal variability rather than the microclimate organisms experience, so distributed temperature and humidity sensors across forest interiors, edges and matrix habitats, crossed with canopy structure from LiDAR or comparable products, would test that mediation directly. The extinction-filtering

interpretation and the temporal boundary require the same kind of evidence in time rather than space. Long-term resurveys, occupancy time series and historical landscape reconstruction should show whether landscapes with attenuated configuration associations are still losing sensitive species, have already passed through a filter, or are responding to restoration.

Extending the response beyond richness would address a limitation shared by every empirical chapter. Abundance data would reveal whether richness losses are preceded by population decline. Compositional turnover would separate ordered loss from replacement. Trait and demographic data would identify which attributes predict sensitivity and whether local extinction, reduced colonisation, recruitment limitation or lowered survival generates the observed pattern. Direct information on species co-occurrence and biotic interactions would likewise be needed to test whether competition, facilitation or trophic relationships mediate the assemblage-level responses observed here; richness data alone cannot identify those pathways. Together, these dimensions would convert conditional richness patterns into more direct mechanistic tests.

Prospective fragmentation studies should identify the ecological contrast before model fitting and sample landscapes spanning the habitat amount by configuration space needed to estimate it, rather than relying only on post hoc diagnosis. Complete predictor independence may be unavailable in real landscapes, but the sampled domain must contain the contrast being interpreted. Replication in systems with different transformation histories, using independent datasets and analytical teams, would then test whether the conditional structures recovered here extend beyond the Atlantic Forest. Where landscape change can be followed prospectively, a before–after control–impact or matched quasi-experimental design would provide a stronger test: repeated assemblage surveys before and after forest loss, subdivision or restoration, paired with comparable landscapes that do not undergo the same change. Such temporal contrasts would not make habitat amount and configuration independent, but they would distinguish responses to landscape change from pre-existing differences among landscapes and provide a more direct test of the conditional response surfaces proposed here.

Conservation implications

The dissertation supports conditional conservation inference in transformed tropical landscapes rather than a single threshold or universal management prescription.

Thermal exposure helps identify where additional conversion will impose the greatest biodiversity cost, which forest amount alone does not. In the Atlantic Forest, land-use effects on richness became more negative where daily temperature range was higher. This suggests that conservation prioritisation should consider not only how much forest remains, but also where conversion exposes assemblages to strong thermal variability across daily and annual ranges. Maintaining canopy structure, reducing edge exposure, and restoring forest cover may be especially important in thermally variable landscapes.

Conservation assessments should not treat all heterogeneity as beneficial. In the transformed Atlantic Forest landscapes analysed here, land-use-responsive vegetational heterogeneity was more consistent with degradation, edge proliferation, and mosaic reorganisation than with ecological opportunity. By contrast, persistent geophysical heterogeneity supported baseline richness as an additive template, while water-energy availability conditioned the greenness-magnitude surface and produced guild-specific direct interactions without identifying an abundance-mediated mechanism. Conservation planning should therefore distinguish persistent geophysical heterogeneity, which can support refugial conditions, from vegetation texture generated by land-use transformation.

Forest amount remains fundamental, yet large remnant area alone does not guarantee faunal or functional integrity (Jorge et al. 2013; Magioli et al. 2021). Chapter 5 showed that amount-adjusted configuration associations in forest-specialist richness became more negative towards higher forest amount, although reduced recoverability towards the low-cover end of the gradient could contribute to the same surface. Higher forest amount here refers to the upper part of each design cell's forest-cover distribution, not to one absolute threshold. Applying the fitted pattern to a particular landscape therefore requires its absolute forest-cover distribution together with the distribution of the relevant configuration metric. Configuration should not be ignored where sensitive assemblages persist, but habitat amount remains the primary conservation quantity, and the result does not justify substituting configuration management for forest retention.

Monitoring based only on total richness can miss guild-specific or compositional change when tolerant and forest-associated taxa respond in opposing directions. Aggregate richness is selectively rather than universally inadequate, so it reports shared direct damage well and loses

the moderation that identifies where damage intensifies. Conservation assessments should therefore incorporate forest-affiliation guilds and, where the question concerns compositional change, beta-diversity decomposition. Broad landscape indicators can guide conservation, but their interpretation depends on how closely they represent the response and mechanism under study. Poor alignment calls for more specific response variables and landscape metrics.

Conclusions

Across the Atlantic Forest analyses, land-use effects depended on thermal variability, heterogeneity origin, productivity context, assemblage composition and forest configuration. Because the first two chapters used compounded land-use surrogates, they identified where effects intensified without isolating the structural pathway.

The fragmentation sequence resolved that limitation recursively. Additive nulls became non-diagnostic when habitat amount and configuration remained coupled in realised predictor space. Conditional response surfaces replaced average effects as the target, yet the fitted surface still required checks of geometry, metric correspondence, response specificity, domain and uncertainty. In the Atlantic Forest application, amount-adjusted configuration associations in forest-specialist richness became more negative towards higher forest amount, while forest amount remained the dominant landscape quantity and positive low-cover intervals converged only in a limited subset of forest-product, metric and formulation combinations.

Land-use effects on biodiversity therefore cannot be reduced to average coefficients, aggregate richness or one spatial scale. Ecological context determines where effects intensify, while response definition, predictor geometry and model estimand determine what those effects can mean. The dissertation's contribution is to show where those two forms of conditionality meet and how to test their consequences without treating model coefficients as ecological facts.

References

Almeida-Gomes M, Prevedello JA, Crouzeilles R (2016) The use of native vegetation as a proxy for habitat may overestimate habitat availability in fragmented landscapes. Landscape Ecology 31:711-719.

Andrén H (1994) Effects of habitat fragmentation on birds and mammals in landscapes with different proportions of suitable habitat: a review. Oikos 71:355-366.

Balmford A (1996) Extinction filters and current resilience: the significance of past selection pressures for conservation biology. Trends in Ecology and Evolution 11:193-196.

Baselga A (2010) Partitioning the turnover and nestedness components of beta diversity. Global Ecology and Biogeography 19:134-143.

Bélisle M (2005) Measuring landscape connectivity: the challenge of behavioral landscape ecology. Ecology 86:1988-1995.

Betts MG, Fahrig L, Hadley AS, et al. (2014) A species-centered approach for uncovering generalities in organism responses to habitat loss and fragmentation. Ecography 37:517-527.

Betts MG, Wolf C, Pfeifer M, et al. (2019) Extinction filters mediate the global effects of habitat fragmentation on animals. Science 366:1236-1239.

Bevan PA, Braga Ferreira G, Ingram DJ, et al. (2025) Regional Biomes outperform broader spatial units in capturing biodiversity responses to land-use change. Ecography 2025:e07318.

Currie DJ (1991) Energy and large-scale patterns of animal- and plant-species richness. The American Naturalist 137:27-49.

Daskalova GN, Myers-Smith IH, Bjorkman AD, et al. (2020) Landscape-scale forest loss as a catalyst of population and biodiversity change. Science 368(6497):1341-1347.

De Frenne P, Zellweger F, Rodríguez-Sánchez F, et al. (2019) Global buffering of temperatures under forest canopies. Nature Ecology and Evolution 3:744-749.

Dennis M, Huck JJ (2025) The continuity-contiguity problem in fragmentation-biodiversity research. Journal of Biogeography 52:e15077.

De Palma A, Sanchez-Ortiz K, Martin PA, et al. (2018) Challenges with inferring how land-use affects terrestrial biodiversity: study design, time, space and synthesis. Advances in Ecological Research 58:163-199.

Didham RK, Kapos V, Ewers RM (2012) Rethinking the conceptual foundations of habitat fragmentation research. Oikos 121:161-170.

Dormann CF, Elith J, Bacher S, et al. (2013) Collinearity: a review of methods to deal with it and a simulation study evaluating their performance. Ecography 36:27-46.

Duncan RP, Kefford BJ (2021) Interactions in statistical models: three things to know. Methods in Ecology and Evolution 12:2287-2297.

Evans KL, Warren PH, Gaston KJ (2005) Species-energy relationships at the macroecological scale: a review of the mechanisms. Biological Reviews 80:1-25.

Ewers RM, Banks-Leite C (2013) Fragmentation impairs the microclimate buffering effect of tropical forests. PLoS ONE 8:e58093.

Ewers RM, Didham RK (2006) Confounding factors in the detection of species responses to habitat fragmentation. Biological Reviews 81:117-142.

Fahrig L (2003) Effects of habitat fragmentation on biodiversity. Annual Review of Ecology, Evolution and Systematics 34:487-515.

Fahrig L (2013) Rethinking patch size and isolation effects: the habitat amount hypothesis. Journal of Biogeography 40:1649-1663.

Fletcher RJ, Betts MG, Damschen EI, et al. (2023) Addressing the problem of scale that emerges with habitat fragmentation. Global Ecology and Biogeography 32:828-841.

Fletcher RJ, Bullock JM, Graham LJ, et al. (2026) Resurrecting habitat fragmentation as a process over time. Trends in Ecology and Evolution, in press.

Fletcher RJ, Burrell NS, Reichert BE, et al. (2016) Divergent perspectives on landscape connectivity reveal consistent effects from genes to communities. Current Landscape Ecology Reports 1:67-79.

Fletcher RJ, Green RE, Bladon EK, et al. (2025) Beyond species richness for biological conservation. Conservation Letters 18:e13124.

Hadley AS, Betts MG (2016) Refocusing habitat fragmentation research using lessons from the last decade. Current Landscape Ecology Reports 1:55-66.

Hasui E, Martensen AC, Uezu A, et al. (2024) Populations across bird species distribution ranges respond differently to habitat loss and fragmentation: implications for conservation strategies. Perspectives in Ecology and Conservation 22:43-54.

Hawkins BA, Field R, Cornell HV, et al. (2003) Energy, water, and broad-scale geographic patterns of species richness. Ecology 84:3105-3117.

Henle K, Davies KF, Kleyer M, et al. (2004) Predictors of species sensitivity to fragmentation. Biodiversity and Conservation 13:207-251.

Hillebrand H, Blasius B, Borer ET, et al. (2018) Biodiversity change is uncoupled from species richness trends: consequences for conservation and monitoring. Journal of Applied Ecology 55:169-184.

Hjort J, Gordon JE, Gray M, et al. (2015) Why geodiversity matters in valuing nature's stage. Conservation Biology 29:630-639.

Hurlbert AH, Stegen JC (2014) When should species richness be energy limited, and how would we know? Ecology Letters 17:401-413.

Hutchinson GE (1957) Concluding remarks. Cold Spring Harbor Symposia on Quantitative Biology 22:415-427.

Jackson HB, Fahrig L (2015) Are ecologists conducting research at the optimal scale? Global Ecology and Biogeography 24:52-63.

Jorge MLSP, Galetti M, Ribeiro MC, et al. (2013) Mammal defaunation as surrogate of trophic cascades in a biodiversity hotspot. Biological Conservation 163:49-57.

Kemppinen J, Lembrechts JJ, Van Meerbeek K, et al. (2024) Microclimate, an important part of ecology and biogeography. Global Ecology and Biogeography 33:e13834

Khaliq I, Hof C, Prinzinger R, et al. (2014) Global variation in thermal tolerances and vulnerability of endotherms to climate change. Proceedings of the Royal Society B 281:20141097.

Koper N, Schmiegelow FKA, Merrill EH (2007) Residuals cannot distinguish between ecological effects of habitat amount and fragmentation: implications for the debate. Landscape Ecology 22:811-820.

Kuussaari M, Bommarco R, Heikkinen RK, et al. (2009) Extinction debt: a challenge for biodiversity conservation. Trends in Ecology and Evolution 24:564-571.

Laurance WF, Useche DC, Rendeiro J, et al. (2012) Averting biodiversity collapse in tropical forest protected areas. Nature 489:290-294.

Lawes MJ, Eeley HAC, Findlay NJ, et al. (2007) Resilient forest faunal communities in South Africa: a legacy of palaeoclimatic change and extinction filtering? Journal of Biogeography 34:1246-1264.

Legendre P (2014) Interpreting the replacement and richness difference components of beta diversity. Global Ecology and Biogeography 23:1324-1334.

Lundberg I, Johnson R, Stewart BM (2021) What is your estimand? Defining the target quantity connects statistical evidence to theory. American Sociological Review 86:532-565.

MacArthur RH (1972) Geographical Ecology. Harper and Row, New York.

Magioli M, Ferraz KMPMB, Chiarello AG, et al. (2021) Land-use changes lead to functional loss of terrestrial mammals in a Neotropical rainforest. Perspectives in Ecology and Conservation 19:161-170.

Martínez-Lanfranco JA (2026a) Habitat amount control is necessary but not sufficient to resolve the fragmentation debate. Conservation Biology, accepted. EcoEvoRxiv. DOI 10.32942/X2KT05.

Martínez-Lanfranco JA (2026b) Habitat-loss-driven predictor coupling limits inference about the independent effects of configuration in additive habitat-amount models: implications for the fragmentation debate. Ecosphere, in review. bioRxiv. DOI 10.64898/2026.04.12.718042.

Matthews TJ, Cottee-Jones HE, Whittaker RJ (2014) Habitat fragmentation and the species-area relationship: a focus on total species richness obscures the impact of habitat loss on habitat specialists. Diversity and Distributions 20:1136-1146.

McKinney ML, Lockwood JL (1999) Biotic homogenization: a few winners replacing many losers in the next mass extinction. Trends in Ecology and Evolution 14:450-453.

Miguet P, Jackson HB, Jackson ND, et al. (2016) What determines the spatial extent of landscape effects on species? Landscape Ecology 31:1177-1194.

Newbold T, Hudson LN, Contu S, et al. (2018) Widespread winners and narrow-ranged losers: land use homogenizes biodiversity in local assemblages worldwide. PLoS Biology 16:e2006841.

Orme CDL, Mayor S, dos Anjos L, et al. (2019) Distance to range edge determines sensitivity to deforestation. Nature Ecology and Evolution 3:886-891.

Pedersen EJ, Miller DL, Simpson GL, et al. (2019) Hierarchical generalized additive models in ecology: an introduction with mgcv. PeerJ 7:e6876.

Prunier JG, Colyn M, Legendre X, et al. (2015) Multicollinearity in spatial genetics: separating the wheat from the chaff using commonality analyses. Molecular Ecology 24:263-283.

Purvis A, Newbold T, De Palma A, et al. (2018) Modelling and Projecting the Response of Local Terrestrial Biodiversity Worldwide to Land Use and Related Pressures: The PREDICTS Project. Advances in Ecological Research 58:1-35.

Püttker T, Crouzeilles R, Almeida-Gomes M, et al. (2020) Indirect effects of habitat loss via habitat fragmentation: a cross-taxa analysis of forest-dependent species. Biological Conservation 241:108368.

Ray-Mukherjee J, Nimon K, Mukherjee S, et al. (2014) Using commonality analysis in multiple regressions: a tool to decompose regression effects in the face of multicollinearity. Methods in Ecology and Evolution 5:320-328.

Rodrigues DP, Skupien FL, Oliveira Sausen J, et al. (2020) Small mammals in fragments of Atlantic Forest: species richness answering to field methods and environment. Journal of Tropical Ecology 36:101-108.

Ruffell J, Banks-Leite C, Didham RK (2016) Accounting for the causal basis of collinearity when measuring the effects of habitat loss versus habitat fragmentation. Oikos 125:117-125.

Rybicki J, Abrego N, Ovaskainen O (2020) Habitat fragmentation and species diversity in competitive communities. Ecology Letters 23:506-517.

Saura S (2021) The habitat amount hypothesis implies negative effects of habitat fragmentation on species richness. Journal of Biogeography 48:11-22.

Semper-Pascual A, Burton C, Baumann M, et al. (2021) How do habitat amount and habitat fragmentation drive time-delayed responses of biodiversity to land-use change? Proceedings of the Royal Society B 288:20202466.

Senior RA, Hill JK, González del Pliego P, et al. (2017) A pantropical analysis of the impacts of forest degradation and conversion on local temperature. Ecology and Evolution 7:7897-7908.

Smith AC, Koper N, Francis CM, et al. (2009) Confronting collinearity: comparing methods for disentangling the effects of habitat loss and fragmentation. Landscape Ecology 24:1271-1285.

Sobral-Souza T, Bosco NS, Candelaria LP, et al. (2024) Spatial bias in sampling small rodents in the Atlantic Forest: a landscape and accessibility perspective. Perspectives in Ecology and Conservation 22:297-305.

Sobral-Souza T, Stropp J, Santos JP, et al. (2021a) Knowledge gaps hamper understanding the relationship between fragmentation and biodiversity. PeerJ 9:e11673

Socolar JB, Gilroy JJ, Kunin WE, et al. (2016) How should beta-diversity inform biodiversity conservation? Trends in Ecology and Evolution 31:67-80

Spake R, Barajas-Barbosa MP, Blowes SA, et al. (2022) Detecting thresholds of ecological change in the Anthropocene. Annual Review of Environment and Resources 47:797-821.

Spake R, Bowler DE, Callaghan CT, et al. (2023) Understanding 'it depends' in ecology: a guide to hypothesising, visualising and interpreting statistical interactions. Biological Reviews 98:983-1002.

Stein A, Gerstner K, Kreft H (2014) Environmental heterogeneity as a universal driver of species richness across taxa, biomes and spatial scales. Ecology Letters 17:866-880

Sunday JM, Bates AE, Kearney MR, et al. (2014) Thermal-safety margins and the necessity of thermoregulatory behavior across latitude and elevation. Proceedings of the National Academy of Sciences 111:5610-5615.

Taylor PD, Fahrig L, Henein K, et al. (1993) Connectivity is a vital element of landscape structure. Oikos 68:571-573.

Tews J, Brose U, Grimm V, et al. (2004) Animal species diversity driven by habitat heterogeneity/diversity: the importance of keystone structures. Journal of Biogeography 31:79-92.

Tredennick AT, Hooker G, Ellner SP, et al. (2021) A practical guide to selecting models for exploration, inference, and prediction in ecology. Ecology 102:e03336.

Tuanmu MN, Jetz W (2015) A global, remote sensing-based characterization of terrestrial habitat heterogeneity. Global Ecology and Biogeography 24:1329-1339.

Tukiainen H, Bailey JJ, Field R, et al. (2017) Combining geodiversity with climate and topography to account for threatened species richness. Conservation Biology 31:364-375.

Valente JJ, Betts MG (2019) Response to fragmentation by avian communities is mediated by species traits. Diversity and Distributions 25:48-60.

Villard MA, Metzger JP (2014) Beyond the fragmentation debate: a conceptual model to predict when habitat configuration really matters. Journal of Applied Ecology 51:309-318.

Zhang H, Chase JM, Liao J (2024) Habitat amount modulates biodiversity responses to fragmentation. Nature Ecology and Evolution 8:1437-1447.

Table 6.1. Shared inferential conditions across the Atlantic Forest empirical chapters and the conclusions retained after each condition is applied. Predictor-space adequacy concerns whether the observed predictor combinations identify the ecological contrast of interest. Response-estimand alignment concerns whether the response measures the quantity being interpreted. Metric-mechanism correspondence concerns whether the operational metric maps onto the proposed process. Domain stability concerns whether the interpretation persists across the spatial, temporal, definitional and geographic domains of application. Statistical detectability is an intervening diagnostic tier and not a fifth condition. Temporal alignment is developed in the text because its consequences differ among chapters.

Inferential condition	**Climate-conditioned land-use effects**	**Origin-dependent heterogeneity**	**Amount-conditioned forest configuration**	**Inference retained**
Predictor-space adequacy	Cropland prevalence and non-natural cover each aggregate several pathways of transformation, and the fitted climate axes stay partly entangled with regional climate, terrain and land-use placement.	Greenness-magnitude variability is a process-coupled predictor whose variation partly follows the land-use gradient it moderates.	Forest amount and configuration are co-generated, so removing the modelled amount dependence defines an operational amount-adjusted component rather than a matched fixed-amount contrast.	The fitted interactions identify conditional associations within the observed predictor spaces. Comparison across products, metrics, formulations and scales shows which conclusions persist, without converting those contrasts into causal interventions.

Inferential condition	Climate-conditioned land-use effects	Origin-dependent heterogeneity	Amount-conditioned forest configuration	Inference retained
Response-estimand alignment	Aggregate richness dilutes the specialist thermal response, while coefficient discordance and low precision bound interpretation of the weaker guilds.	Guild-opposed effects cancel in aggregate richness, and non-forest richness is a directional comparator, not a formal negative control.	Restricted total richness is the arithmetic sum of its three guild components, so the weaker aggregate association localises rather than replicates the specialist result. Guild composition and composition-dependent sampling remain inseparable from the specialist–non-forest mirror.	Guild partitioning reveals response structure that aggregate richness obscures. The repeated specialist concentration holds within the shared response architecture, while abundance, turnover and demographic processes remain unresolved.

Inferential condition	Climate-conditioned land-use effects	Origin-dependent heterogeneity	Amount-conditioned forest configuration	Inference retained
Metric-mechanism correspondence	Daily and annual temperature range index one shared atmospheric variability axis rather than realised microclimate, and their concordance does not separate diurnal from seasonal timescales.	Greenness-magnitude variability and configurational greenness diversity measure different forms of vegetation texture, and water-energy conditioning does not measure demographic buffering.	Configuration metrics represent different structural dimensions. The specialist–non-forest mirror was absent under edge density, and amount-conditioned associations remain consistent with extinction filtering without observing it directly.	Mechanism-linked predictions separate candidate pathways through predicted direction, guild specificity and robustness contrasts. They rank interpretations without identifying the intervening mechanism.

Inferential condition	Climate-conditioned land-use effects	Origin-dependent heterogeneity	Amount-conditioned forest configuration	Inference retained
Domain stability	Scale coherence does not identify one ecological scale of canopy buffering, and the result stays defined within the climate axes and land-use surrogates analysed.	Ordered scale changes do not establish a transition between demographic constraint and niche diversification, and vegetation texture may carry different meanings in other landscape contexts.	Forest-cover positions are relative to each design cell, configuration may follow a different change history from amount, and extent profiles confound ecology with spatial support and predictor averaging.	The core conclusions persist across broad portions of the product, metric, formulation and scale matrix within the Atlantic Forest, while metric-specific departures delimit their scope. Absolute thresholds, process scales and transferability require direct testing.

Complete list of literature cited

Allouche O, Kalyuzhny M, Moreno-Rueda G, Pizarro M, Kadmon R (2012) Area-heterogeneity tradeoff and the diversity of ecological communities. Proceedings of the National Academy of Sciences 109:17495-17500.

Almeida-Gomes M, Prevedello JA, Crouzeilles R (2016) The use of native vegetation as a proxy for habitat may overestimate habitat availability in fragmented landscapes. Landscape Ecology 31:711-719.

Alroy J (2017) Effects of habitat disturbance on tropical forest biodiversity. Proceedings of the National Academy of Sciences 114:6056-6061.

Amatulli G, Domisch S, Tuanmu M-N, Parmentier B, Ranipeta A, Malczyk J, Jetz W (2018) A suite of global, cross-scale topographic variables for environmental and biodiversity modeling. Scientific Data 5:180040.

Anderson MJ, Crist TO, Chase JM, Vellend M, Inouye BD, Freestone AL, Sanders NJ, Cornell HV, Comita LS, Davies KF, Harrison SP, Kraft NJB, Stegen JC, Swenson NG (2011) Navigating the multiple meanings of β diversity: a roadmap for the practicing ecologist. Ecology Letters 14:19–28

Andrén H (1994) Effects of habitat fragmentation on birds and mammals in landscapes with different proportions of suitable habitat: a review. Oikos 71:355-366.

Arel-Bundock V, Greifer N, Heiss A (2024) How to interpret statistical models using marginaleffects for R and Python. Journal of Statistical Software 111:1-32.

Arroyo-Rodríguez V, Fahrig L, Tabarelli M, et al. (2020) Designing optimal human-modified landscapes for forest biodiversity conservation. Ecology Letters 23:1404-1420.

Arroyo-Rodríguez V, Saldaña-Vázquez RA, Fahrig L, Santos BA (2017) Does forest fragmentation cause an increase in forest temperature? Ecological Research 32:81-88.

Balmford A (1996) Extinction filters and current resilience: the significance of past selection pressures for conservation biology. Trends in Ecology and Evolution 11:193-196.

Banks-Leite C, Betts MG, Ewers RM, Orme CDL, Pigot AL (2022) The macroecology of landscape ecology. Trends in Ecology and Evolution 37:480-487.

Banks-Leite C, Pardini R, Tambosi LR, Pearse WD, Bueno AA, Bruscagin RT, Condez TH, Dixo M, Igari AT, Martensen AC, Metzger JP (2014) Using ecological thresholds to evaluate the costs and benefits of set-asides in a biodiversity hotspot. Science 345:1041-1045.

Bar-Massada A, Wood EM (2014) The richness-heterogeneity relationship differs between heterogeneity measures within and among habitats. Ecography 37:528-535.

Barr DJ, Levy R, Scheepers C, Tily HJ (2013) Random effects structure for confirmatory hypothesis testing: Keep it maximal. Journal of Memory and Language 68:255-278.

Baselga A (2010) Partitioning the turnover and nestedness components of beta diversity. Global Ecology and Biogeography 19:134-143.

Baston D (2023) exactextractr: fast extraction from raster datasets using polygons. R package version 0.10.0.

Bélisle M (2005) Measuring landscape connectivity: the challenge of behavioral landscape ecology. Ecology 86:1988-1995.

Ben-Hur E, Kadmon R (2020) An experimental test of the area-heterogeneity tradeoff. Proceedings of the National Academy of Sciences 117:4815-4822.

Betts MG, Fahrig L, Hadley AS, et al. (2014) A species-centered approach for uncovering generalities in organism responses to habitat loss and fragmentation. Ecography 37:517-527.

Betts MG, Wolf C, Pfeifer M, Banks-Leite C, Arroyo-Rodríguez V, Ribeiro DB, et al. (2019) Extinction filters mediate the global effects of habitat fragmentation on animals. Science 366:1236-1239.

Betts MG, Wolf C, Ripple WJ, Phalan B, Millers KA, Duarte A, Butchart SHM, Levi T (2017) Global forest loss disproportionately erodes biodiversity in intact landscapes. Nature 547:441–444

Bevan PA, Braga Ferreira G, Ingram DJ, et al. (2025) Regional Biomes outperform broader spatial units in capturing biodiversity responses to land-use change. Ecography 2025:e07318.

Bolker BM, Brooks ME, Clark CJ, Geange SW, Poulsen JR, Stevens MHH, White J-SS (2009) Generalized linear mixed models: a practical guide for ecology and evolution. Trends in Ecology and Evolution 24:127-135.

Bosco L, Cushman SA, Wan HY, et al. (2021) Fragmentation effects on woodlark habitat selection depend on habitat amount and spatial scale. Animal Conservation 24:84–94

Bouchet PJ, Meeuwig JJ, Salgado Kent CP, et al. (2015) Topographic determinants of mobile vertebrate predator hotspots: current knowledge and future directions. Biological Reviews 90:699-728.

Bovendorp RS, Villar N, de Abreu-Junior EF, et al. (2017) Atlantic small-mammal: a dataset of communities of rodents and marsupials of the Atlantic forests of South America. Ecology 98:2226.

Brook BW, Sodhi NS, Bradshaw CJA (2008) Synergies among extinction drivers under global change. Trends in Ecology and Evolution 23:453-460.

Brooks ME, Kristensen K, van Benthem KJ, et al. (2017) glmmTMB balances speed and flexibility among packages for zero-inflated generalized linear mixed modeling. The R Journal 9:378-400.

Burnham KP, Anderson DR (2002) Model Selection and Multimodel Inference: A Practical Information-Theoretic Approach, 2nd edn. Springer-Verlag, New York.

Burnham KP, Anderson DR, Huyvaert KP (2011) AIC model selection and multimodel inference in behavioral ecology: some background, observations, and comparisons. Behavioral Ecology and Sociobiology 65:23-35.

Carnaval AC, Moritz C (2008) Historical climate modelling predicts patterns of current biodiversity in the Brazilian Atlantic forest. Journal of Biogeography 35:1187-1201.

Carnicer J, Brotons L, Herrando S, Sol D (2013) Improved empirical tests of area-heterogeneity tradeoffs. Proceedings of the National Academy of Sciences 110:E2858-E2860.

Chase JM, Jeliazkov A, Ladouceur E, Viana DS (2020) Biodiversity conservation through the lens of metacommunity ecology. Annals of the New York Academy of Sciences 1469:86-104.

Chase JM, Kraft NJB, Smith KG, Vellend M, Inouye BD (2011) Using null models to disentangle variation in community dissimilarity from variation in α-diversity. Ecosphere, 2(2), art24.

Chase JM, Liebergesell M, Sagouis A, et al (2020) FragSAD: a database of diversity and species abundance distributions from habitat fragments. Ecology 100:e02775.

Clavel J, Julliard R, Devictor V (2011) Worldwide decline of specialist species: toward a global functional homogenization? Frontiers in Ecology and the Environment 9:222-228.

Coelho MTP, Barreto E, Rangel TF, Diniz-Filho JAF, Wüest RO, Bach W, Skeels A, McFadden IR, Roberts DW, Pellissier L, Zimmermann NE (2023) The geography of climate and the global patterns of species diversity. Nature 622:537-544.

Crase B, Liedloff AC, Wintle BA (2012) A new method for dealing with residual spatial autocorrelation in species distribution models. Ecography 35:879-888.

Currie DJ (1991) Energy and large-scale patterns of animal- and plant-species richness. The American Naturalist 137:27-49.

Currie DJ, Mittelbach GG, Cornell HV, et al. (2004) Predictions and tests of climate-based hypotheses of broad-scale variation in taxonomic richness. Ecology Letters 7:1121-1134.

Daly L, Gabriel J, Hajdasz AC, et al. (2025) Connectivity of forest patches via wooded corridors increases biodiversity at low, but not high, forest amounts. Conservation Letters 18:e13154

Daskalova GN, Myers-Smith IH, Bjorkman AD, et al. (2020) Landscape-scale forest loss as a catalyst of population and biodiversity change. Science 368(6497):1341-1347.

De Camargo RX, Boucher-Lalonde V, Currie DJ (2018) At the landscape level, birds respond strongly to habitat amount but weakly to fragmentation. Diversity and Distributions 24:629–639

De Frenne P, Zellweger F, Rodríguez-Sánchez F, et al. (2019) Global buffering of temperatures under forest canopies. Nature Ecology and Evolution 3:744-749.

De Palma A, Sanchez-Ortiz K, Martin PA, Chadwick A, Gilbert G, Bates AE, Börger L, Contu S, Hill SLL, Purvis A (2018) Challenges with inferring how land-use affects terrestrial

biodiversity: study design, time, space and synthesis. Advances in Ecological Research 58:163-199.

de Souza Leite M, Boesing AL, Metzger JP, Prado PI (2022) Matrix quality determines the strength of habitat loss filtering on bird communities at the landscape scale. Journal of Applied Ecology 59:2790–2802

Dennis M, Huck J, Holt C, et al. (2026) Beyond the patch: leveraging functional habitat delineation in fragmentation-biodiversity research. Landscape Ecology 41:37

Dennis M, Huck JJ (2025) The continuity-contiguity problem in fragmentation-biodiversity research. Journal of Biogeography 52:e15077.

Dennis M, Huck JJ, Holt CD, da Conceição Bispo P, McHenry E, Speak A, James P (2024) Land-cover gradients determine alternate drivers of mammalian species richness in fragmented landscapes. Landscape Ecology 39:146.

Devictor V, Clavel J, Julliard R, Lavergne S, Mouillot D, Thuiller W, Venail P, Villéger S, Mouquet N (2010) Defining and measuring ecological specialization. Journal of Applied Ecology 47:15-25.

Devictor V, Julliard R, Jiguet F (2008) Distribution of specialist and generalist species along spatial gradients of habitat disturbance and fragmentation. Oikos 117:507-514.

Didham RK, Kapos V, Ewers RM (2012) Rethinking the conceptual foundations of habitat fragmentation research. Oikos 121:161-170.

Donald PF, Arendarczyk B, Spooner F, Buchanan GM (2019) Loss of forest intactness elevates global extinction risk in birds. Animal Conservation 22:341-347.

Dormann CF, Elith J, Bacher S, Buchmann C, Carl G, Carré G, Marquéz JRG, Gruber B, Lafourcade B, Leitão PJ, Münkemüller T, McClean C, Osborne PE, Reineking B, Schröder B, Skidmore AK, Zurell D, Lautenbach S (2013) Collinearity: a review of methods to deal with it and a simulation study evaluating their performance. Ecography 36:27-46.

Duncan RP, Kefford BJ (2021) Interactions in statistical models: three things to know. Methods in Ecology and Evolution 12:2287-2297.

Edwards DP, Socolar JB, Mills SC, Burivalova Z, Koh LP, Wilcove DS (2019) Conservation of Tropical Forests in the Anthropocene. Current Biology 29:R1008-R1020.

Evans KL, Warren PH, Gaston KJ (2005) Species-energy relationships at the macroecological scale: a review of the mechanisms. Biological Reviews 80:1-25.

Ewers RM, Banks-Leite C (2013) Fragmentation impairs the microclimate buffering effect of tropical forests. PLoS One 8:e58093.

Ewers RM, Didham RK (2006) Confounding factors in the detection of species responses to habitat fragmentation. Biological Reviews 81:117-142.

Fahrig L (1997) Relative effects of habitat loss and fragmentation on population extinction. Journal of Wildlife Management 61: 603-610.

Fahrig L (1998) When does fragmentation of breeding habitat affect population survival? Ecological Modelling 105:273-292.

Fahrig L (1999) Forest loss and fragmentation: which has the greater effect on persistence of forest-dwelling animals? In: Rochelle JA, Lehmann LA, Wisniewski J (eds) Forest fragmentation: wildlife and management implications. Brill, Leiden, pp 87–95

Fahrig L (2002) Effect of habitat fragmentation on the extinction threshold: a synthesis. Ecological Applications 12:346–353

Fahrig L (2003) Effects of habitat fragmentation on biodiversity. Annual Review of Ecology, Evolution, and Systematics 34:487-515.

Fahrig L (2013) Rethinking patch size and isolation effects: the habitat amount hypothesis. Journal of Biogeography 40:1649-1663.

Fahrig L (2015) Just a hypothesis: a reply to Hanski. Journal of Biogeography 42:993–994

Fahrig L (2017) Ecological responses to habitat fragmentation per se. Annual Review of Ecology, Evolution, and Systematics 48:1-23.

Fahrig L (2021) What the habitat amount hypothesis does and does not predict: a reply to Saura. Journal of Biogeography 48:1530-1535.

Fahrig L, Arroyo-Rodríguez V, Bennett JR, et al. (2019) Is habitat fragmentation bad for biodiversity? Biological Conservation 230:179-186.

Fahrig L, Arroyo-Rodríguez V, Cazetta E, Ford A, Lancaster J, Ranius T (2021) Landscape connectivity. In: Perry G, Francis R, Millington J, Minor E (eds) The Routledge Handbook of Landscape Ecology. Taylor and Francis, pp 67–88.

Fahrig L, Baudry J, Brotons L, Burel FG, Crist TO, Fuller RJ, Sirami C, Siriwardena GM, Martin JL (2011) Functional landscape heterogeneity and animal biodiversity in agricultural landscapes. Ecology Letters 14:101–112

Fahrig L, Galán-Acedo C, Edwards BPM, Habrich AK, Martin AE, Albert CH, Arroyo-Rodríguez V, Binley AD, Cole JR, Lockhart JK, Nielsen SE, Riva F, Schmoeller MPR, Tscharntke T (2026) Why controlling for habitat amount is critical for resolving the fragmentation debate. Conservation Biology e70245

Fieberg JR, Vitense K, Johnson DH (2020) Resampling-based methods for biologists. PeerJ 8:e9089.

Figueiredo MSL, Barros CS, Delciellos AC, et al. (2017) Abundance of small mammals in the Atlantic Forest (ASMAF): a data set for analyzing tropical community patterns. Ecology 98:2981.

Figueiredo MSL, Weber MM, Brasileiro CA, et al. (2021) Tetrapod diversity in the Atlantic Forest: maps and gaps. In: Marques MCM, Grelle CEV (eds) The Atlantic Forest: history, biodiversity, threats and opportunities of the mega-diverse forest. Springer, Cham, pp 185-203.

Fischer J, Lindenmayer DB (2007) Landscape modification and habitat fragmentation: a synthesis. Global Ecology and Biogeography 16:265-280.

Flather CH, Bevers M (2002) Patchy reaction-diffusion and population abundance: the relative importance of habitat amount and arrangement. The American Naturalist 159:40-56.

Fletcher RJ Jr, Betts MG, Damschen EI, Hefley TJ, Hightower J, Smith TA, Fortin MJ, Haddad NM (2023) Addressing the problem of scale that emerges with habitat fragmentation. Global Ecology and Biogeography 32:828-841.

Fletcher RJ Jr, Bullock JM, Graham LJ, Kortessis N, Watts K (2026) Resurrecting habitat fragmentation as a process over time. Trends in Ecology & Evolution [in press]. https://doi.org/10.1016/j.tree.2026.03.009

Fletcher RJ Jr, Burrell N, Reichert BE, Vasudev D (2016) Divergent perspectives on landscape connectivity reveal consistent effects from genes to communities. Current Landscape Ecology Reports 1:67–79

Fletcher RJ Jr, Didham RK, Banks-Leite C, et al. (2018) Is habitat fragmentation good for biodiversity? Biological Conservation 226:9-15.

Fletcher RJ Jr, Fortin M-J (2018) Spatial Ecology and Conservation Modeling. Springer, Cham.

Fletcher RJ Jr, Green RE, Bladon EK, Atkinson PW, Phalan BT, Williams D, Visconti P, Balmford A (2025) Beyond species richness for biological conservation. Conservation Letters 18:e13124.

Fletcher RJ, Reichert BE, Holmes K (2018) The negative effects of habitat fragmentation operate at the scale of dispersal. Ecology 99:2176-2186.

Fletcher RJ, Ries L, Battin J, Chalfoun AD (2007) The role of habitat area and edge in fragmented landscapes: definitively distinct or inevitably intertwined? Canadian Journal of Zoology 85:1017-1030.

Foley JA, DeFries R, Asner GP, et al. (2005) Global consequences of land use. Science 309:570-574.

Frazier AE, Kedron P (2017) Landscape metrics: past progress and future directions. Current Landscape Ecology Reports 2:63-72.

Freckleton RP (2002) On the misuse of residuals in ecology: regression of residuals vs. multiple regression. Journal of Animal Ecology 71: 542-545.

Frishkoff LO, Karp DS, Flanders JR, et al. (2016) Climate change and habitat conversion favour the same species. Ecology Letters 19:1081-1090.

Ganuza C, Redlich S, Uhler J, et al. (2022) Interactive effects of climate and land use on pollinator diversity differ among taxa and scales. Science Advances 8:eabm9359.

Gardner RH, Milne BT, Turner MG, O'Neill RV (1987) Neutral models for the analysis of broad-scale landscape pattern. Landscape Ecology 1:19–28

Gengler NW, Acevedo MA, Branch LC (2024) Habitat configuration influences mammal populations at a wider spatial extent than habitat composition: a meta-analysis of forest mammal datasets. Landscape Ecology 39:2.

Geres LS, Richter T, Seidl R, et al. (2025) Macro- and microclimate interactively shape species diversity of multiple taxa in mountain landscapes. Ecography 2025:e07984.

Gonçalves-Souza T, Chase JM, Haddad NM, Vancine MH, Didham RK, Melo FLP, Aizen MA, Bernard E, Chiarello AG, Faria D, Gibb H, de Lima MG, Magnago LFS, Mariano-Neto E, Nogueira AA, Nemésio A, Passamani M, Pinho BX, Rocha-Santos L, et al. (2025) Species turnover does not rescue biodiversity in fragmented landscapes. Nature 640:702-706.

Gonzalez A, Loreau M (2009) The causes and consequences of compensatory dynamics in ecological communities. Annual Review of Ecology, Evolution, and Systematics 40:393-414.

Gunderson AR (2023) Habitat degradation exacerbates the effects of anthropogenic warming by removing thermal refuges. Global Change Biology 29:6157-6158.

Haddad NM, Brudvig LA, Clobert J, et al. (2015) Habitat fragmentation and its lasting impact on Earth's ecosystems. Science Advances 1:e1500052.

Haddad NM, Gonzalez A, Brudvig LA, Burt MA, Levey DJ, Damschen EI (2017a) Experimental evidence does not support the Habitat Amount Hypothesis. Ecography 40:48–55

Haddad NM, Holt RD, Fletcher RJ Jr, Loreau M, Clobert J (2017b) Connecting models, data, and concepts to understand fragmentation's ecosystem-wide effects. Ecography 40:1–8

Hadley AS, Betts MG (2016) Refocusing habitat fragmentation research using lessons from the last decade. Current Landscape Ecology Reports 1:55-66.

Halstead KE, Alexander JD, Hadley AS, et al. (2019) Using a species-centered approach to predict bird community responses to habitat fragmentation. Landscape Ecology 34:1919–1935

Hansen MC, Potapov PV, Moore R, et al. (2013) High-resolution global maps of 21st-century forest cover change. Science 342:850-853.

Hanski I (2015) Habitat fragmentation and species richness. Journal of Biogeography 42:989-993.

Harrison XA (2014) Using observation-level random effects to model overdispersion in count data in ecology and evolution. PeerJ 2:e616.

Hart T, Bayne EM (2025) Habitat amount–fragmentation interactions drive Canada Warbler dynamics across spatial scales. Avian Conservation and Ecology 20(2):15

Hartig F (2022) DHARMa: residual diagnostics for hierarchical (multi-level / mixed) regression models. R package version 0.4.6.

Hasui E, Martensen AC, Uezu A, et al. (2024) Populations across bird species distribution ranges respond differently to habitat loss and fragmentation: implications for conservation strategies. Perspectives in Ecology and Conservation 22:43-54.

Hasui E, Metzger JP, Pimentel RG, et al. (2018) ATLANTIC BIRDS: a data set of bird species from the Brazilian Atlantic Forest. Ecology 99:497.

Hatfield JH, Barlow J, Joly CA, et al. (2020) Mediation of area and edge effects in forest fragments by adjacent land use. Conservation Biology 34:395-404.

Hawkins BA, Field R, Cornell HV, et al. (2003) Energy, water, and broad-scale geographic patterns of species richness. Ecology 84:3105-3117.

Heidrich L, Bae S, Friess N, et al. (2020) Heterogeneity-diversity relationships differ between and within trophic levels in temperate forests. Nature Ecology and Evolution 4:1204-1212.

Henle K, Davies KF, Kleyer M, Margules C, Settele J (2004) Predictors of species sensitivity to fragmentation. Biodiversity and Conservation 13:207-251.

Hesselbarth MHK, Sciaini M, With KA, Wiegand K, Nowosad J (2019) landscapemetrics: an open-source R tool to calculate landscape metrics. Ecography 42:1648-1657.

Hijmans RJ (2022) terra: spatial data analysis. R package version 1.6-17.

Hijmans RJ, Cameron SE, Parra JL, et al. (2005) Very high resolution interpolated climate surfaces for global land areas. International Journal of Climatology 25:1965-1978.

Hillebrand H, Blasius B, Borer ET, et al. (2018) Biodiversity change is uncoupled from species richness trends: consequences for conservation and monitoring. Journal of Applied Ecology 55:169-184.

Hjort J, Gordon JE, Gray M, Hunter ML Jr (2015) Why geodiversity matters in valuing nature's stage. Conservation Biology 29:630-639.

Hortal J, Carrascal LM, Triantis KA, Thébault E, Meiri S, Sfenthourakis S (2013) Species richness can decrease with altitude but not with habitat diversity. Proceedings of the National Academy of Sciences 110:E2149-E2150.

Hua F, Wang W, Nakagawa S, Liu S, Miao X, Yu L, Du Z, Abrahamczyk S, Arias-Sosa LA, Buda K, Budka M (2024) Ecological filtering shapes the impacts of agricultural deforestation on biodiversity. Nature Ecology & Evolution 8:251–266

Hurlbert AH, Stegen JC (2014) When should species richness be energy limited, and how would we know? Ecology Letters 17:401-413.

Hutchinson GE (1957) Concluding remarks. Cold Spring Harbor Symposia on Quantitative Biology 22:415-427.

Ismaeel A, Tai APK, Santos EG, et al. (2024) Patterns of tropical forest understory temperatures. Nature Communications 15:549.

IUCN (2012) IUCN Habitats Classification Scheme, Version 3.1. International Union for Conservation of Nature, Gland, Switzerland and Cambridge, UK.

Jackson HB, Fahrig L (2015) Are ecologists conducting research at the optimal scale? Global Ecology and Biogeography 24:52-63.

Jorge MLSP, Galetti M, Ribeiro MC, et al. (2013) Mammal defaunation as surrogate of trophic cascades in a biodiversity hotspot. Biological Conservation 163:49-57.

Jucker T, Jackson TD, Zellweger F, et al. (2020) A research agenda for microclimate ecology in human-modified tropical forests. Frontiers in Forests and Global Change 2:92.

Kadmon R, Allouche O (2007) Integrating the effects of area, isolation, and habitat heterogeneity on species diversity: a unification of island biogeography and niche theory. The American Naturalist 170:443-454.

Kemppinen J, Lembrechts JJ, Van Meerbeek K, et al. (2024) Microclimate, an important part of ecology and biogeography. Global Ecology and Biogeography 33:e13834.

Kerr JT, Packer L (1997) Habitat heterogeneity as a determinant of mammal species richness in high-energy regions. Nature 385:252-254.

Khaliq I, et al. (2014) Interspecific variation in responses to thermal environments across vertebrate ectotherms and endotherms. Oecologia 174:1277-1287.

Khaliq I, Hof C, Prinzinger R, et al. (2014) Global variation in thermal tolerances and vulnerability of endotherms to climate change. Proceedings of the Royal Society B 281:20141097.

Koper N, Schmiegelow FKA, Merrill EH (2007) Residuals cannot distinguish between ecological effects of habitat amount and fragmentation: implications for the debate. Landscape Ecology 22:811-820.

Kramer JMF, Zwiener VP, Camana M, de Lima RAF, Müller SC (2026) Effects of deforestation and fragmentation on tree beta diversity and endemism across landscapes in the Atlantic Forest biodiversity hotspot. Global Change Biology 32:e70855.

Kramer JMF, Zwiener VP, Müller SC (2023) Biotic homogenization and differentiation of plant communities in tropical and subtropical forests. Conservation Biology, 37, e14025.

Kuussaari M, Bommarco R, Heikkinen RK, et al. (2009) Extinction debt: a challenge for biodiversity conservation. Trends in Ecology and Evolution 24:564-571.

Laanisto L, Tamme R, Hiiesalu I, Szava-Kovats R, Gazol A, Pärtel M (2013) Microfragmentation concept explains non-positive environmental heterogeneity–diversity relationships. Oecologia 171:217–226

Laurance WF, Useche DC, Rendeiro J, et al. (2012) Averting biodiversity collapse in tropical forest protected areas. Nature 489:290-294.

Lawes MJ, Eeley HAC, Findlay NJ, et al. (2007) Resilient forest faunal communities in South Africa: a legacy of palaeoclimatic change and extinction filtering? Journal of Biogeography 34:1246-1264.

Legendre P (2014) Interpreting the replacement and richness difference components of beta diversity. Global Ecology and Biogeography 23:1324-1334.

Lehtilä K, Vinter T, Dinnetz P (2020) Plant response to habitat amount and configuration in Swedish forests. Diversity and Distributions 26:329–339

Levin SA (1992) The problem of pattern and scale in ecology. Ecology 73:1943-1967.

Li H, Wu J (2004) Use and misuse of landscape indices. Landscape Ecology 19:389-399.

Lima F, Beca G, Muylaert RL, et al. (2017) ATLANTIC-CAMTRAPS: a dataset of medium and large terrestrial mammal communities in the Atlantic Forest of South America. Ecology 98:2979.

Lundberg I, Johnson R, Stewart BM (2021) What is your estimand? Defining the target quantity connects statistical evidence to theory. American Sociological Review 86:532-565.

MacArthur RH (1972) Geographical Ecology. Harper and Row, New York.

Magioli M, Ferraz KMPMB, Chiarello AG, et al. (2021) Land-use changes lead to functional loss of terrestrial mammals in a Neotropical rainforest. Perspectives in Ecology and Conservation 19:161-170.

Mantyka-Pringle CS, Martin TG, Rhodes JR (2012) Interactions between climate and habitat loss effects on biodiversity: a systematic review and meta-analysis. Global Change Biology 18:1239-1252. Correction (2013) Global Change Biology 19:1642-1644.

Marques MCM, Grelle CEV (eds) (2021) The Atlantic Forest: history, biodiversity, threats and opportunities of the mega-diverse forest. Springer, Cham.

Marques MCM, Trindade W, Bohn A, Grelle CEV (2021) The Atlantic Forest: an introduction to the megadiverse forest of South America. In: Marques MCM, Grelle CEV (eds) The Atlantic Forest: history, biodiversity, threats and opportunities of the mega-diverse forest. Springer, Cham, pp 3-23.

Martínez-Lanfranco JA (2026a) Habitat amount control is necessary but not sufficient to resolve the fragmentation debate. Conservation Biology, accepted. EcoEvoRxiv: https://doi.org/10.32942/X2KT05

Martínez-Lanfranco JA (2026b) Habitat-loss-driven predictor coupling limits inference about the independent effects of configuration in additive habitat-amount models: implications for the fragmentation debate. Ecosphere, in review. bioRxiv: https://doi.org/10.64898/2026.04.12.718042

Martínez-Lanfranco JA (2026c) A patch-number null does not rule out conditional configuration effects. Conservation Letters 19:e70063.

Martínez-Lanfranco JA (2026d) Configuration coefficients can mislead fragmentation-based conservation inference from beta diversity. Conservation Biology, in review. EcoEvoRxiv, https://ecoevorxiv.org/repository/view/14049/

Matthews TJ, Cottee-Jones HE, Whittaker RJ (2014) Habitat fragmentation and the species-area relationship: a focus on total species richness obscures the impact of habitat loss on habitat specialists. Diversity and Distributions 20: 1136-1146.

McGarigal K, Cushman SA (2002) Comparative evaluation of experimental approaches to the study of habitat fragmentation effects. Ecological Applications 12:335–345

McGarigal K, Cushman SA, Ene E (2012) FRAGSTATS v4: Spatial Pattern Analysis Program for Categorical and Continuous Maps. University of Massachusetts, Amherst.

McGarigal K, Wan HY, Zeller KA, Timm BC, Cushman SA (2016) Multi-scale habitat selection modelling: a review and outlook. Landscape Ecology 31: 1161-1175.

McKinney ML, Lockwood JL (1999) Biotic homogenization: a few winners replacing many losers in the next mass extinction. Trends in Ecology and Evolution 14:450-453.

Mendes CB, Prevedello JA (2020) Does habitat fragmentation affect landscape-level temperatures? A global analysis. Landscape Ecology 35:1743-1756.

Metzger MJ, Bunce RGH, Jongman RHG, Sayre R, Trabucco A, Zomer R (2013) A high-resolution bioclimate map of the world: a unifying framework for global biodiversity research and monitoring. Global Ecology and Biogeography 22:630-638.

Meyer WB, Turner BL (1992) Human population growth and global land-use/cover change. Annual Review of Ecology and Systematics 23:39-61.

Meza-Elizalde MC, Armenteras-Pascual D (2021) Edge influence on the microclimate and vegetation of fragments of a north Amazonian forest. Forest Ecology and Management 498:119546.

Miguet P, Jackson HB, Jackson ND, Martin AE, Fahrig L (2016) What determines the spatial extent of landscape effects on species? Landscape Ecology 31: 1177-1194.

Miller-Rushing AJ, Primack RB, Devictor V, Corlett RT, Cumming GS, Loyola R, Maas B, Pejchar L (2019) How does habitat fragmentation affect biodiversity? A controversial question at the core of conservation biology. Biological Conservation 232:271-273.

Montràs-Janer T, Suggitt AJ, Fox R, et al. (2024) Anthropogenic climate and land-use change drive short- and long-term biodiversity shifts across taxa. Nature Ecology & Evolution 8:739-751.

Mori AS, Furukawa T, Sasaki T (2013) Response diversity determines the resilience of ecosystems to environmental change. Biological Reviews 88:349-364.

Moura MR, Argôlo AJ, Costa HC (2017b) Historical and contemporary correlates of snake biogeographical subregions in the Atlantic Forest hotspot. Journal of Biogeography 44:640-650.

Moura MR, Costa HC, Argôlo AJS, et al. (2017a) Environmental constraints on the compositional and phylogenetic beta-diversity of tropical forest snake assemblages. Journal of Animal Ecology 86:1192-1204.

Moura MR, Villalobos F, Costa GC, Garcia PCA (2016) Disentangling the role of climate, topography and vegetation in species richness gradients. PLoS ONE 11:e0152468.

Muylaert RDL, Stevens RD, Esbérard CEL, et al. (2017) ATLANTIC BATS: a data set of bat communities from the Atlantic Forests of South America. Ecology 98:3227.

Muylaert RL, Vancine MH, Bernardo R, et al (2018) Uma nota sobre os limites territoriais da Mata Atlântica. Oecologia Australis 22:302-311.

Myers JA, Chase JM, Jiménez I, Jørgensen PM, Araujo-Murakami A, Paniagua-Zambrana N, Seidel R (2013) Beta-diversity in temperate and tropical forests reflects dissimilar mechanisms of community assembly. Ecology Letters 16:151-157.

Nakagawa S, Schielzeth H (2013) A general and simple method for obtaining R^2 from generalized linear mixed-effects models. Methods in Ecology and Evolution 4:133-142.

Neel MC, McGarigal K, Cushman SA (2004) Behavior of class-level landscape metrics across gradients of class aggregation and area. Landscape Ecology 19:435-455.

Newbold T (2018) Future effects of climate and land-use change on terrestrial vertebrate community diversity under different scenarios. Proceedings of the Royal Society B 285:20180792.

Newbold T, Hudson LN, Contu S, et al. (2018) Widespread winners and narrow-ranged losers: land use homogenizes biodiversity in local assemblages worldwide. PLoS Biology 16:e2006841.

Newbold T, Hudson LN, Hill SLL, et al. (2015) Global effects of land use on local terrestrial biodiversity. Nature 520:45-50.

Newbold T, Oppenheimer P, Etard A, Williams JJ (2020) Tropical and Mediterranean biodiversity is disproportionately sensitive to land-use and climate change. Nature Ecology and Evolution 4:1630-1638.

Northrup JM, Rivers JW, Yang Z, Betts MG (2019) Synergistic effects of climate and land-use change influence broad-scale avian population declines. Global Change Biology 25:1561-1575.

O'Hara RB, Kotze DJ (2010) Do not log-transform count data. Methods in Ecology and Evolution 1:118-122.

Oliver TH, Morecroft MD (2014) Interactions between climate change and land use change on biodiversity: attribution problems, risks, and opportunities. Wiley Interdisciplinary Reviews: Climate Change 5:317-335.

Orme CDL, Mayor S, dos Anjos L, et al. (2019) Distance to range edge determines sensitivity to deforestation. Nature Ecology and Evolution 3:886-891.

Ovaskainen O, Soininen J (2011) Making more out of sparse data: hierarchical modeling of species communities. Ecology 92:289–295

Paradis E, Schliep K (2019) ape 5.0: an environment for modern phylogenetics and evolutionary analyses in R. Bioinformatics 35:526-528.

Parks SA, Carroll C, Dobrowski SZ, Allred BW (2020) Human land uses reduce climate connectivity across North America. Global Change Biology 26:2944-2955.

Pärtel M, Bennett JA, Zobel M (2016) Macroecology of biodiversity: disentangling local and regional effects. New Phytologist 211:404-410.

Pebesma E (2018) Simple features for R: standardized support for spatial vector data. The R Journal 10:439-446.

Pedersen EJ, Miller DL, Simpson GL, Ross N (2019) Hierarchical generalized additive models in ecology: an introduction with mgcv. PeerJ 7:e6876

Pfeifer M, Lefebvre V, Peres CA, et al. (2017) Creation of forest edges has a global impact on forest vertebrates. Nature 551:187-191.

Popovic G, Mason TJ, Drobniak SM, et al. (2024) Four principles for improved statistical ecology. Methods in Ecology and Evolution 15:266–281

Prevedello JA, Winck GR, Weber MM, Nichols E, Sinervo B (2019) Impacts of forestation and deforestation on local temperature across the globe. PLoS ONE 14:e0213368.

Prunier JG, Colyn M, Legendre X, et al. (2015) Multicollinearity in spatial genetics: separating the wheat from the chaff using commonality analyses. Molecular Ecology 24:263-283.

Purvis A, Newbold T, De Palma A, et al. (2018) Modelling and Projecting the Response of Local Terrestrial Biodiversity Worldwide to Land Use and Related Pressures: The PREDICTS Project. Advances in Ecological Research 58:1-35.

Püttker T, Crouzeilles R, Almeida-Gomes M, et al. (2020) Indirect effects of habitat loss via habitat fragmentation: a cross-taxa analysis of forest-dependent species. Biological Conservation 241:108368.

R Core Team (2024) R: A Language and Environment for Statistical Computing. R Foundation for Statistical Computing, Vienna, Austria. https://www.R-project.org/.

Räsänen A, Kuitunen M, Hjort J, Vaso A, Kuitunen T, Lensu A (2016) The role of landscape, topography, and geodiversity in explaining vascular plant species richness in a fragmented landscape. Boreal Environment Research 21:53-70.

Ray-Mukherjee J, Nimon K, Mukherjee S, Morris D, Slotow R, Hamer M (2014) Using commonality analysis in multiple regressions: a tool to decompose regression effects in the face of multicollinearity. Methods in Ecology and Evolution 5:320-328.

Read QD, Zarnetske PL, Record S, Dahlin KM, Costanza JK, Finley AO, Gaddis KD, et al (2020) Beyond counts and averages: relating geodiversity to dimensions of biodiversity. Global Ecology and Biogeography 29:696-710.

Redlich S, Zhang J, Benjamin C, et al. (2022) Disentangling effects of climate and land use on biodiversity and ecosystem services - a multi-scale experimental design. Methods in Ecology and Evolution 13:514-527.

Ribeiro MC, Metzger JP, Martensen AC, Ponzoni FJ, Hirota MM (2009) The Brazilian Atlantic Forest: how much is left, and how is the remaining forest distributed? Implications for conservation. Biological Conservation 142:1141-1153.

Ricklefs RE (1987) Community diversity: relative roles of local and regional processes. Science 235:167-171.

Ricklefs RE (2008) Disintegration of the ecological community. The American Naturalist 172:741-750.

Riva F, Fahrig L (2022) The disproportionately high value of small patches for biodiversity conservation. Conservation Letters 15:e12881

Riva F, Fahrig L (2023) Obstruction of biodiversity conservation by minimum patch size criteria. Conservation Biology 37:e14092

Riva F, Koper N, Fahrig L (2024) Overcoming confusion and stigma in habitat fragmentation research. Biological Reviews 99:1411-1424.

Rodrigues DP, Skupien FL, Oliveira Sausen J, et al. (2020) Small mammals in fragments of Atlantic Forest: species richness answering to field methods and environment. Journal of Tropical Ecology 36:101-108.

Rueda M, Hawkins BA, Moreno-Rueda G, Vidanes RM, Ferrero M, Rodríguez MÁ (2013) Does fragmentation increase extinction thresholds? A European-wide test with seven forest birds. Global Ecology and Biogeography 22:1282–1292

Ruffell J, Banks-Leite C, Didham RK (2016) Accounting for the causal basis of collinearity when measuring the effects of habitat loss versus habitat fragmentation. Oikos 125:117-125.

Rybicki J, Abrego N, Ovaskainen O (2020) Habitat fragmentation and species diversity in competitive communities. Ecology Letters 23:506-517.

Rybicki J, Hanski I (2013) Species-area relationships and extinctions caused by habitat loss and fragmentation. Ecology Letters 16:27-38.

Santos JP, Freitas AVL, Brown KS, et al (2018) ATLANTIC BUTTERFLIES: a data set of fruit-feeding butterfly communities from the Atlantic Forests. Ecology 99:2875.

Saura S (2021) The habitat amount hypothesis implies negative effects of habitat fragmentation on species richness. Journal of Biogeography 48:11-22.

Scharnhorst VS, Ockermüller E, Becsi B, et al. (2026) Generalist 'Winners' replace specialist 'Losers': interactions between regional anthropogenic impacts and climate warming drive long-term wild bee declines. Diversity and Distributions 32:e70151.

Scheffers BR, et al. (2016) The broad footprint of climate change from genes to biomes to people. Science 354:aaf7671.

Schielzeth H (2010) Simple means to improve the interpretability of regression coefficients. Methods in Ecology and Evolution 1:103-113.

Seiferling I, Proulx R, Wirth C (2014) Disentangling the environmental-heterogeneity-species-diversity relationship along a gradient of human footprint. Ecology 95:2084-2095.

Semper-Pascual A, Burton C, Baumann M, et al. (2021) How do habitat amount and habitat fragmentation drive time-delayed responses of biodiversity to land-use change? Proceedings of the Royal Society B 288:20202466.

Senior RA, Hill JK, Edwards DP (2019) Global loss of climate connectivity in tropical forests. Nature Climate Change 9:623-626.

Senior RA, Hill JK, González del Pliego P, Goode LK, Edwards DP (2017) A pantropical analysis of the impacts of forest degradation and conversion on local temperature. Ecology and Evolution 7:7897-7908.

Sloan S, Jenkins CN, Joppa LN, et al. (2014) Remaining natural land cover in the global biodiversity hotspots. Biological Conservation 177:12-24.

Smith AC, Fahrig L, Francis CM (2011) Landscape size affects the relative importance of habitat amount, habitat fragmentation, and matrix quality on forest birds. Ecography 34:103–113

Smith AC, Koper N, Francis CM, Fahrig L (2009) Confronting collinearity: comparing methods for disentangling the effects of habitat loss and fragmentation. Landscape Ecology 24:1271-1285.

Smith C, Baker JCA, Spracklen DV (2023) Tropical deforestation causes large reductions in observed precipitation. Nature 615:270-274.

Sobral-Souza T, Bosco NS, Candelaria LP, et al. (2024) Spatial bias in sampling small rodents in the Atlantic Forest: a landscape and accessibility perspective. Perspectives in Ecology and Conservation 22:297-305.

Sobral-Souza T, Stropp J, Santos JP, et al. (2021a) Knowledge gaps hamper understanding the relationship between fragmentation and biodiversity. PeerJ 9:e11673

Socolar JB, Gilroy JJ, Kunin WE, Edwards DP (2016) How should beta-diversity inform biodiversity conservation? Trends in Ecology and Evolution 31:67-80.

Solórzano A, Brasil LSCA, Oliveira RR (2021) The Atlantic Forest ecological history: from pre-colonial times to the Anthropocene. In: Marques MCM, Grelle CEV (eds) The Atlantic Forest: history, biodiversity, threats and opportunities of the mega-diverse forest. Springer, Cham, pp 25-44.

Souza CM Jr, Shimbo JZ, Rosa MR, et al. (2020) Reconstructing three decades of land use and land cover changes in Brazilian biomes with Landsat archive and Earth Engine. Remote Sensing 12:2735.

Souza Y, Gonçalves F, Lautenschlager L, et al. (2019) ATLANTIC MAMMALS: a data set of assemblages of medium- and large-sized mammals of the Atlantic Forest of South America. Ecology 100:e02785.

Spake R, Barajas-Barbosa MP, Blowes SA, et al. (2022) Detecting thresholds of ecological change in the Anthropocene. Annual Review of Environment and Resources 47:797-821.

Spake R, Bowler DE, Callaghan CT, Blowes SA, Doncaster CP, Antao LH, Nakagawa S, McElreath R, Chase JM (2023) Understanding 'it depends' in ecology: a guide to hypothesising, visualising and interpreting statistical interactions. Biological Reviews 98:983–1002

Stark G, Ma L, Zeng Z-G, et al. (2023) Cool shade and not-so-cool shade: how habitat loss may accelerate thermal stress under current and future climate. Global Change Biology 29:6201-6216.

Stein A, Gerstner K, Kreft H (2014) Environmental heterogeneity as a universal driver of species richness across taxa, biomes and spatial scales. Ecology Letters 17:866-880.

Stein A, Kreft H (2015) Terminology and quantification of environmental heterogeneity in species-richness research. Biological Reviews 90:815-836.

Stoffel MA, Nakagawa S, Schielzeth H (2021) partR2: partitioning R2 in generalized linear mixed models. PeerJ 9:e11414.

Storch D, Bohdalková E, Okie J (2018) The more-individuals hypothesis revisited: the role of community abundance in species richness regulation and the productivity-diversity relationship. Ecology Letters 21:920-937.

Sunday JM, Bates AE, Kearney MR, et al. (2014) Thermal-safety margins and the necessity of thermoregulatory behavior across latitude and elevation. Proceedings of the National Academy of Sciences 111:5610-5615.

Sweeney CP, Jarzyna MA (2022) Assessing the synergistic effects of land use and climate change on terrestrial biodiversity: are generalists always the winners? Current Landscape Ecology Reports 7:41-48.

Swift TL, Hannon SJ (2010) Critical thresholds associated with habitat loss: a review of the concepts, evidence, and applications. Biological Reviews 85:35-53.

Tabarelli M, Aguiar AV, Ribeiro MC, Metzger JP, Peres CA (2010) Prospects for biodiversity conservation in the Atlantic Forest: lessons from aging human-modified landscapes. Biological Conservation 143:2328-2340.

Tamme R, Hiiesalu I, Laanisto L, Szava-Kovats R, Pärtel M (2010) Environmental heterogeneity, species diversity and co-existence at different spatial scales. Journal of Vegetation Science 21:796-801.

Taylor PD, Fahrig L, Henein K, Merriam G (1993) Connectivity is a vital element of landscape structure. Oikos 68:571-573.

Teng SN, Svenning J-C, Santana J, Reino L, Abades S, Xu C (2020) Linking landscape ecology and macroecology by scaling biodiversity in space and time. Current Landscape Ecology Reports 5:25-34.

Tews J, Brose U, Grimm V, et al. (2004) Animal species diversity driven by habitat heterogeneity/diversity: the importance of keystone structures. Journal of Biogeography 31:79-92.

Title PO, Bemmels JB (2018) ENVIREM: an expanded set of bioclimatic and topographic variables increases flexibility and improves performance of ecological niche modeling. Ecography 41:291-307.

Tourani M, Sollmann R, Kays R, et al. (2023) Maximum temperatures determine the habitat affiliations of North American mammals. Proceedings of the National Academy of Sciences 120:e2304411120.

Trabucco A, Zomer R (2019) Global High-Resolution Soil-Water Balance. figshare Dataset. https://doi.org/10.6084/m9.figshare.7707605.v3

Trabucco A, Zomer RJ (2010) Global Soil Water Balance Geospatial Database. CGIAR Consortium for Spatial Information. Published online, available from the CGIAR-CSI GeoPortal at https://csidotinfo.wordpress.com

Tredennick AT, Hooker G, Ellner SP, Adler PB (2021) A practical guide to selecting models for exploration, inference, and prediction in ecology. Ecology 102:e03336

Trew BT, Edwards DP, Lees AC, et al. (2024) Novel temperatures are already widespread beneath the world's tropical forest canopies. Nature Climate Change 14:753-759.

Trzcinski MK, Fahrig L, Merriam G (1999) Independent effects of forest cover and fragmentation on the distribution of forest breeding birds. Ecological Applications 9:586–593

Tscharntke T, Tylianakis JM, Rand TA, et al. (2012) Landscape moderation of biodiversity patterns and processes - eight hypotheses. Biological Reviews 87:661-685.

Tuanmu M-N, Jetz W (2014) A global 1-km consensus land-cover product for biodiversity and ecosystem modelling. Global Ecology and Biogeography 23:1031-1045.

Tuanmu M-N, Jetz W (2015) A global, remote sensing-based characterization of terrestrial habitat heterogeneity for biodiversity and ecosystem modelling. Global Ecology and Biogeography 24:1329-1339.

Tukiainen H, Bailey JJ, Field R, Kangas K, Hjort J (2017) Combining geodiversity with climate and topography to account for threatened species richness. Conservation Biology 31:364-375.

Turner MG (1989) Landscape ecology: the effect of pattern on process. Annual Review of Ecology and Systematics 20:171-197.

Valente JJ, Betts MG (2019) Response to fragmentation by avian communities is mediated by species traits. Diversity and Distributions 25:48-60.

Valente JJ, Gannon DG, Hightower J, Kim H, Leimberger KG, Macedo R, Rousseau JS, Weldy MJ, Zitomer RA, Fahrig L, Fletcher RJ, Wu J, Betts MG (2023) Toward conciliation in the habitat fragmentation and biodiversity debate. Landscape Ecology 38:2717-2730.

Valente JJ, Rivers JW, Yang Z, et al. (2023b) Fragmentation effects on an endangered species across a gradient from the interior to edge of its range. Conservation Biology 37:e14091

Vancine MH, Duarte K da S, de Souza YS, et al. (2018) ATLANTIC AMPHIBIANS: a data set of amphibian communities from the Atlantic Forests of South America. Ecology 99:1692.

Vancine MH, Muylaert RL, Niebuhr BB, et al. (2024) The Atlantic Forest of South America: spatiotemporal dynamics of the vegetation and implications for conservation. Biological Conservation 291:110499.

Vetter D, Hansbauer MM, Végvári Z, Storch I (2011) Predictors of forest fragmentation sensitivity in Neotropical vertebrates: a quantitative review. Ecography 34: 1-8.

Villard M-A, Metzger JP (2014) Beyond the fragmentation debate: a conceptual model to predict when habitat configuration really matters. Journal of Applied Ecology 51:309-318.

Wan J-Z, Wang C-J, Marquet PA (2023) Environmental heterogeneity as a driver of terrestrial biodiversity on a global scale. Progress in Physical Geography 47:912-930.

Wang X, Blanchet FG, Koper N (2014) Measuring habitat fragmentation: an evaluation of landscape pattern metrics. Methods in Ecology and Evolution 5: 634-646.

Watling JI, Arroyo-Rodríguez V, Pfeifer M, et al. (2020) Support for the habitat amount hypothesis from a global synthesis of species density studies. Ecology Letters 23:674-681.

Weeks TL, Betts MG, Pfeifer M, et al. (2023) Climate-driven variation in dispersal ability predicts responses to forest fragmentation in birds. Nature Ecology & Evolution 7:1079–1091

Whittaker RH (1972) Evolution and measurement of species diversity. Taxon 21:213–251.

Wiens JA (1989) Spatial scaling in ecology. Functional Ecology 3:385-397.

Williams JJ, Bates AE, Newbold T (2020) Human-dominated land uses favour species affiliated with more extreme climates, especially in the tropics. Ecography 43:391-405.

Williams JJ, Freeman R, Spooner F, et al. (2022) Vertebrate population trends are influenced by interactions between land use, climatic position, habitat loss and climate change. Global Change Biology 28:797-815.

Williams JJ, Newbold T (2020) Local climatic changes affect biodiversity responses to land use: a review. Diversity and Distributions 26:76-92.

Williams JJ, Newbold T (2021) Vertebrate responses to human land use are influenced by their proximity to climatic tolerance limits. Diversity and Distributions 27:1308-1323

Willmer JNG, Püttker T, Prevedello JA (2022) Global impacts of edge effects on species richness. Biological Conservation 272:109654.

With KA, Crist TO (1995) Critical thresholds in species' responses to landscape structure. Ecology 76:2446–2459

Wood SN (2004) Stable and efficient multiple smoothing parameter estimation for generalized additive models. Journal of the American Statistical Association 99:673-686.

Wood SN (2017) Generalized Additive Models: An Introduction with R, 2nd edn. CRC Press, Boca Raton FL.

Zeileis A, Hothorn T (2002) Diagnostic checking in regression relationships. R News 2(3):7-10.

Zhang H, Chase JM, Liao J (2024) Habitat amount modulates biodiversity responses to fragmentation. Nature Ecology and Evolution 8:1437-1447.

Zwiener VP, Padial AA, Marques MCM (2020) The mechanisms explaining tree species richness and composition are convergent in a megadiverse hotspot. Biodiversity and Conservation 29:799-815.

APPENDIX 0.S1. Species assemblages and spatial aggregation

This appendix documents the data foundation shared by the three Atlantic Forest empirical chapters — the sources and structure of the sampled species assemblages, their spatiotemporal distribution, the multi-level aggregation framework, the taxonomic universes underlying the response strata, and observed species richness across scales in the Atlantic Forest (AF). The General Introduction sets out the design in summary; this appendix reports it in full. Two classes of quantity appear here and are sourced differently. Upstream data-provenance quantities computed at the data-assembly stage — per-taxon and per-method sample sizes, gamma diversity, species-pool coverage, effort units, and year ranges — are reported from the curated data documentation. Quantities downstream of aggregation — assemblage-record and location counts, and richness summaries — are recomputed from the aggregated data and are not carried forward from earlier analyses.

0.S1.1 Sampled assemblages and data sources

We assembled 2,986 uniquely georeferenced species assemblages derived from standardised field protocols with reported survey effort. We restricted the compilation to assemblages sampled from 1980 onward, because earlier records are sparse and unevenly distributed among animal groups; the retained data span 1981 to 2018. We organised these by animal group (TAXA, n = 7), sampling method (METHOD, n = 16), and year (YEAR, n = 36); this structure reflects the way the primary datasets were originally compiled and published (Table 0.S1.1). The sampled assemblages span the Atlantic Forest in space and time (Figure 0.S1.1).

The published datasets combine coherent taxonomic groups at different hierarchical levels — bats (Chiroptera), birds (Aves), snakes (Serpentes), amphibians (Amphibia), and butterflies (Lepidoptera) — with morpho-ecological subsets defined by life history and sampling protocol (for example, mid–large versus small terrestrial mammals, and bat subsets). We refer to these seven dataset-defined animal groups as TAXA in the model structure, and we model TAXA = 7. Sampling heterogeneity arises both among taxa, through differing protocols and effort units (bird mist-netting in net-hours versus bat mist-netting in Straube–Bianconi units), and within taxa, when subsets of a community are targeted under a shared protocol (entire bat assemblages versus leaf-nosed phyllostomid subsets). This heterogeneity underlies the method-level sampling

structure. We excluded opportunistic and presence-only records, and for assemblages reported over multiple years without year-specific resolution we assigned the endpoint year for consistency.

The seven animal groups comprise the following.

- Mid- and large-bodied mammals (> 1 kg) — 397 assemblages, 1987–2018, comprising 211 from camera traps (camera-nights) and 186 from mixed active-search and line-transect methods (115 in survey-days, 71 in survey-kilometres). The mid–large mammal data draw on two Atlantic-series data papers (Lima et al. 2017; Souza et al. 2019).
- Small mammals — 418 assemblages, 1982–2015, comprising 324 from live traps, 66 from mixed methods, and 28 from pitfall traps, with effort in trap-nights (Bovendorp et al. 2017; Figueiredo et al. 2017).
- Bats — 174 assemblages, 1991–2016, sampled by mist nets in Straube–Bianconi units, comprising 49 that target leaf-nosed species only (Phyllostomidae) and 125 that target full bat assemblages (Muylaert et al. 2017).
- Passerines — 604 assemblages, 1981–2016, comprising 323 from point counts (number of points), 212 from mist nets (net-hours), and 69 from linear transects (hours) (Hasui et al. 2018).
- Snakes — 216 assemblages, 1984–2015, mixed methods standardised by number of sites surveyed. These are available as published richness estimates per assemblage rather than site-by-species matrices (Moura et al. 2017a, b).
- Anurans — 1,060 assemblages, 1982–2016, comprising 762 active, 182 mixed, and 116 passive, with effort in surveyed months (Vancine et al. 2018).
- Fruit-feeding butterflies — 117 assemblages, 1986–2018, mixed methods standardised by site counts (Santos et al. 2018).

These per-taxon counts sum to 2,986 curated assemblages across the Atlantic Forest.

0.S1.2 Gamma diversity and species-pool coverage

For mammals, birds, anurans, and butterflies we coded each dataset as a site-by-species binary matrix and derived observed richness as the number of distinct species recorded per sampling

unit. For snakes, available only as published richness estimates, we retained those values directly; where more than one snake assemblage fell within a single response unit we retained the maximum of their published estimates, because published richness totals cannot be combined as species lists. The observed taxon-level gamma diversity (γ) across the curated datasets comprises 70 mid–large mammals (77 including genus-level "sp." records), 116 small mammals, 97 bats, 428 passerines, 198 snakes, 495 anurans, and 278 butterflies. Because the seven groups are mutually exclusive, these values sum to 1,682 species across the compilation as a whole. These represent a substantial share of the AF species pool — approximately 71% of terrestrial mammals, 80% of bats, 74% of passerines, 71% of snakes, and 70% of anurans among tetrapods, and 60% of fruit-feeding butterflies. For taxa sampled by multiple methods, we also report method-level γ, which exposes differences in taxonomic coverage among protocols (Table 0.S1.1).

Species-pool quantities are reported as curated data-provenance summaries from the assembly stage. Because the downstream modelling datasets retain richness counts rather than full species-identity matrices for all taxa, these quantities are not recomputed from the aggregation pipeline. Snake richness in particular is analysis-derived (published estimates; Moura et al. 2017a, b) rather than matrix-derived, and we treat it as such throughout.

0.S1.3 Response strata and taxonomic universe

We classified each species by forest affiliation from species-level habitat records in the IUCN Red List of Threatened Species, coded under the IUCN Habitats Classification Scheme (Version 3.1; IUCN 2012). Following Donald et al. (2019), we treated species whose listed habitats fall entirely within level-1 class 1 (Forest) as forest specialists, species listed for forest together with one or more non-forest level-1 classes (e.g. savanna, grassland, wetland) as forest-facultative, and species with no forest habitat listed as non-forest. This classification allows us to construct response strata over two taxonomic universes. The full universe comprises all seven dataset-defined groups. The restricted universe comprises five habitat-classifiable groups and excludes butterflies and snakes for different reasons. Snakes are available only as analysis-derived richness estimates rather than species-level records, so individual species cannot be matched to IUCN habitat records. Fruit-feeding butterflies are resolved to species, but were not comprehensively assessed on the IUCN Red List, so no IUCN habitat records were available to

classify them. Neither taxon could therefore be assigned forest affiliation under the IUCN scheme. We therefore analyse five response strata across these two universes — total richness over the full universe (total_full), total richness over the restricted universe (total_restr), and the three mutually exclusive forest-affiliation guilds within the restricted universe, namely forest-specialist, forest-facultative, and non-forest richness. The restricted-total richness and the three guilds share an identical five-group universe and an identical set of assemblage records, and the guilds partition the restricted total exactly (total_restr = specialist + facultative + non-forest). We verified this construction from the data. For every restricted-universe assemblage record drawn from the five habitat-classifiable dataset-defined animal groups, restricted-total richness equals forest-specialist + forest-facultative + non-forest richness exactly. Rounding can affect only the displayed mean summaries in Table 0.S1.2, not the partition identity. The restricted-universe datasets contain the identical five-group set (mid–large mammals, small mammals, bats, passerines, anurans), and the two taxa dropped from the full universe are butterflies and snakes.

Across the five taxa carrying habitat-affiliation coding, we classified 1,199 species into forest specialists (302; 25.2%), forest-facultative species (702; 58.5%), and non-forest species (195; 16.3%). The specialist fraction ranges from 14.1% in anurans to 38.8% in small mammals (Figure 0.6). These proportions describe the classified species pool from which we construct the forest-specialist response; they are not model-weighted assemblage proportions. For four of the five taxa the classified pool matches the gamma diversity reported in Section S1.2 exactly, with all taxa resolved to species (anurans 495, bats 97, small mammals 116, passerines 428). Mid–large mammals are the exception. The classified pool holds 56 resolved species and seven genus-level records (63 in total), fewer than the 70 resolved species in the curated mid–large mammal gamma diversity (77 including the same genus-level records). We report the classified pool here because we construct the forest-specialist response from richness counts over this classified pool rather than from species identities, so its difference from the total observed gamma diversity does not affect the models.

0.S1.4 Multi-scale aggregation and sample-size attrition

We overlaid systematic hexagonal grids equivalent in area to circular buffers of 10, 25, and 40 km radius across the AF (Figure 0.4), defining the spatial units for aggregating assemblages by TAXA × METHOD at local, intermediate, and broader landscape contexts. The analyses

reported in the empirical chapters use year-aggregated responses; a year-retained variant was constructed at the 10-km grain for temporal sensitivity assessment only. We distinguish two counts, which are not interchangeable. Assemblage records — the distinct taxa × method × unit samples that constitute the modelling sample size — and locations — the distinct sampling coordinates, which are fewer because co-located samples (the same coordinates surveyed by different methods or for different taxa) are separate records but one location. At the 10-km grain, the full universe comprises 1,689 assemblage records across 1,651 locations; the restricted and specialist universes each comprise 1,360 records across 1,323 locations. Aggregation to coarser grains reduces both counts (full universe, 1,306 records at 25 km, 1,069 at 40 km), because coarser grids define fewer spatial units. The full universe carries more records and locations than the restricted universe at every grain, because restricting to the five classifiable taxa removes records — and sites sampled only for the excluded taxa. Hierarchical aggregation inherently produces sample-size attrition, from the curated datasets through the 10, 25, and 40 km grids (Figure 0.5 by animal group; Figure 0.S1.2 by sampling method).

Because taxa may sample distinct regions of the Atlantic Forest, and therefore distinct climatic space, we document per-taxon geographic coverage at contrasting levels of aggregation (Figure 0.S1.3). Point occurrences are shown for the curated data alongside the same records after aggregation to 40-km hexagons, with marginal density curves summarising the distribution of records along latitude and longitude independently rather than jointly. The corresponding climate-space overlap is reflected in the predictor distributions reported in the chapter-specific appendices.

0.S1.5 Richness summaries

We report observed richness summaries per taxon and per response stratum across aggregation scales. The richness distributions are right-skewed for every taxon — the mean exceeds the median in all cases — so we report the median alongside the mean, with the observed minimum and maximum, rather than the mean and standard deviation, which would misrepresent the skew and imply values below the observed floor. The taxa differ markedly in typical richness. At the 10-km grain in the full universe, median richness ranges from 7 in small mammals to 53 in butterflies. They differ also in tail length, which the median and mean separate from typical richness. Passerines carry the longest right tail (maximum richness 195, roughly 4.5 times their

median), whereas butterflies show the highest typical richness but a comparatively short tail (maximum 120, near 2.3 times their median).

Within the habitat-classifiable universe the guild partition is uneven. Forest-facultative richness carries most of the restricted total in every group, with medians of 30 in passerines, 13 in anurans, 9 in mid–large mammals, 7 in bats and 3 in small mammals. Forest-specialist richness is lower and more variable, from medians of 9 in passerines and 5 in bats down to 0 in anurans, though maxima reach 69 in passerines and 17 in both bats and small mammals. Non-forest richness is lowest throughout, with medians of 0 or 1 in every group and maxima of 23 in passerines, 19 in anurans, and 5 or fewer in the three remaining groups, so that stratum carries the least information and the widest relative uncertainty. Median richness rises modestly from the 10-km to the 40-km grain in every stratum, because coarser units consolidate more source assemblages.

0.S1.6 References

Bovendorp RS, Villar N, de Abreu-Junior EF, et al (2017) Atlantic small-mammal. A dataset of communities of rodents and marsupials of the Atlantic forests of South America. Ecology 98:2226.

Donald PF, Arendarczyk B, Spooner F, Buchanan GM (2019) Loss of forest intactness elevates global extinction risk in birds. Animal Conservation 22:341–347.

Figueiredo MSL, Barros CS, Delciellos AC, et al (2017) Abundance of small mammals in the Atlantic Forest (ASMAF). A data set for analyzing tropical community patterns. Ecology 98:2981.

Hasui E, Metzger JP, Pimentel RG, et al (2018) ATLANTIC BIRDS. A data set of bird species from the Brazilian Atlantic Forest. Ecology 99:497.

IUCN (2012) IUCN Habitats Classification Scheme (Version 3.1). International Union for Conservation of Nature, Gland, Switzerland and Cambridge, UK.

Lima F, Beca G, Muylaert RL, et al (2017) ATLANTIC-CAMTRAPS. A dataset of medium and large terrestrial mammal communities in the Atlantic Forest of South America. Ecology 98:2979.

Moura MR, Costa HC, Argôlo AJS, et al (2017a) Environmental constraints on the compositional and phylogenetic beta-diversity of tropical forest snake assemblages. Journal of Animal Ecology 86:1192–1204.

Moura MR, Argôlo AJ, Costa HC (2017b) Historical and contemporary correlates of snake biogeographical subregions in the Atlantic Forest hotspot. Journal of Biogeography 44:640–650.

Muylaert RDL, Stevens RD, Esbérard CEL, et al (2017) ATLANTIC BATS. A data set of bat communities from the Atlantic Forests of South America. Ecology 98:3227.

Santos JP, Freitas AVL, Brown KS, et al (2018) ATLANTIC BUTTERFLIES. A data set of fruit-feeding butterfly communities from the Atlantic Forests. Ecology 99:2875.

Souza Y, Gonçalves F, Lautenschlager L, et al (2019) ATLANTIC MAMMALS. A data set of assemblages of medium- and large-sized mammals of the Atlantic Forest of South America. Ecology 100:e02785.

Vancine MH, Duarte K da S, de Souza YS, et al (2018) ATLANTIC AMPHIBIANS. A data set of amphibian communities from the Atlantic Forests of South America. Ecology 99:1692.

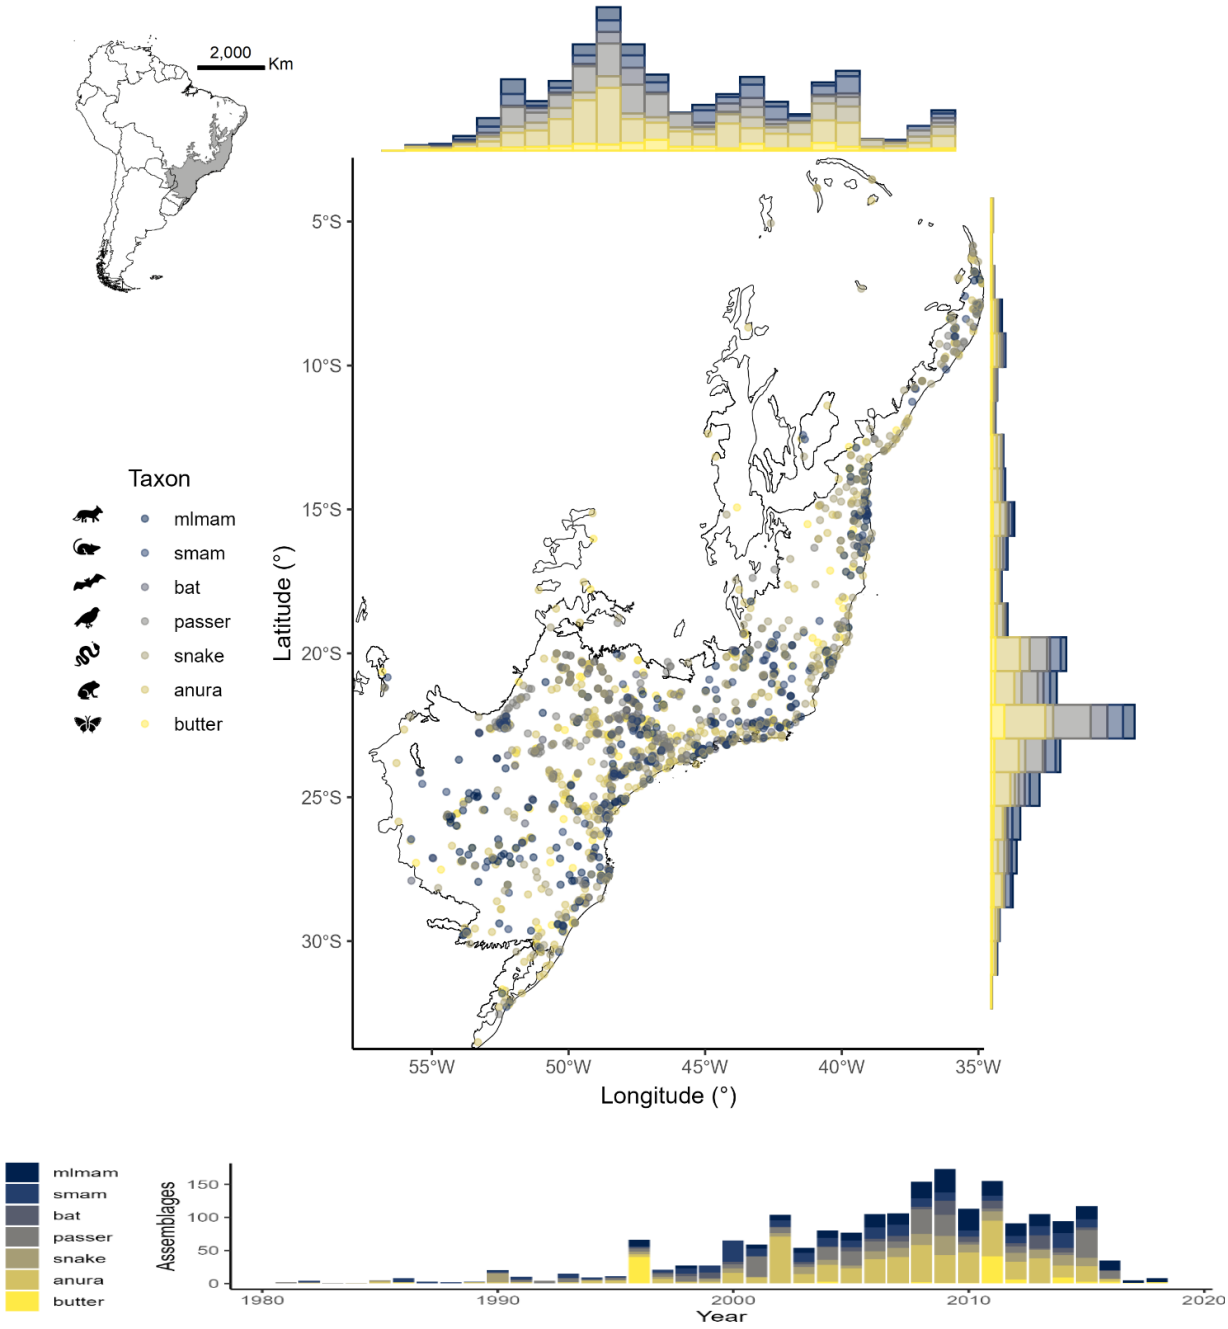


Figure 0.S1.1. Geographic and temporal distribution of curated species assemblages across the Atlantic Forest. Points show the locations of all 2,986 curated assemblages before spatial aggregation, with marginal histograms giving sampling frequency across latitude and longitude (top panel), and the yearly distribution of sampling by animal group between 1981 and 2018 (bottom panel). Sampling is concentrated in the centre and south of the biome, and no single animal group spans the full period, so the compilation describes a realised sampling domain rather than a spatially or temporally balanced census. Figure 0.3 shows the records separated by animal group.

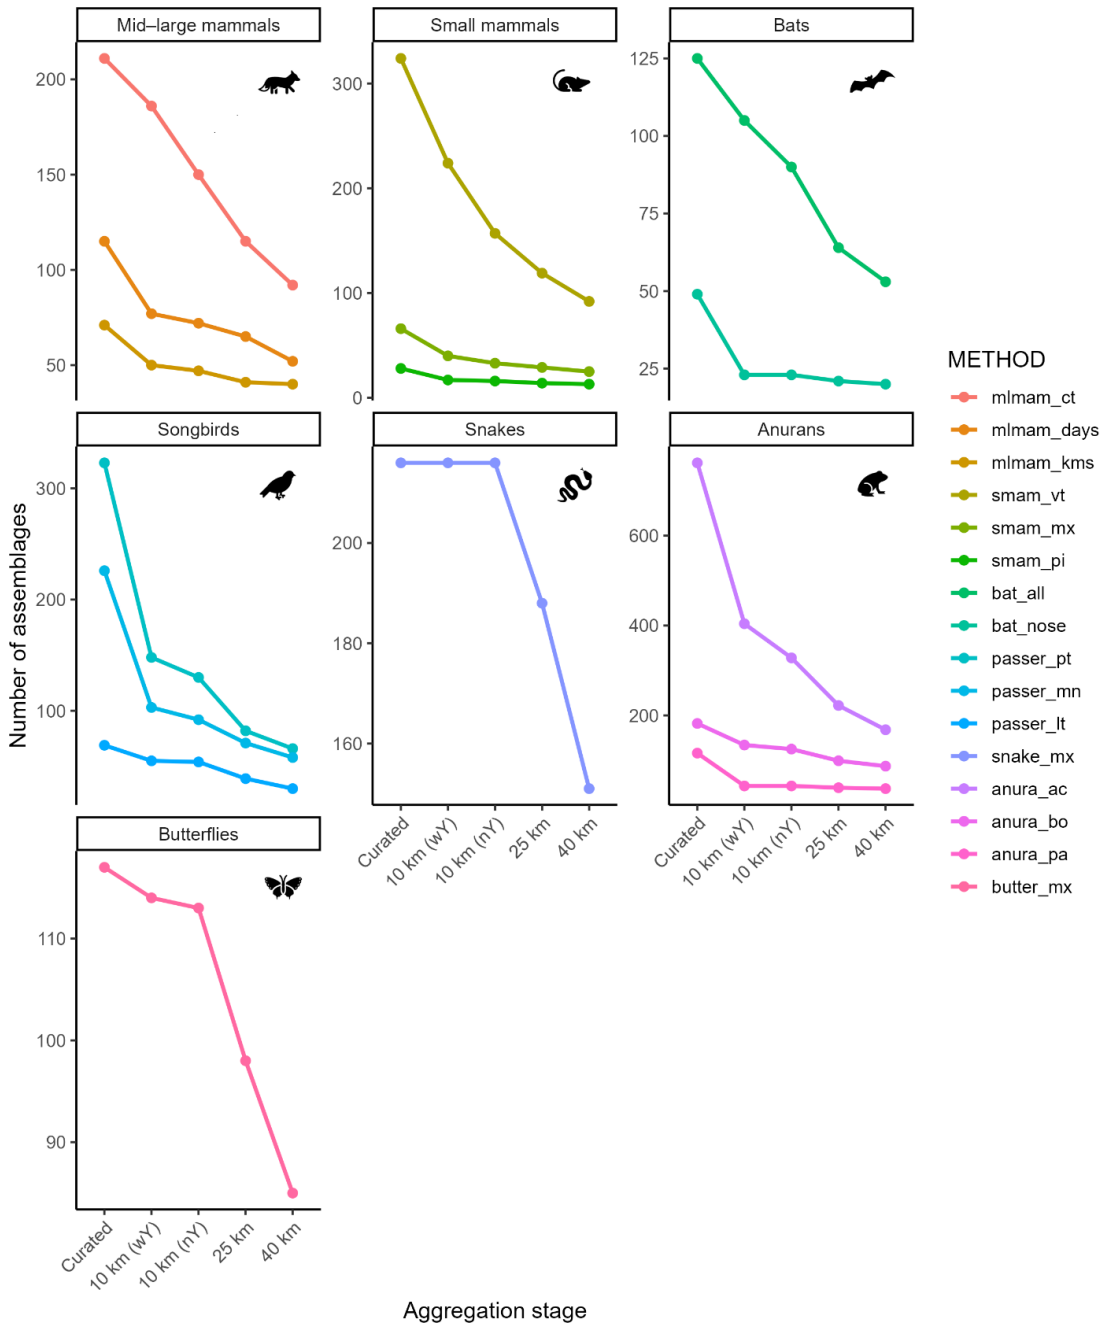


Figure 0.S1.2. Sample-size attrition by sampling method across aggregation stages. Number of assemblages retained at each stage, from the fully disaggregated curated data through the 10 km grain to the broadest 40 km aggregation. Stages labelled 10 km (wY) and 10 km (nY) distinguish the year-retained and year-aggregated variants; the analyses reported in this dissertation use the year-aggregated form throughout. Attrition is steepest for methods sampled densely relative to the grain, because more source assemblages consolidate into each modelling record. Figure 0.5 shows the same attrition summarised by animal group.

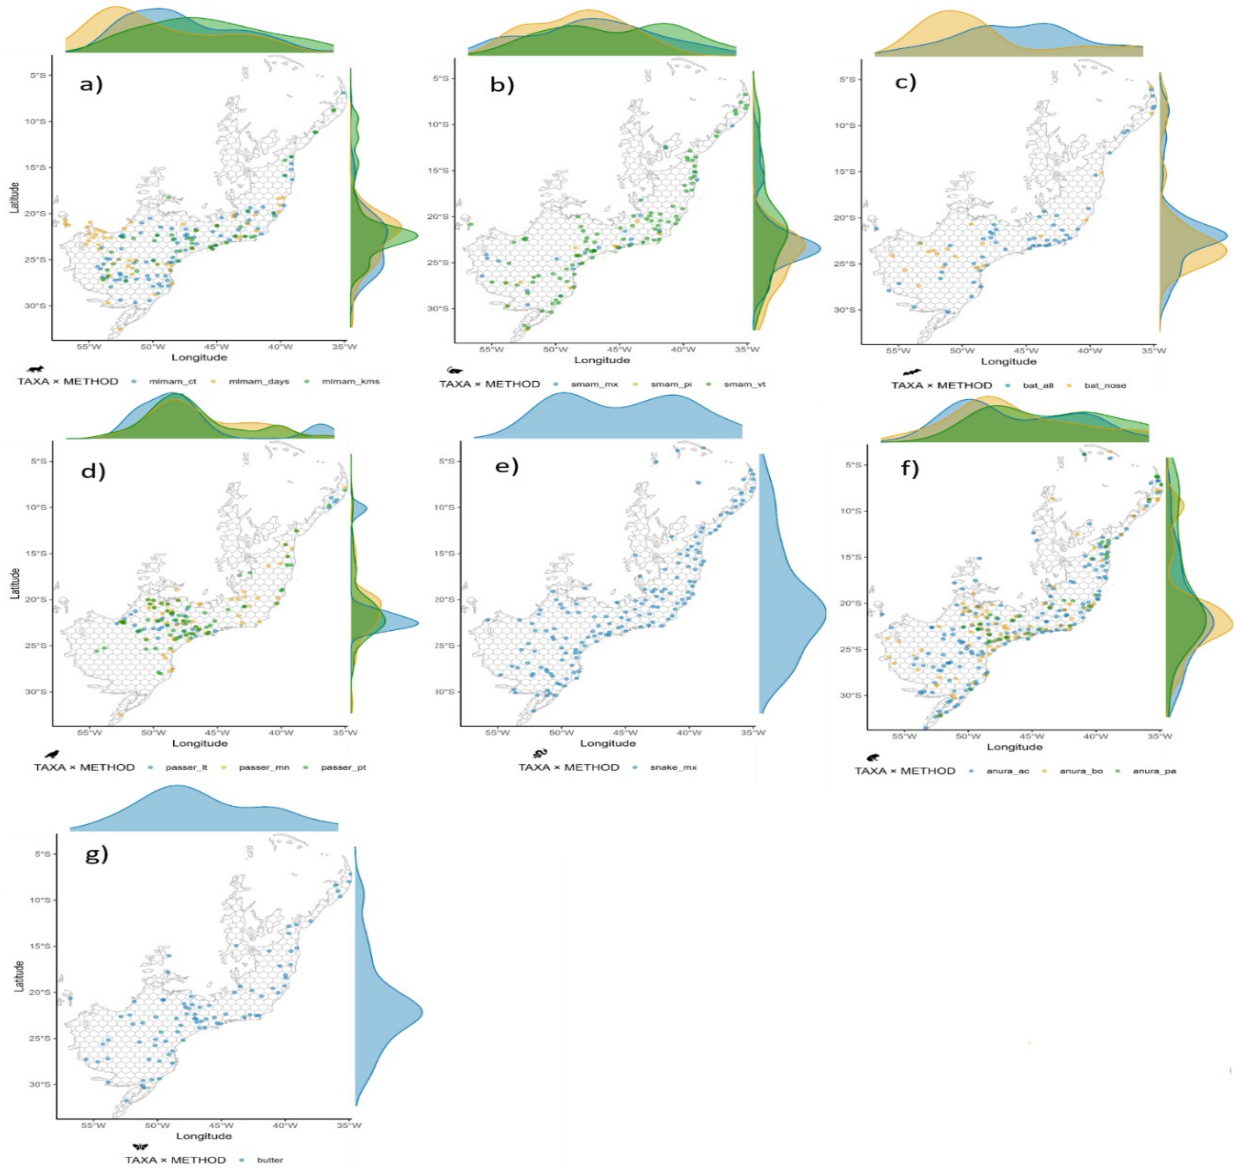


Figure 0.S1.3. Spatial distribution of curated species assemblages after aggregation to 40-km hexagonal response units. Panels show mid–large mammals (a), small mammals (b), bats (c), passerines (d), snakes (e), anurans (f) and fruit-feeding butterflies (g), with points coloured by sampling method. Marginal density curves summarise the latitudinal and longitudinal distributions of records for each method independently, not jointly. Comparison with the unaggregated distributions in Figure 0.3 shows the spatial consequence of aggregation: records consolidate within occupied hexagons and coverage contracts toward the densely sampled centre and south of the biome, while the broad latitudinal span is retained.

Table 0.S1.1. Curated assemblages by animal group and sampling method. Number of assemblages (n) for each TAXA × METHOD combination before spatial aggregation, with the effort units reported by the source dataset, the temporal range of sampling, and the data source. Taxon-level and method-level gamma diversity (γ) give the total number of species recorded within each subset. Effort units are not comparable across methods, which is why sampling method is retained as a second organising level alongside animal group. The seven animal groups are mutually exclusive, so taxon-level values sum to 1,682 species across the compilation as a whole.

Animal group	TAXA (γ_T)	Sampling protocol/scope	Effort unit(s)	METHOD (γ_M)	Samples (n)	YEAR range
Mid–large mammals	mlmam (70)	Camera traps	Camera-nights	mlmam_ct (48)	211	2002–2018
		Mixed	Days surveyed	mlmam_days (58)	115	1996–2018
		Mixed	Km surveyed	mlmam_kms (54)	71	1987–2017
Small mammals	smam (116)	Live traps	Trap-nights	smam_vt (102)	324	1982–2015
		Mixed	Trap-nights	smam_mx (66)	66	1987–2015
		Pitfall traps	Trap-nights	smam_pi (56)	28	2005–2015
Bats bat (97)		Mist nets (all)	Straube–Bianconi	bat_all (93)	125	1991–2016
		Mist nets (guild scope)	Straube–Bianconi	bat_nose (51)	49	1997–2013
Songbirds	passer (428)	Point counts	Number of points	passer_pt (363)	323	1991–2015
		Mist nets	Net-hours	passer_mn (299)	212	1981–2013
		Line transects	Hours surveyed	passer_lt (314)	69	2001–2016
Snakes	snake (198)	Mixed	Site-based count	snake_mx (198)	216	1984–2015
Anurans	anura (495)	Active	Months surveyed	anura_ac (462)	762	1982–2016
		Mixed (active and passive)	Months surveyed	anura_bo (356)	182	2000–2015
		Passive	Months surveyed	anura_pa (151)	116	2000–2016
Butterflies	butter (278)	Mixed methods	Site-based count	butter_mx (278)	117	1986–2018

Table 0.S1.2. Observed species richness across aggregation scales by animal group and sampling method. Minimum, maximum, mean (α) and standard deviation of observed richness for each TAXA × METHOD combination at each aggregation level, with the number of modelling records (n) at that level. Levels labelled 10 km (wY) and 10 km (nY) distinguish the year-retained and year-aggregated variants at the 10 km grain; the analyses reported in this dissertation use the year-aggregated form throughout. Taxon-level and method-level gamma diversity (γ) give the total number of species recorded within each subset. Record counts decline across levels because coarser grids define fewer spatial units, while mean richness rises as more source assemblages consolidate into each record.

				Species richness			
TAXA (γ)	METHOD (γ)	Level	Samples (n)	Min	Max	Mean	SD
Mid-large mammals (70)	mlmam_ct (48)	10 km (wY)	186	1	30	10.8	5.8
		10 km (nY)	150	1	30	11.1	6.3
		25 km	115	1	33	12.2	6.8
		40 km	92	1	33	13.4	7.2
	mlmam_days (58)	10 km (wY)	77	1	29	12.5	7.1
		10 km (nY)	72	1	29	13.2	7.0
		25 km	65	1	29	13.8	7.2
		40 km	52	1	29	15.5	7.3
	mlmam_kms (54)	10 km (wY)	50	1	24	11.1	5.9
		10 km (nY)	47	1	24	11.4	6.1
		25 km	41	1	24	12.1	6.5
		40 km	40	1	24	12.4	6.7
Small mammals (116)	smam_vt (102)	10 km (wY)	224	1	22	6.6	4.1
		10 km (nY)	157	1	23	7.7	5.0
		25 km	119	1	31	8.4	5.8
		40 km	92	1	29	9.9	6.5
	smam_mx (66)	10 km (wY)	40	1	24	9.4	4.9
		10 km (nY)	33	1	25	9.6	5.4
		25 km	29	1	25	10.0	5.5
		40 km	25	3	24	11.4	5.9
	smam_pi (56)	10 km (wY)	17	3	22	10.9	5.4
		10 km (nY)	16	3	25	10.9	5.9
		25 km	14	3	25	11.1	6.8
		40 km	13	3	25	11.8	6.6
Bats (97)	bat_all (93)	10 km (wY)	105	1	53	14.2	8.5
		10 km (nY)	90	1	53	15.1	8.8
		25 km	64	1	53	17.0	10.0
		40 km	53	1	53	18.3	10.6
		10 km (wY)	23	4	20	11.5	4.6

	bat_nose (51)	10 km (nY)	23	4	20	11.5	4.6
		25 km	21	4	20	11.9	4.9
		40 km	20	4	26	11.8	5.5
Songbirds (428)	passer_pt (363)	10 km (wY)	148	8	116	53.3	24.5
		10 km (nY)	130	8	129	54.9	26.8
		25 km	82	13	156	64.9	31.5
		40 km	66	14	159	69.6	32.4
	passer_mn (299)	10 km (wY)	103	1	85	27.3	18.0
		10 km (nY)	92	1	85	29.9	18.8
		25 km	71	1	87	33.6	20.6
		40 km	58	1	98	38.4	22.5
	passer_lt (314)	10 km (wY)	55	11	195	51.3	30.6
		10 km (nY)	54	11	195	51.5	31.0
		25 km	39	19	212	57.0	35.3
		40 km	30	19	212	64.2	39.0
Snakes (198)	snake_mx (198)	10 km (wY)	216	5	47	16.7	9.3
		10 km (nY)	216	5	47	16.6	9.3
		25 km	188	5	47	17.4	9.5
		40 km	151	5	47	17.9	9.8
Anurans (495)	anura_ac (462)	10 km (wY)	404	1	48	18.5	9.3
		10 km (nY)	328	1	61	20.0	10.6
		25 km	222	1	70	23.5	13.1
		40 km	168	1	71	26.4	14.9
	anura_bo (356)	10 km (wY)	134	5	67	21.4	10.0
		10 km (nY)	125	6	68	22.3	10.0
		25 km	99	8	68	24.6	11.2
		40 km	87	8	68	26.3	11.3
	anura_pa (151)	10 km (wY)	43	2	35	11.0	6.8
		10 km (nY)	43	2	35	11.0	6.8
		25 km	39	2	35	11.3	7.4
		40 km	37	2	39	11.9	7.8
Butterflies (278)	butter_mx (278)	10 km (wY)	114	15	120	58.7	25.9
		10 km (nY)	113	15	120	58.7	26.2
		25 km	98	15	120	61.8	27.9
		40 km	85	22	121	65.7	28.5

APPENDIX 0.S2. Thesis-wide glossary of recurring terms

This appendix defines terms that recur across chapters. Every entry either appears repeatedly in the thesis as working vocabulary, or is used sparingly but resolves a distinction that matters across chapters. Terminology specific to the diagnostic framework of Chapter 3 is defined in Appendix 3.S1, and predictor abbreviations are listed with each chapter's figures and tables. The definitions are reading aids and do not replace the fuller descriptions in the General Introduction and the chapter Methods.

***The biodiversity data and its spatial structure*.** These terms distinguish what was observed in the field from what enters a model, and the two spatial scales at which each analysis operates.
Aggregation
Used in two senses. Spatially, aggregation is the consolidation of source assemblages into hexagonal response units at a given grain, which is what turns assemblages into modelling records. Taxonomically, aggregation is the pooling of forest-affiliation guilds into restricted total richness, or of all animal groups into full-universe richness. The two are independent. A response can be spatially aggregated at 10 km while remaining guild-resolved. Where the text refers to what aggregation conceals, the taxonomic sense is meant.
Aggregation cancellation
The loss or attenuation of a guild-level signal when guilds are pooled and their contributions have opposing directions. Distinguished from dilution, where a response concentrated strongly in one guild becomes weaker in aggregate richness without requiring opposing signs. Both are ways in which aggregate richness can obscure guild-level response structure.
Full universe
All seven dataset-defined animal groups and all 2,986 assemblages; the basis for full-universe total richness.
Habitat-classifiable universe

The five vertebrate groups for which species-level IUCN habitat records permit forest-affiliation coding, and the 2,653 assemblages they contribute. A subset of the full universe, not a separate compilation.
Modelling record
The unit of analysis. Source assemblages of the same animal group, sampled by the same method, that fall within the same spatial unit are consolidated into one record. A modelling record therefore represents an animal group by sampling method by spatial unit combination rather than an individual field site.
Predictor extent
The spatial scale at which the surrounding landscape is described: concentric extents of 10 to 100 km radius, centred on the centroid of the assemblages aggregated into each record. Changing the predictor extent changes the environmental context attached to a response, not the response itself.
Response grain
The spatial scale at which species are counted: hexagonal units equivalent in area to circles of 10, 25 or 40 km radius. Changing the response grain changes which biodiversity observation is constructed, and therefore what quantity is being estimated.
Response stratum
One richness response fitted separately within the shared analysis, meaning full-universe total, restricted total, forest-specialist, forest-facultative or non-forest richness. The strata derive from overlapping data and should not be treated as independent samples.
Sampled location
A distinct set of sampling coordinates. Locations are fewer than modelling records, because one location surveyed for more than one animal group, or by more than one method, contributes more than one record.
Scale coherence

Agreement of an effect across the eighteen grain-by-extent combinations. Because these are nested, overlapping windows on the same assemblages, recurrence across them indicates coherence across scales rather than independent replication.
Scale matrix
The set of valid response-grain by predictor-extent combinations evaluated in a chapter, eighteen in the Atlantic Forest analyses, since predictor extent is never smaller than response grain. The cells overlap in data and should not be read as independent replicates. Recurrence across them indicates scale coherence rather than repeated confirmation.
Source assemblage
A single curated record of the species observed at one location, for one animal group, under one standardised sampling protocol with reported survey effort. The compilation holds 2,986 of them. A source assemblage is not a modelling observation.
Specificity diagnostic
A response stratum used as an interpretive comparison to assess whether a pattern is specific to the mechanism or response group proposed for it. Non-forest richness serves this role: it is not a negative control, because land use can affect non-forest species, and the expectation is different sensitivity rather than an absent effect.
Land cover, habitat and configuration. These terms separate what a landscape contains from how it is arranged, and both from the biological concept neither measures directly.
Configuration
The spatial arrangement of land cover — subdivision, isolation, edge exposure, connectivity — as distinct from how much of it there is. Composition and configuration are separable in principle but often covary in real landscapes, because land-cover transformation can change both how much of a focal cover type remains and how it is spatially arranged.
Environmental filtering

The process by which environmental conditions exclude species whose requirements those conditions do not meet, so the assemblage present at a location is a subset of the regional pool. Land-use transformation is expected to filter directionally. Species differ in forest dependence, so conversion removes forest-associated species before matrix-tolerant ones. A filtering interpretation requires evidence that the loss is ordered by species attributes, not simply that richness declined.
Fragmentation per se
The effect of configuration on biodiversity at a held-comparable habitat amount: the same amount of habitat, differently arranged.
Guild
Used here for the three mutually exclusive forest-affiliation classes, forest-specialist, forest-facultative and non-forest, assigned from species-level habitat information. These are habitat-affiliation classes rather than trophic or functional guilds. The classification describes affiliation with listed forest and non-forest habitat types, and stops short of demonstrating physiological habitat obligacy.
Habitat
Environmental conditions meeting the biological requirements of a focal taxon or response group. Habitat is taxon-specific. Taxon-specific habitat availability or suitability is not measured directly for the landscapes analysed here. Where an analysis requires habitat amount explicitly, a land-cover class is used as its proxy, so the resulting inference is conditional on that proxy rather than on measured habitat suitability.
Habitat amount
How much habitat a landscape contains. In the amount-conditioned fragmentation analysis it is operationalised as forest cover within a fixed extent, taking forest as the land-cover proxy for the habitat of forest-associated taxa.
Habitat obligacy

The degree to which a species genuinely requires a habitat type to persist, as distinct from being recorded there. The forest-affiliation classes used in this dissertation derive from listed habitat associations in the IUCN Red List and describe exclusivity among listed habitat types. A species classed as forest-specialist is listed only for forest, which is not the same as demonstrated physiological or demographic dependence on forest conditions. The classification is a repeatable assignment rule rather than experimental evidence of requirement.
Land-cover composition
How much of each land-cover type a landscape contains, irrespective of how it is arranged. Cropland prevalence and non-natural cover represent land-cover composition in the climate and heterogeneity analyses; forest cover represents it in the fragmentation analysis. Not to be confused with assemblage composition, which refers to which species are present.
Land-use surrogate
A mapped composition variable standing in for land-use transformation. Cropland prevalence and non-natural cover are used here. They index landscape composition along a transformation gradient, not management intensity, and are not interchangeable: they differ in what they include and in their fitted functional form.
Matrix
In landscape-ecological usage, the non-focal land cover surrounding habitat remnants, whose permeability and contrast can affect organisms. In expressions such as *scale matrix* or *coefficient matrix*, matrix instead means an array of analytical combinations or results. The two uses are unrelated.
From question to inference. These terms name the steps by which an ecological question becomes an interpreted result. The distinctions among them are what allow a fitted coefficient to be assessed against the quantity it was meant to estimate.
Additive and interaction models
An additive model enters predictors as separate terms and, on the model scale, holds the effect of one constant across values of the other. An interaction model adds a product term, letting the fitted

effect of one vary along the other. The distinction is inferential as well as statistical. An additive fragmentation model imposes a constant configuration effect across the habitat-amount gradient, while an interaction model allows it to change, and that is what makes the competing frameworks distinguishable.
Biodiversity response
The biological quantity measured in an analysis. Species richness is the principal response in the Atlantic Forest empirical analyses, while other chapters also consider distinct diversity quantities. The response must be distinguished from the environmental transformation or predictor used to explain it.
Concordance
Agreement in sign between two quantities that should agree if a coefficient is geometrically stable. Structure-coefficient concordance compares a term's coefficient with its structure coefficient and is part of the primary support criterion. Raw–residualised concordance compares the two fragmentation formulations and is a robustness axis, not a support criterion. Sign concordance across scale cells describes whether a supported effect holds direction across the scale matrix.
Conditional
Used in several senses. In the interaction framework, a conditional effect is an effect evaluated at a specified value of another predictor, and a conditional response surface is the set of those effects across the moderating gradient. In mixed-model variance reporting, conditional R^2 describes variance accounted for by fixed and random effects together. In a conditional Poisson GLMM, conditional refers to the response distribution given the random effects. In conditional prediction, it refers to predictions that incorporate specified random effects. Elsewhere it keeps its ordinary meaning of *given*. The interaction sense carries the ecological conditionality argument of this dissertation and is distinct from the mixed-model uses.
Ecological conditionality

Dependence of a biodiversity response on ecological context. It covers where, under which environmental conditions, at what scale, and for which assemblage components the effect of land use or configuration changes. Chapters 1, 2 and 5 examine different forms of it.
Effect and association
In the observational analyses an effect is a fitted model contrast, the change in fitted response associated with a predictor under the stated model and conditioning set. It carries no causal claim. *Association* is used where that limitation is being emphasised. The dissertation distinguishes throughout between a pattern consistent with an ecological mechanism and evidence that identifies that mechanism.
Estimand
The ecological quantity or contrast that the question requires, made precise by the design. Fragmentation per se is an estimand: the effect of configuration at a given habitat amount. The estimand must be distinguished from what a fitted specification actually defines — an additive amount-plus-configuration model implies a configuration effect constant along the amount gradient, whereas an interaction model implies one that varies with it — and from the value the data return for it.
Estimate
The value the realised data return for the quantity a fitted model defines, with its uncertainty. An estimate can be precise and still fail to estimate the intended estimand if the realised predictor space does not contain the required contrast, if the response does not represent the ecological quantity being interpreted, or if the metric does not represent the proposed mechanism.
Fixed and random effects
Fixed effects are the systematic predictor terms whose coefficients carry the focal relationships. Random effects allow specified groups to depart from that common structure. A random intercept lets group baselines differ, and a random slope lets the relationship with a predictor differ among groups. In Chapters 1, 2 and 5, TAXA and METHOD represent animal-group and sampling-method grouping structure. The observation-level random effect serves a different purpose. It gives

each modelling observation its own intercept to absorb extra-Poisson variation, and represents no biological grouping.
Inferential conditionality
Dependence of an ecological interpretation on the conditions under which the estimate was obtained, namely realised predictor geometry, response–estimand alignment, metric–mechanism correspondence and domain stability. An estimate can be statistically well defined while the interpretation assigned to it remains conditional on these properties of the data and design.
Main effect in an interaction model
The coefficient of one interacting predictor when the other equals its reference value. Because the focal continuous predictors are standardised, a land-use or configuration main effect is the fitted effect at the **mean** of the moderator, not an average effect across the moderator's observed range.
Marginal
Used in three distinct senses. In response-surface analyses, a marginal effect or marginal slope is the local change in the fitted response with respect to a focal predictor, evaluated at stated values of the others. In an interaction model it varies along the moderator gradient, as $\beta_1 + \beta_3 A$. Its scale depends on the analysis. Some Chapter 1 displays evaluate marginal slopes on the response scale, while Chapters 4 and 5 define the fragmentation marginal effect on the model scale and show response-scale predictions separately. In coefficient diagnostics, a marginal association is the unadjusted relationship between one predictor and the response, before the others enter. In mixed-model variance reporting, marginal R^2 is the share of modelled variance attributed to the fixed effects. These meanings are not interchangeable.
Model scale and response scale
The model scale is where a generalized model's linear predictor and coefficients are defined, which for the Poisson log-link models used here is the log scale. The response scale is expected species richness after the inverse link. A coefficient or marginal slope on the model scale and a predicted richness contrast on the response scale describe the same fitted model but are not numerically interchangeable.

Moderator
A predictor whose value changes the size or direction of another predictor's relationship with the response, fitted as an interaction term. Climate, heterogeneity, productivity and forest amount are moderators of land-use or configuration effects in this dissertation. A control and a moderator do different work. A control enters additively, to remove a competing explanation. A moderator enters multiplicatively, because the question is whether the focal effect depends on it.
Partial coefficient
A predictor's coefficient after accounting for the other terms in the model, representing the fitted change associated with that predictor while the others are held fixed. Its ecological interpretation depends on the realised predictor geometry. An estimable partial coefficient does not guarantee that the held-fixed comparison is well represented by sampled landscapes.
Predictor geometry
The realised joint structure of the predictors across the sampled landscapes. It includes the ranges sampled, how they overlap, which combinations occur and which are rare or absent. Geometry is a property of the data, fixed before any model is fitted, and it determines which comparisons the data contain. Geometry determines how directly the sampled data support the comparison represented by a coefficient. Two predictors moderately correlated pairwise may still share a dominant gradient that leaves the required contrast unsampled, which is why conventional collinearity screens are necessary but not sufficient.
Predictor standardisation
In the Atlantic Forest empirical analyses of Chapters 1, 2 and 5, focal continuous predictors are z-transformed before fitting, so coefficients are expressed per one standard deviation of the corresponding predictor rather than per natural unit. Standardisation also places zero at the predictor mean. In an interaction between standardised predictors, the main effect of one is evaluated at the mean of the other.
Primary support

The evidence criterion used for focal interaction results in Chapters 1, 2 and 5. An interaction is primary-supported in a given cell when it improves fit over its matched reduced or additive comparator by more than two AICc units **and** the product term's structure coefficient is concordant in sign with its coefficient. Both conditions must hold. Primary support is distinct from a coefficient being non-zero, a p-value below 0.05, or an interval excluding zero. Those quantities are retained and reported separately where available. Support is evaluated within each formulation and never pooled across them.
Residual spatial autocovariate
A covariate derived from the spatial pattern remaining in the residuals of the corresponding non-spatial model, then entered as a fixed effect in the fitted model to account for broad-scale residual spatial dependence. It is a defensive statistical adjustment. It does not stand for an ecological predictor or a measured process, and it offers no guarantee that spatial confounding has been removed. The geographic smooth used to derive it is fitted to residuals after the environmental model, and not estimated jointly with the focal environmental gradients. The resulting fitted values are then included alongside those predictors in the final model.
Residualisation
Removing from one predictor the component predicted by another and fitting the residual in place of the original. In Chapter 3 it is used diagnostically, to examine how coefficient allocation changes as shared predictor variation is reassigned. In Chapter 5 it defines a different operational estimand, the effect of more-than-expected fragmentation at a given forest amount after nonlinear amount control, and is carried as a complementary robustness formulation. Residualisation does not create the same-amount, different-configuration contrasts that causal identification would require.
Response surface
The fitted relationship between a biodiversity response and a focal predictor across the range of one or more other predictors. In an interaction model the slope of the focal predictor can change along the moderator gradient, producing a conditional response surface. An additive model instead constrains that slope to be constant across the gradient. Competing fragmentation frameworks

predict different response-surface geometries along the habitat-amount gradient, which is what makes them empirically distinguishable.
Structure coefficient
The correlation between a predictor and the model's fitted linear predictor. It describes how strongly and in what direction a predictor aligns with the fitted response gradient, regardless of how much unique variance its partial coefficient retains once correlated predictors enter. Used alongside partial coefficients and marginal associations to diagnose coefficient distortion under predictor coupling.
Suppression
A coefficient geometry in which a predictor tracks the fitted response closely but contributes little unique variance once a correlated predictor is included, so its partial coefficient collapses toward zero or reverses sign while its structure coefficient stays high. Suppression is a property of the estimation rather than the ecology. It signals that a partial coefficient has been redistributed to a correlated partner, and that a near-zero coefficient cannot be read as absence of association.

APPENDIX 1.S1. Chapter-specific response documentation

This chapter uses the shared Atlantic Forest assemblage compilation, forest-affiliation classification and multiscale response design described in the General Introduction and Appendix 0.S1. Here we document only chapter-specific response construction and diagnostics.

1.S1.1 Total-versus-restricted-universe benchmark

We assess whether excluding butterflies and snakes alters the total-richness signal by comparing univariate land-use and climate coefficients between the full and restricted universes (Table 1.S1.1). The coefficients agree in sign at 99.1% of predictor–scale combinations (107 of 108) and differ by a mean absolute value of 0.011, so the univariate total-richness signal does not depend on the two taxa lacking habitat-affiliation data. The single sign disagreement involves annual temperature range, whose coefficients in the restricted universe are numerically near zero (eight of its combinations fall below the near-zero threshold), so the disagreement is uninformative rather than substantive. We report this univariate comparison as a pre-control subset-invariance baseline; the restricted total and the three forest-affiliation guilds are all fitted as interaction strata in the main analysis. The paired comparison of the adjusted thermal-moderation pattern between the full and restricted universes is reported with the main results.

Table 1.S1.1. Univariate total-versus-restricted-universe benchmark by predictor and scale.

Predictor	Response grain	Full universe median	Restricted universe median	Median shift	Full universe IQR	Restricted universe IQR
CROP	10 km	-0.011	-0.022	-0.010	1.574	1.572
CROP	25 km	-0.009	-0.013	-0.004	1.524	1.520
CROP	40 km	-0.027	-0.027	-0.000	1.441	1.445
NNCOV	10 km	+0.205	+0.204	-0.001	1.731	1.752
NNCOV	25 km	+0.194	+0.174	-0.020	1.586	1.619
NNCOV	40 km	+0.160	+0.162	+0.002	1.648	1.675

Predictor	Response grain	Full universe median	Restricted universe median	Median shift	Full universe IQR	Restricted universe IQR
MAXT	10 km	+0.057	+0.065	+0.008	1.557	1.561
MAXT	25 km	+0.032	+0.042	+0.010	1.520	1.510
MAXT	40 km	+0.059	+0.058	-0.001	1.538	1.544
ATR	10 km	+0.309	+0.293	-0.016	1.121	0.991
ATR	25 km	+0.300	+0.291	-0.009	1.124	1.016
ATR	40 km	+0.279	+0.261	-0.018	1.099	1.012
DTR	10 km	+0.097	+0.168	+0.070	1.702	1.661
DTR	25 km	+0.161	+0.248	+0.087	1.680	1.690
DTR	40 km	+0.146	+0.242	+0.096	1.713	1.682
PDRY	10 km	+0.285	+0.291	+0.006	1.625	1.623
PDRY	25 km	+0.287	+0.291	+0.004	1.629	1.628
PDRY	40 km	+0.279	+0.283	+0.004	1.629	1.625
TOPOD	10 km	-0.205	-0.202	+0.003	1.569	1.583
TOPOD	25 km	-0.141	-0.153	-0.012	1.538	1.554
TOPOD	40 km	-0.074	-0.079	-0.005	1.558	1.563
ALPHA	10 km	-0.034	-0.056	-0.023	1.611	1.652
ALPHA	25 km	-0.128	-0.145	-0.017	1.587	1.663
ALPHA	40 km	-0.130	-0.125	+0.005	1.609	1.657
AET	10 km	-0.018	-0.061	-0.043	1.215	1.271
AET	25 km	-0.006	-0.056	-0.050	1.227	1.252
AET	40 km	+0.016	-0.030	-0.046	1.203	1.220

APPENDIX 1.S2. Predictor definitions, data provenance, and predictor geometry

This appendix documents the predictor set analysed in the chapter, the source and specification of each predictor, and the geometric relationships among predictors that justify the modelling choices. Predictor definitions and specifications are descriptive data-provenance material. The geometry diagnostics — collinearity across scales, variance inflation, nonlinearity, and environmental-space structure — are recomputed from the current predictor data rather than carried forward from earlier analyses. The a priori expectations that these predictors enter — the moderation hypotheses and their support logic — are documented separately in Appendix 1.S3.

1.S2.1 Predictor set and definitions

We analyse two land-use surrogates and four climate moderators, with two additive controls. Cropland prevalence (CROP) isolates the agricultural component of anthropogenic transformation, the pressure most directly associated with extensive forest conversion. Non-natural cover (NNCOV), the complement of natural land cover, integrates the total extent of human land use, including cropland, urban areas, and infrastructure. The two surrogates capture complementary dimensions. Landscapes dominated by agriculture show CROP approximately equal to NNCOV, whereas landscapes with substantial non-agricultural conversion show NNCOV exceeding CROP. We therefore treat them as complementary rather than interchangeable, and we fit them in separate models.

The four climate moderators represent two thermal axes and one hydric axis. Daily temperature range (DTR) and annual temperature range (ATR) quantify short-term and seasonal thermal variability, respectively. Maximum temperature of the warmest month (MAXT) quantifies peak-heat exposure. Precipitation of the driest quarter (PDRY) quantifies dry-season hydric stress; we invert this predictor during data assembly so that higher PDRY denotes drier dry seasons and thus greater hydric stress, and we interpret interactions involving PDRY relative to that orientation. We treat the thermal-variability metrics (DTR, ATR) and the peak-heat metric (MAXT) as competing rather than nested thermal hypotheses, and we do not privilege any one axis a priori.

Actual evapotranspiration (AET) and topographic complexity (TOPOD) enter the models as additive controls rather than as focal moderators. AET represents the water–energy balance that

structures baseline productivity, and TOPOD represents terrain-driven environmental buffering. The Priestley–Taylor coefficient (ALPHA) appears only as a proxy for AET in the energy-axis sensitivity check reported in Appendix 1.S5.3, not as a focal predictor. Table 1.S2.1 lists each predictor, its operational meaning, and its role in the analysis, and Figures 1.S2.1 to 1.S2.3 map the native surface of each predictor across the biome before aggregation to response units.

Table 1.S2.1. Predictor set, operational meaning, and analytical role.

Predictor	**Family**	**Operational meaning**	**Role in analysis**
CROP	Land-use pressure	Cropland prevalence (% cropland cover)	Focal land-use surrogate (agriculture-specific)
NNCOV	Land-use pressure	Non-natural cover (% non-natural land)	Focal land-use surrogate (total human land use)
DTR	Thermal variability	Mean diurnal temperature range (Bio02)	Focal climate moderator
ATR	Thermal variability	Annual temperature range (Bio07)	Focal climate moderator
MAXT	Thermal extreme	Maximum temperature of the warmest month (Bio05)	Focal climate moderator (competing thermal axis)
PDRY	Hydric stress	Dry-season precipitation scantness (Bio17, inverted — higher = drier)	Focal climate moderator
AET	Productivity / water–energy	Actual evapotranspiration (realised water–energy flux)	Additive control

Predictor	Family	Operational meaning	Role in analysis
ALPHA	Productivity / water–energy	Priestley–Taylor coefficient	AET proxy (energy-axis sensitivity check, Appendix 1.S5.3)
TOPOD	Abiotic-template control	Terrain ruggedness index	Additive control

1.S2.2 Data provenance and specifications

Each predictor derives from a documented source product at a stated temporal anchor and native spatial resolution, and all predictors are z-standardised before fitting (Table 1.S2.2). Predictors are aggregated to the response units at each of the analysed extents; the column templates that carry each predictor through the pipeline are documented in the released data dictionary.

Table 1.S2.2. Predictor specifications: source product, temporal anchor, native resolution.

Predictor	Description	Temporal anchor	Native resolution	Source
CROP	Cropland prevalence (%)	1992–2006 consensus	1 km	Tuanmu and Jetz 2014
NNCOV	Non-natural cover (% land)	2020	30 m	Souza et al. 2020 (MapBiomas)
DTR	Mean diurnal temperature range	1970–2000	~1 km	Hijmans et al. 2005; Title and Bemmels 2018
ATR	Annual temperature range	1970–2000	~1 km	Hijmans et al. 2005

Predictor	Description	Temporal anchor	Native resolution	Source
MAXT	Max temperature of warmest month	1970–2000	~1 km	Hijmans et al. 2005
PDRY	Precipitation of driest quarter (inverted — higher = drier)	1970–2000	~1 km	Hijmans et al. 2005
AET	Actual evapotranspiration	Derived	Derived	Title and Bemmels 2018
ALPHA	Priestley–Taylor coefficient	Trabucco baseline	1 km	Trabucco and Zomer 2010
TOPOD	Terrain ruggedness index	Static	1 km	Amatulli et al. 2018

1.S2.3 Predictor collinearity across scales

Several predictor pairs are strongly collinear across the 18 grain × extent combinations (Table 1.S2.3). The two thermal-variability metrics are strongly correlated but distinct (DTR–ATR, mean $r = 0.824$; range 0.806–0.842); the nonlinearity screen in S2.5 shows little evidence that either is a nonlinear re-expression of the other. The two energy metrics are constructionally linked and correspondingly strongly correlated (ALPHA–AET, mean $r = 0.815$), because ALPHA enters the construction of AET. The two land-use surrogates are correlated (NNCOV–CROP, mean $r = 0.800$), consistent with cropland forming a large component of total non-natural cover, though the two retain distinct information where non-agricultural conversion is extensive. We address these collinearities structurally rather than by residualisation — we fit the coupled thermal metrics (DTR, ATR) and the coupled energy metrics (ALPHA, AET) in separate models, so that no single model carries two strongly correlated predictors along one axis. We also excluded precipitation seasonality a priori because it represents an overlapping seasonal-hydric axis, and we retain PDRY as the focal dry-season-stress moderator; the geometry reported

here is therefore for the nine retained predictors, so precipitation seasonality does not appear in the diagnostics below.

Table 1.S2.3. Correlation stability across the 18 grain × extent combinations (nine retained predictors), sorted by mean strength.

Predictor pair	Scales	Mean r	Min	Max
DTR–ATR	18	0.824	0.806	0.842
ALPHA–AET	18	0.815	0.795	0.837
NNCOV–CROP	18	0.800	0.742	0.827
ALPHA–PDRY	18	−0.746	−0.757	−0.731
ALPHA–MAXT	18	−0.597	−0.645	−0.542
AET–PDRY	18	−0.577	−0.601	−0.553
ALPHA–CROP	18	−0.566	−0.613	−0.530
ALPHA–NNCOV	18	−0.539	−0.592	−0.474
TOPOD–CROP	18	−0.539	−0.605	−0.451
AET–ATR	18	0.534	0.368	0.685
PDRY–NNCOV	18	0.428	0.382	0.466
DTR–NNCOV	18	0.425	0.341	0.476
TOPOD–NNCOV	18	−0.420	−0.494	−0.373
MAXT–CROP	18	0.403	0.322	0.461
TOPOD–MAXT	18	−0.374	−0.388	−0.362
AET–CROP	18	−0.342	−0.368	−0.323

Predictor pair	Scales	Mean r	Min	Max
PDRY–CROP	18	0.317	0.300	0.338
MAXT–NNCOV	18	0.310	0.247	0.385
TOPOD–ALPHA	18	0.293	0.276	0.314

1.S2.4 Variance inflation and suppressor geometry

Variance inflation factors confirm that the two thermal-variability metrics cannot co-occur in a single model (Table 1.S2.4). In the full-climate configuration the thermal-variability metrics carry the highest inflation (DTR = 9.57, ATR = 8.12), so co-fitting them would produce unstable coefficient estimates; the land-use surrogates (CROP = 4.10, NNCOV = 3.98), dry-season precipitation scantness (PDRY = 3.16), the peak-heat metric (MAXT = 1.35), and the additive controls (AET = 2.32, TOPOD = 1.74) show low inflation. The diagnostic configurations show why we did not co-fit the two thermal-variability metrics as simultaneous moderators — their inflation stayed elevated whenever both appeared in the same diagnostic set (DTR = 5.48, ATR = 6.12 in the land-use × thermal configuration), and every factor fell below 3 only where one of the pair was absent (the land-use × water and control configurations). In the fitted interaction models we analysed each thermal-variability moderator separately, so no focal model carried both strongly correlated thermal-variability predictors together, and the reported interaction coefficients are not inflated by within-axis collinearity.

Table 1.S2.4. Variance inflation factors by model configuration (representative scale). Dashes mark predictors absent from a configuration; * marks VIF > 5.

Predictor	Full climate	LU × thermal	LU × water	Controls
CROP	4.10	2.20	1.77	—
NNCOV	3.98	—	—	—

Predictor	Full climate	LU × thermal	LU × water	Controls
MAXT	1.35	1.34	—	—
ATR	8.12*	6.12*	—	—
DTR	9.57*	5.48*	—	—
PDRY	3.16	—	1.66	—
AET	2.32	1.60	1.58	2.95
ALPHA	—	—	—	3.20
TOPOD	1.74	1.60	1.56	1.15

Correlation and variance inflation do not capture every gradient-sharing problem. Two predictors can be only moderately correlated yet share a dominant spatial gradient that distorts their additive partial coefficients — a suppressor geometry in which one predictor inflates or reverses another's coefficient without a corresponding structure-coefficient shift. We therefore flag suppressor relationships among the predictors directly, and we carry this diagnostic into the interaction models through the structure-coefficient concordance criterion (Appendix 1.S3), which excludes interaction terms whose product-term coefficient and structure coefficient disagree in sign. These suppressor relationships identify where the additive partial coefficients require that safeguard.

1.S2.5 Nonlinearity among the predictors

Relationships among the predictors are frequently nonlinear (Table 1.S2.5). At the representative scale, generalised additive fits outperform linear fits by more than 0.05 in explained variance for 36 of 72 ordered predictor–predictor relationships, with a maximum gain of 0.26. This is a property of predictor geometry — the predictors relate nonlinearly to one another — and it reinforces that the linear collinearity diagnostics understate their coupling. The screen also discriminates strong correlation from redundancy in form. The two ordered DTR–ATR relationships increased explained variance by only 0.028 and 0.051 when we allowed nonlinear

fits, against a maximum gain of 0.26 elsewhere in the same screen. Neither metric is therefore well described as a nonlinear re-expression of the other, and their strong correlation (mean r = 0.824) is consistent with a shared thermal-variability gradient rather than one metric restating the other. We therefore fitted them in separate models to keep their partial coefficients interpretable. We entered AET and TOPOD as linear additive terms, so controls-robust support demonstrates robustness to that stated control specification but does not exhaust nonlinear predictor coupling or unmodelled land-use × control moderation. Whether the land-use surrogates require polynomial terms in the richness models is a separate question, resolved by model selection and reported in Appendix 1.S4.

Table 1.S2.5. Strongest predictor–predictor nonlinearity at the representative scale (top 12 of the 36 ordered relationships exceeding the 0.05 gain threshold, of 72 tested). Gain = GAM R^2 − linear R^2.

Relationship (x → y)	**Linear R^2**	**GAM R^2**	**Gain**
ATR → CROP	0.009	0.271	0.262
ATR → PDRY	0.016	0.264	0.248
AET → MAXT	0.065	0.288	0.223
ATR → ALPHA	0.038	0.248	0.211
ATR → AET	0.238	0.442	0.204
ALPHA → NNCOV	0.350	0.526	0.176
PDRY → DTR	0.079	0.255	0.175
ATR → TOPOD	0.006	0.181	0.175
ATR → MAXT	0.009	0.178	0.169
AET → NNCOV	0.084	0.253	0.169

Relationship (x → y)	Linear R^2	GAM R^2	Gain
ATR → NNCOV	0.013	0.182	0.169
TOPOD → PDRY	0.000	0.166	0.166

1.S2.6 Source references (predictor data)

Amatulli G, Domisch S, Tuanmu M-N, et al (2018) A suite of global, cross-scale topographic variables for environmental and biodiversity modeling. Scientific Data 5:180040.

Hijmans RJ, Cameron SE, Parra JL, et al (2005) Very high resolution interpolated climate surfaces for global land areas. International Journal of Climatology 25:1965–1978.

Souza CM Jr, Shimbo JZ, Rosa MR, et al (2020) Reconstructing three decades of land use and land cover changes in Brazilian biomes with Landsat archive and Earth Engine. Remote Sensing 12:2735.

Title PO, Bemmels JB (2018) ENVIREM. An expanded set of bioclimatic and topographic variables increases flexibility and improves performance of ecological niche modeling. Ecography 41:291–307.

Trabucco A, Zomer RJ (2010) Global Soil Water Balance Geospatial Database. CGIAR Consortium for Spatial Information. Published online, available from the CGIAR-CSI GeoPortal at https://csidotinfo.wordpress.com

Tuanmu M-N, Jetz W (2014) A global 1-km consensus land-cover product for biodiversity and ecosystem modelling. Global Ecology and Biogeography 23:1031–1045.

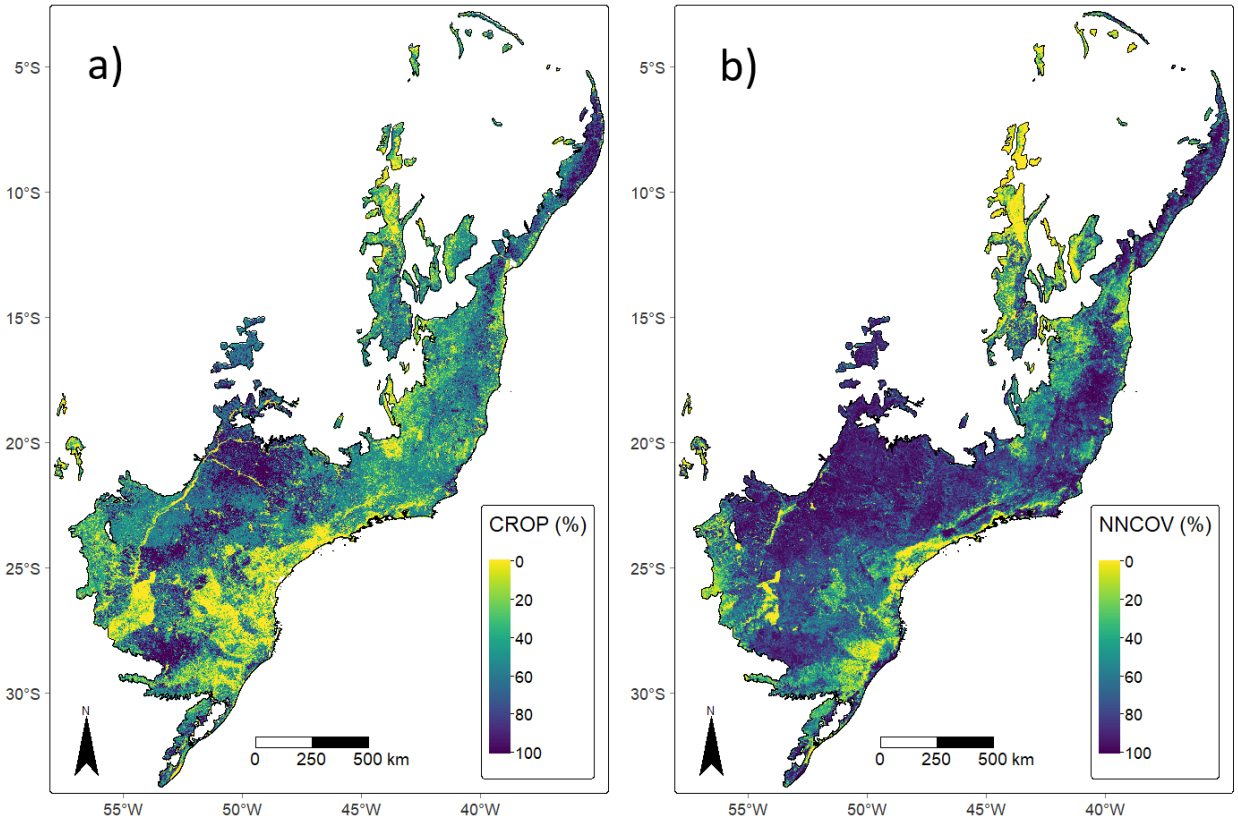


Figure 1.S2.1. Land-use surrogates across the Atlantic Forest. (a) Cropland prevalence (CROP), the percentage of cultivated and managed vegetation per 1-km cell from the global consensus land-cover product. (b) Non-natural cover (NNCOV), the percentage of transformed cover per cell, derived as 100 minus natural vegetation cover from the Atlantic Forest-specific 30-m annual land-cover product resampled to 1 km, so that higher values indicate greater transformation. Both panels show the native predictor surfaces before aggregation to response units, and colour scales run from low (dark) to high (light) values. The two surrogates are related but not interchangeable: cropland forms extensive contiguous tracts in the interior south and west, whereas non-natural cover additionally captures pasture, urban, mining and other anthropogenic classes, giving it broader coverage along the coast and in the north. Full provenance is given in Table 1.S2.2.

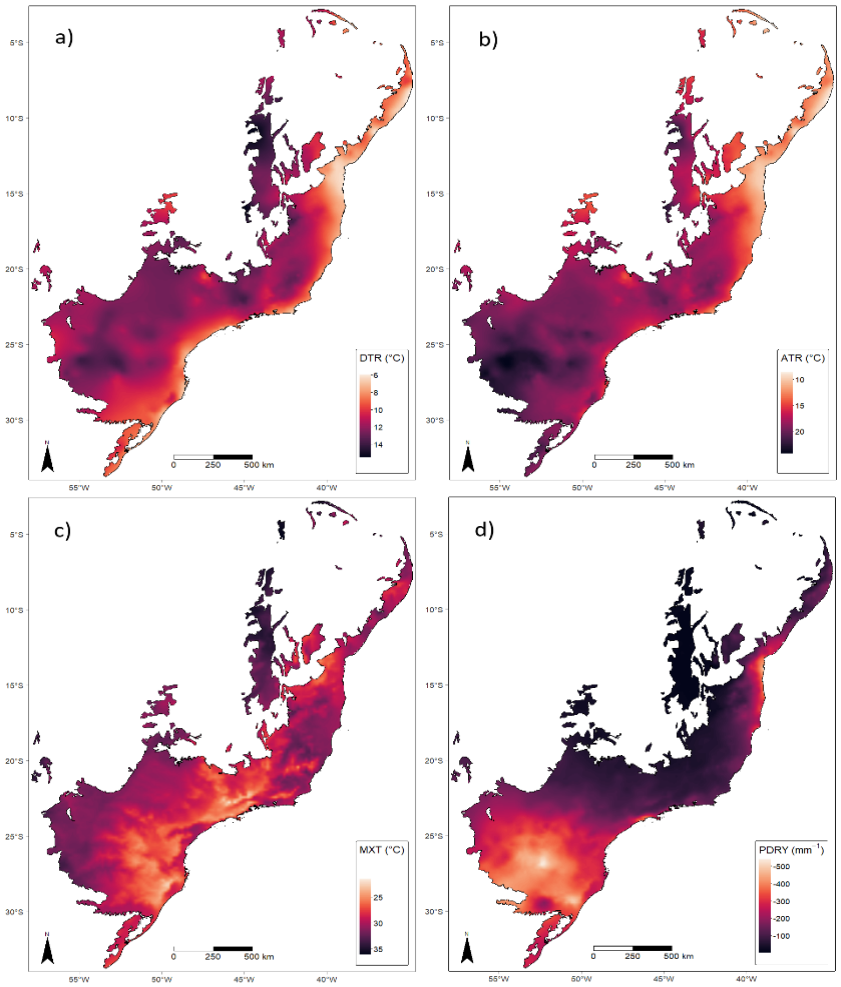


Figure 1.S2.2. Climate moderators across the Atlantic Forest. (a) Daily temperature range (DTR), the mean diurnal difference between maximum and minimum temperature. (b) Annual temperature range (ATR), the difference between the maximum temperature of the warmest month and the minimum of the coldest. (c) Maximum temperature of the warmest month (MAXT). (d) Precipitation of the driest quarter (PDRY), shown here with the axis inverted so that higher values denote drier dry seasons, matching the orientation used in the analyses. All panels show the native predictor surfaces before aggregation to response units. The two thermal-variability axes are strongly correlated across the biome (mean r = 0.824; Table 1.S2.3), both increasing inland and toward the south, which is why they are fitted in separate models and interpreted as one thermal-variability axis rather than as separable daily and annual mechanisms. Maximum heat and dry-season stress vary along partly different geographic gradients, the former peaking in the interior north and the latter in the south. Full provenance is given in Table 1.S2.2.

.

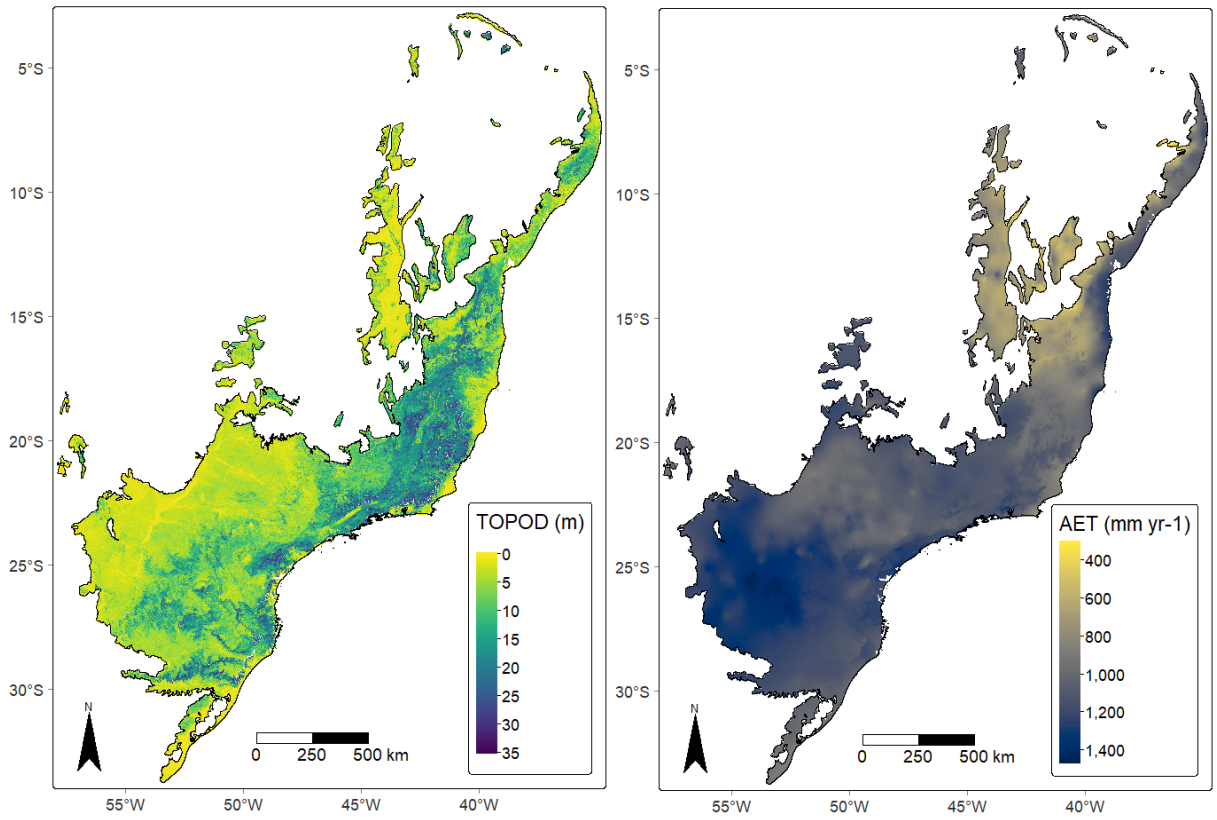


Figure 1.S2.3. Environmental controls across the Atlantic Forest. Topographic ruggedness (TOPOD), the terrain ruggedness index, indexing local relief. Actual evapotranspiration (AET), the realised water–energy flux, derived as the product of annual potential evapotranspiration and the Priestley–Taylor coefficient before standardisation. Both panels show the native predictor surfaces before aggregation to response units. These two enter the Tier 2 models as additive controls rather than as moderators, so that supported land-use × climate interactions can be evaluated after conditioning on broad water–energy and terrain structure. Both are spatially structured across the biome — AET declining southward and TOPOD concentrated along the coastal escarpment — which is why the spatial autocovariate is recomputed after they enter the model. Full provenance is given in Table 1.S2.2.

APPENDIX 1.S3. Predictor roles and support-classification rules

Table 1.1 states the chapter hypotheses, mechanisms, and alternative explanations. Here we document the predictor-level expectations, citation basis, and operational rules used to classify primary support, controlled support, and robustness. This appendix therefore serves as the predictor-role and classification reference rather than as a second hypothesis table. The predictor definitions and geometry are documented in Appendix 1.S2, and the model specification and diagnostic adequacy in Appendix 1.S4.

1.S3.1 Expected predictor roles

Each predictor entered the analysis with an a priori expected role grounded in ecological mechanism (Table 1.S3.1). Land-use pressure, thermal variability, thermal extremes, and hydric stress define the focal axes of expected biodiversity loss and land-use amplification, whereas productivity and terrain enter as additive controls. For the marginal expectations, the expected direction refers to the anticipated univariate association before conditioning on other predictors.

Table 1.S3.1. Predictor families, ecological mechanisms, expected roles, and key references.

Predictor family	Ecological mechanism	Expected role / direction	Key references
Land-use pressure (CROP, NNCOV)	Habitat filtering of forest-dependent species through matrix hostility, reduced forest extent and continuity, and resource simplification.	Negative, across taxa and scales.	Newbold et al. 2015

Predictor family	Ecological mechanism	Expected role / direction	Key references
Thermal variability (DTR, ATR)	Forest canopies buffer thermal variability; canopy loss increases exposure to atmospheric variability. DTR and ATR index proximate and broader dimensions of that exposure.	Negative marginal effects; negative land-use × climate interactions (variability amplifies land-use damage).	De Frenne et al. 2019; Ewers and Banks-Leite 2013; Jucker et al. 2020; Ismaeel et al. 2024; Williams and Newbold 2020
Thermal extreme (MAXT)	Acute peak-heat exposure; canopy loss can raise peak temperatures. MAXT competes with the variability moderators as the operative thermal axis (Mantyka-Pringle et al. 2012).	Negative land-use × MAXT interaction if peak heat compounds land-use damage; the contrast with DTR/ATR discriminates the thermal axis.	Sunday et al. 2014; Khaliq et al. 2014; Senior et al. 2017; Tourani et al. 2023; Mantyka-Pringle et al. 2012

Predictor family	Ecological mechanism	Expected role / direction	Key references
Hydric stress (PDRY)	Dry-season water stress compounds canopy loss; species relying on forest microclimatic buffering are most exposed where canopy is removed.	Negative land-use × PDRY interaction, with PDRY oriented so higher values denote drier dry seasons (see Appendix 1.S2).	Smith et al. 2023; Frishkoff et al. 2016; Oliver and Morecroft 2014
Productivity / water–energy (AET, ALPHA)	Higher realised productivity provides demographic subsidies, sustaining richness. Retained as additive control.	Positive marginal associations; AET the primary control, ALPHA a proxy-sensitivity alternative only.	Hawkins et al. 2003; Currie et al. 2004; Title and Bemmels 2018; Trabucco and Zomer 2010
Abiotic-template control (TOPOD)	Terrain structures both land-use placement (flat areas preferentially converted) and climatic exposure (rugged areas provide microrefugia). Retained as an additive control.	Positive marginal association.	Stein et al. 2014; Hjort et al. 2015; Amatulli et al. 2018

1.S3.2 Support and robustness classification

We classify a land-use × climate interaction as primary-supported at a grain × extent combination when it shows a parsimony gain of ΔAICc > 2 over the additive model and product-term structure-coefficient concordance (the interaction coefficient and its structure coefficient share sign). We treat the likelihood-ratio test ($p < 0.05$) as a secondary consistency check. Scale coherence — a consistent pattern of support across the scale matrix — is an interpretive pattern, not a per-cell support criterion. We then evaluate separately whether the coefficient direction matches the a priori expectation — for DTR, ATR, and PDRY, expected-direction support requires a negative interaction, whereas for MAXT a negative interaction would support peak-heat amplification and a positive or surrogate-dependent interaction indicates that MAXT is not a separable compound-damage axis.

We evaluated all land-use × climate interactions across the 18 grain × extent scale matrix. We treated repeated recovery across this matrix as scale coherence rather than as independent replication, while recognising that DTR, ATR, PDRY, and MAXT index climatic pathways that may vary in the grains and extents at which support is recovered.

We assess robustness by refitting each primary-supported interaction with actual evapotranspiration (AET) and topographic complexity (TOPOD) as additive controls, and we classify an interaction as controlled-supported when it retains a parsimony gain of ΔAICc > 2 over the controlled additive model — that is, when the interaction still earns its place once productivity and terrain are accounted for.

Alongside controlled support we report whether each interaction retains its sign and structure-coefficient concordance under the controls, and we classify interactions into three cases. A robust interaction retains the controlled ΔAICc gain. An absorbed interaction retains coefficient direction and structure-coefficient concordance after controls but loses the controlled ΔAICc gain, indicating that the controls absorb the interaction's parsimony support without reversing its direction. A control-sensitive interaction reverses sign or loses product-term structure-coefficient concordance under the controls. We treat such cases as specification-sensitive and do not infer a stable ecological direction from either layer alone.

The peak-heat interaction (LU × MAXT) enters as a competing thermal-axis test — a surrogate-consistent negative interaction would support peak-heat amplification, whereas surrogate-

dependent, positive, or geographically aligned support indicates that MAXT is not the separable amplification axis in this design.

Because DTR and ATR index distinct but strongly correlated components of thermal variability and are fitted in separate models (Appendix 1.S2), we interpret comparable DTR and ATR interactions as concordant evidence for a thermal-variability amplification axis rather than as separable daily- and annual-timescale mechanisms.

APPENDIX 1.S4. Model specification, model-building diagnostics, and adequacy

This appendix documents the model structure, the specification decisions that precede interaction testing, and the diagnostics that establish model adequacy. The predictor definitions and geometry are in Appendix 1.S2, the a priori expectations and support logic in Appendix 1.S3, and the full results in Appendix 1.S5. All quantities are recomputed from the current pipeline outputs; none are carried forward from earlier analyses.

1.S4.1 Model formula and random-effect structure

We model observed species richness with a Poisson generalized linear mixed model that carries an observation-level random effect to absorb overdispersion (Poisson-OLRE). For a land-use surrogate LU and a moderator M, the maximal candidate interaction model takes the form

$$\text{resp} \sim \text{EFFsb} + \text{LU} + \text{M} + \text{LU:M} + \text{space}$$
$$+ (1 \mid \text{method}) + (1 + \text{LU} + \text{M} \mid \text{taxa}) + (1 \mid \text{obs})$$

fitted with a Poisson family and a log link. The fixed effects comprise sampling effort (EFFsb), the land-use surrogate, the moderator, their interaction, and a spatial autocovariate (space) that places the models on consistent spatial footing. The random structure comprises a method-level intercept, taxon-level random slopes for the land-use surrogate and moderator (shown here at their maximal, correlated form), and an observation-level intercept that provides the overdispersion term. We selected the taxon-level random-slope structure for each model from a candidate set ranging from a random intercept only to correlated slopes for both terms (Section S4.4), so the fitted random-slope structure varies across models. We assess distributional adequacy using scaled quantile residuals rather than adopting an alternative count family as the baseline specification (Section S4.5).

This autocovariate summarises the broad-scale spatial structure remaining in the residuals, so the reported fixed effects are estimated after accounting for spatial dependence.

For this chapter, LU takes the two land-use surrogates (CROP, NNCOV) and M takes the four climate moderators (DTR, ATR, MAXT, PDRY), yielding the land-use × climate interaction terms tested across the 18 grain × extent combinations. The additive counterpart of each model drops the interaction term; the controlled counterpart adds actual evapotranspiration and

topographic complexity as additive controls (+ AET + TOPOD), and provides the basis for the controlled-support assessment reported in Appendix 1.S3 and Appendix 1.S5. The moderator-free additive land-use models used for Figure 1.3 are distinct from these LU + M additive counterparts and omit the climate moderator entirely.

1.S4.2 Baseline specification decision audit

To evaluate sensitivity to effort formulation, we anchored the audit on the reported site-based formulation. We classified a reported coefficient as directionally determinate when the ratio of its absolute value to its standard error reached two. For each determinate site-based coefficient we classified the unit-based alternative as determinate or indeterminate and as sign-concordant or sign-discordant, and we calculated magnitude ratios only for jointly determinate, sign-concordant pairs. This determinacy rule gates a nuisance-specification sensitivity comparison. It does not define biological support and it does not replace the chapter's primary criterion of a parsimony gain above two AICc units together with product-term structure-coefficient concordance.

Determinacy in this specification audit describes whether the audited coefficient carries a stable directional signal under alternative effort formulations. It does not imply that forest-specialist coefficients are estimated more precisely than forest-facultative coefficients, because the comparison involves different coefficient magnitudes and standard-error structures across strata.

We paired 36 cells per response stratum, comprising two land-use surrogates across the 18 grain by extent scale combinations. The determinacy gate isolated the stable directional comparison as intended. Among site-based determinate coefficients we found no sign discordance under the unit-based formulation in any response stratum. The four sign disagreements in the full 180-cell matrix all occurred among site-based indeterminate non-forest coefficients and therefore did not enter the agreement or magnitude summaries. Magnitude ratios below one indicate that the unit-based formulation returned smaller absolute coefficients than the reported site-based formulation.

Too few non-forest coefficients were determinate to compute a meaningful agreement statistic, consistent with the precision limitation documented in Appendix 1.S5. One non-forest pair was jointly determinate and sign-concordant, with a magnitude ratio of 1.492, and we did not summarise magnitude agreement for that stratum because only one pair was jointly determinate.

Table 1.S4.1. Five-stratum effort-formulation audit. Paired cells are two land-use surrogates across 18 grain by extent combinations. Determinate denotes an absolute coefficient at least twice its standard error on the reported site-based formulation. Magnitude ratios are unit-based over site-based absolute coefficients, computed only for jointly determinate, sign-concordant pairs.

Response stratum	Paired cells	EFFsb determinate	EFFub determinate and concordant	EFFub indeterminate and concordant	Discordant among EFFsb-determinate	Magnitude ratio, median [range]
Total richness, full universe	36	36	36	0	0	0.839 [0.714–0.968]
Total richness, restricted universe	36	36	36	0	0	0.802 [0.628–0.933]
Forest-specialist	36	36	32	4	0	0.849 [0.776–0.938]
Forest-facultative	36	36	27	9	0	0.814 [0.558–0.947]
Non-forest	36	3	1	2	0	not summarised

We evaluated temporal aggregation separately for total richness at the 10-km grain, the only scope for which we fitted the year-retained comparison. Across the 14 land-use by extent comparisons, year-retained and year-aggregated coefficients agreed in sign in all 14 cases, with a median absolute-magnitude ratio of 0.945 and a range of 0.888 to 0.968. Because direction was

stable across both audits, we carried the site-based effort formulation and year-aggregated responses forward as the baseline specification for the interaction analyses.

1.S4.3 Polynomial-term structure

We allowed polynomial terms for focal predictors during model building and retained them only where model comparison supported them, rather than imposing linear responses throughout. Across the 720 fitted interaction models (144 per response stratum), the selected structure carries a land-use polynomial in 197 and a moderator polynomial in 37 (Table 1.S4.2). The strata differ. Land-use polynomials are frequent in the aggregate and restricted-total responses (62 and 66 of 144) and in forest-facultative richness (38), less frequent in forest specialists (31), and absent in non-forest richness (0), whereas moderator polynomials are rare overall and concentrate in forest specialists (20) and non-forest richness (14), with three in the restricted total. The concentration of polynomial terms on the land-use surrogates, rather than on the moderators, is consistent with threshold-like richness responses to land conversion.

Table 1.S4.2. Selected polynomial structure by stratum (counts over the 144 interaction models per stratum).

Stratum	Interaction models	Land-use polynomial	Moderator polynomial
total_full	144	62	0
total_restr	144	66	3
forest-specialist	144	31	20
forest-facultative	144	38	0
non-forest	144	0	14
All	720	197	37

1.S4.4 Random-slope selection and convergence

We select the random-slope structure for each model by comparing candidate structures on AICc and retaining the best-supported structure. The candidate structures range from a random intercept only, through random slopes for the land-use surrogate or the moderator alone, to correlated random slopes for both. The selected structures differ between strata (Table 1.S4.3). Forest-specialist models almost always select correlated random slopes for both terms (128 of 144) or a land-use slope alone (16), indicating substantial among-taxon variation in both land-use and moderator responses. The aggregate and weaker-guild strata select more varied structures — moderator slopes alone dominate the full and restricted totals (94 and 69 of 144) and forest-facultative richness (72), whereas land-use slopes alone dominate non-forest richness (69) — indicating that among-taxon slope variation is present but less consistently for both terms outside the forest-specialist stratum. This procedure is AICc-based and does not privilege any structure a priori; we retain only converged fits, and convergence is monitored alongside residual adequacy (Section S4.5).

Additive main-effect models used TAXA random intercepts rather than TAXA random slopes. Random-slope land-use models were AICc-favoured in some forest-specialist cells but were weakly identified, inflating fixed-effect standard errors with little change in the land-use coefficient. We therefore used the random-intercept structure for the additive main-effect models that support Figure 1.3, while retaining the selected random-slope structures for the LU × climate interaction models.

We treated taxon-level random-slope results as model-structure diagnostics rather than biological estimates of clade-specific vulnerability. Survey methods were nested within dataset-defined groups, so group-level slope heterogeneity could not be fully separated from method-specific land-use responses. This limitation affects interpretation of random-slope heterogeneity, not the forest-affiliation guild contrasts reported in the main text.

Specialist non-natural-cover contrasts remained modestly sensitive to random-effect structure even after adopting the selected nonlinear form. We therefore treat the random-slope sensitivity as unresolved model-structure variation rather than as evidence against the main forest-dependence ordering.

Three controlled models across the 720 interaction cells fail the convergence and positive-definite-Hessian check, all in non-forest dry-season models, and in each the full and reduced fits disagree. The controlled model comparison is therefore not computable for those cells, and we record it as missing rather than as an absence of support, so it does not enter the controlled-support counts. One of the three is the single non-forest cell that meets the primary support criterion, whose controlled status we therefore report as not determined (Table 1.2).

Table 1.S4.3. Selected random-slope structure by stratum (counts over the 144 interaction models per stratum).

Selected random-slope structure	**total_full**	**total_restr**	**forest-specialist**	**forest-facultative**	**non-forest**
Intercept only	15	7	0	23	1
Land-use slope	17	41	16	11	69
Moderator slope	94	69	0	72	18
Correlated land-use + moderator	18	27	128	38	56
Total	144	144	144	144	144

1.S4.5 Residual diagnostic adequacy

We assess distributional adequacy with scaled quantile residuals (DHARMa) computed from the fitted models, rather than by adopting an alternative count family as the baseline specification. For each of the 720 interaction models — 144 per response stratum across the five strata — we

simulate residuals and test uniformity (Kolmogorov–Smirnov), dispersion, and outlier frequency at $\alpha = 0.05$; every model fits without error (Table 1.S4.4). The diagnostics separate by component and by stratum. Dispersion is adequate in every model of every stratum (100%), so the observation-level random effect captures overdispersion as intended, and outlier frequency is adequate in most models (81.2–93.8% across strata). Uniformity separates the aggregate from the guild-restricted strata. The test is non-significant (adequate) in 77.8% of full-total models and 9.7% of restricted-total models, but in essentially none of the three guild-restricted strata — forest-specialist, forest-facultative, and non-forest richness (all 0.0%).

We tested whether the specialist uniformity rejections reflected an inadequate count family by refitting a representative specialist model under negative-binomial, zero-inflated-Poisson, and zero-inflated-negative-binomial alternatives. In that representative check all four families rejected uniformity while retaining adequate dispersion and outlier behaviour, so the departure was family-invariant there and did not support replacing the Poisson-OLRE baseline with a negative-binomial or zero-inflated family. Because DHARMa uniformity tests are amplified by sample size, we judged the departure by the magnitude of the Kolmogorov–Smirnov statistic rather than its p-value alone (Hartig 2022). The aggregate and guild-restricted strata differ in this magnitude. The full-total residuals sit at a well-behaved level (median $D = 0.030$ across the 144 models), whereas the forest-specialist residuals depart further (median $D = 0.096$; range 0.067–0.122) — above a well-behaved reference ($D \approx 0.04$) but below the departures that accompany genuine dispersion or zero-inflation problems ($D \approx 0.19$); the forest-facultative and non-forest strata show the same qualitative uniformity departure as forest specialists. We do not attribute the specialist departure to a specific cause. The representative refit did not implicate count-family choice, and the source is not identified from these diagnostics. Because dispersion remains adequate across all strata and outlier behaviour remains adequate in most models, the departure is unlikely to materially affect the interaction coefficients, their direction and structure-coefficient concordance, or the ΔAICc-based support classification. We therefore retain the Poisson-OLRE specification and report the forest-specialist residual departure as a moderate, unattributed limitation.

We assess residual spatial independence with the same procedure as the single-predictor diagnostics. For the load-bearing interactions, we compute response-scale residuals and estimate Moran's I distance-based spline correlograms with 95% bootstrap envelopes (100 resamples).

We found no systematic geographic clustering in the load-bearing interaction diagnostics; the observed correlograms remained within the bootstrap envelope around zero across distance. Because these envelopes are estimated from the data rather than from an asymptotic null, they are not inflated by the large number of assemblages. A weak short-distance correlation remains within the envelope rather than indicating residual dependence. These diagnostics indicate that the spatial autocovariate captures residual spatial dependence adequately, so the reported interactions do not reflect unmodelled spatial structure. Temporal non-independence does not arise for the reported models, which use year-aggregated responses; the single-predictor diagnostics, which included year-retained variants, likewise show no residual temporal autocorrelation.

Finally, we compare the land-use and land-use × climate coefficients with and without the spatial autocovariate, to verify that the reported effects do not arise from spatial adjustment. Across the load-bearing forest-specialist interactions, the interaction sign remains stable in 22 of 24 models, with a modest median change in coefficient magnitude (approximately 9%). The thermal-variability interactions (land-use × DTR and land-use × ATR) retain or strengthen under spatial adjustment, whereas the peak-heat interaction (non-natural cover × MAXT) weakens — consistent with maximum temperature's alignment with the broad geographic structure that the residual autocovariate partly absorbs. We used these comparisons to verify that the core thermal-variability result and the primary cropland × dry-season signal did not arise as artefacts of spatial adjustment, whereas the weakening of the non-natural-cover × MAXT interaction under the autocovariate is consistent with its interpretation as a geography-aligned thermal-axis contrast.

Table 1.S4.4. Residual-diagnostic adequacy by response stratum: percentage of the 144 models per stratum with non-significant (adequate) scaled-residual tests at $\alpha = 0.05$. All 720 models fit without error.

Response stratum	Models	Uniformity (KS)	Dispersion	Outlier
total_full	144	77.8%	100%	81.2%
total_restr	144	9.7%	100%	93.8%

Response stratum	Models	Uniformity (KS)	Dispersion	Outlier
forest-specialist	144	0.0%	100%	93.8%
forest-facultative	144	0.0%	100%	93.8%
non-forest	144	0.0%	100%	88.2%

1.S4.6 Control-sensitivity audit of the interaction layer

Adding conditioning variables can reverse a focal coefficient when the controls redistribute variance shared with that term. Such a reversal does not by itself establish a change in the underlying ecological relationship. We therefore audited every fitted interaction cell for sign change between the primary and controlled layers rather than interpreting the controlled coefficient alone. We recorded, for all 720 cells, the primary coefficient, the controlled coefficient, and the product-term structure coefficient in both layers, so that direction is decided by agreement across layers rather than by either layer in isolation.

Eighty-two of 720 cells change sign under the joint AET + TOPOD control set. The changes are concentrated rather than diffuse. Non-forest richness contributes 44, forest-specialist richness 18, forest-facultative richness 11, and the two aggregate anchors 9 between them. Non-forest coefficients are numerically near zero throughout, so sign is uninformative in that stratum and we draw no conclusion from its instability.

Thermal-variability moderation is not control-sensitive. None of the 72 forest-specialist daily-temperature-range and annual-temperature-range cells changes sign, and none of the aggregate-stratum thermal cells does either. A single forest-facultative non-natural-cover cell changes sign on a coefficient near zero, and it does not meet the primary support criterion. The specialist concentration reported in the main text therefore rests on coefficients whose direction is stable across both layers.

Maximum temperature is the control-sensitive axis. Every guild stratum changes character between layers. Forest-specialist cropland by maximum-temperature coefficients are predominantly negative before conditioning and predominantly positive after it, while forest-

facultative and non-forest coefficients move the opposite way. Applying the across-layer agreement rule resolves the specialist case by scale rather than by layer. At the coarsest response grain the uncontrolled structure coefficient is already positive while the uncontrolled coefficient is negative, so those cells are structure-coefficient-discordant and fail the primary criterion, whereas the controlled coefficient turns positive and becomes concordant. At finer grains the pattern reverses. Neither layer resolves the direction on its own.

Because forest-specialist cropland by maximum-temperature met the support criterion at no scale in either layer, we treat this instability as a diagnostic and do not carry it into the reported forest-specialist result. The supported forest-specialist maximum-temperature cells occur under non-natural cover and remain positive in both layers. We therefore report maximum temperature as a surrogate-dependent axis aligned with broad geographic and land-use context, and we treat the cropland by maximum-temperature partial as a suppression-sensitive diagnostic rather than as an estimate. We distinguish this control-sensitivity audit analytically from the spatial-adjustment comparison in Section S4.5 because the two checks alter different model components. They yield compatible diagnostic patterns, but they use the same data and model family and therefore do not provide independent confirmation.

We do not compare interaction magnitudes across climate moderators or land-use surrogates. Cropland and non-natural-cover models sit on different polynomial surfaces, and the maximum-temperature set additionally pools opposite-signed and differently supported cells. We report direction, structure-coefficient concordance, and scale coherence instead.

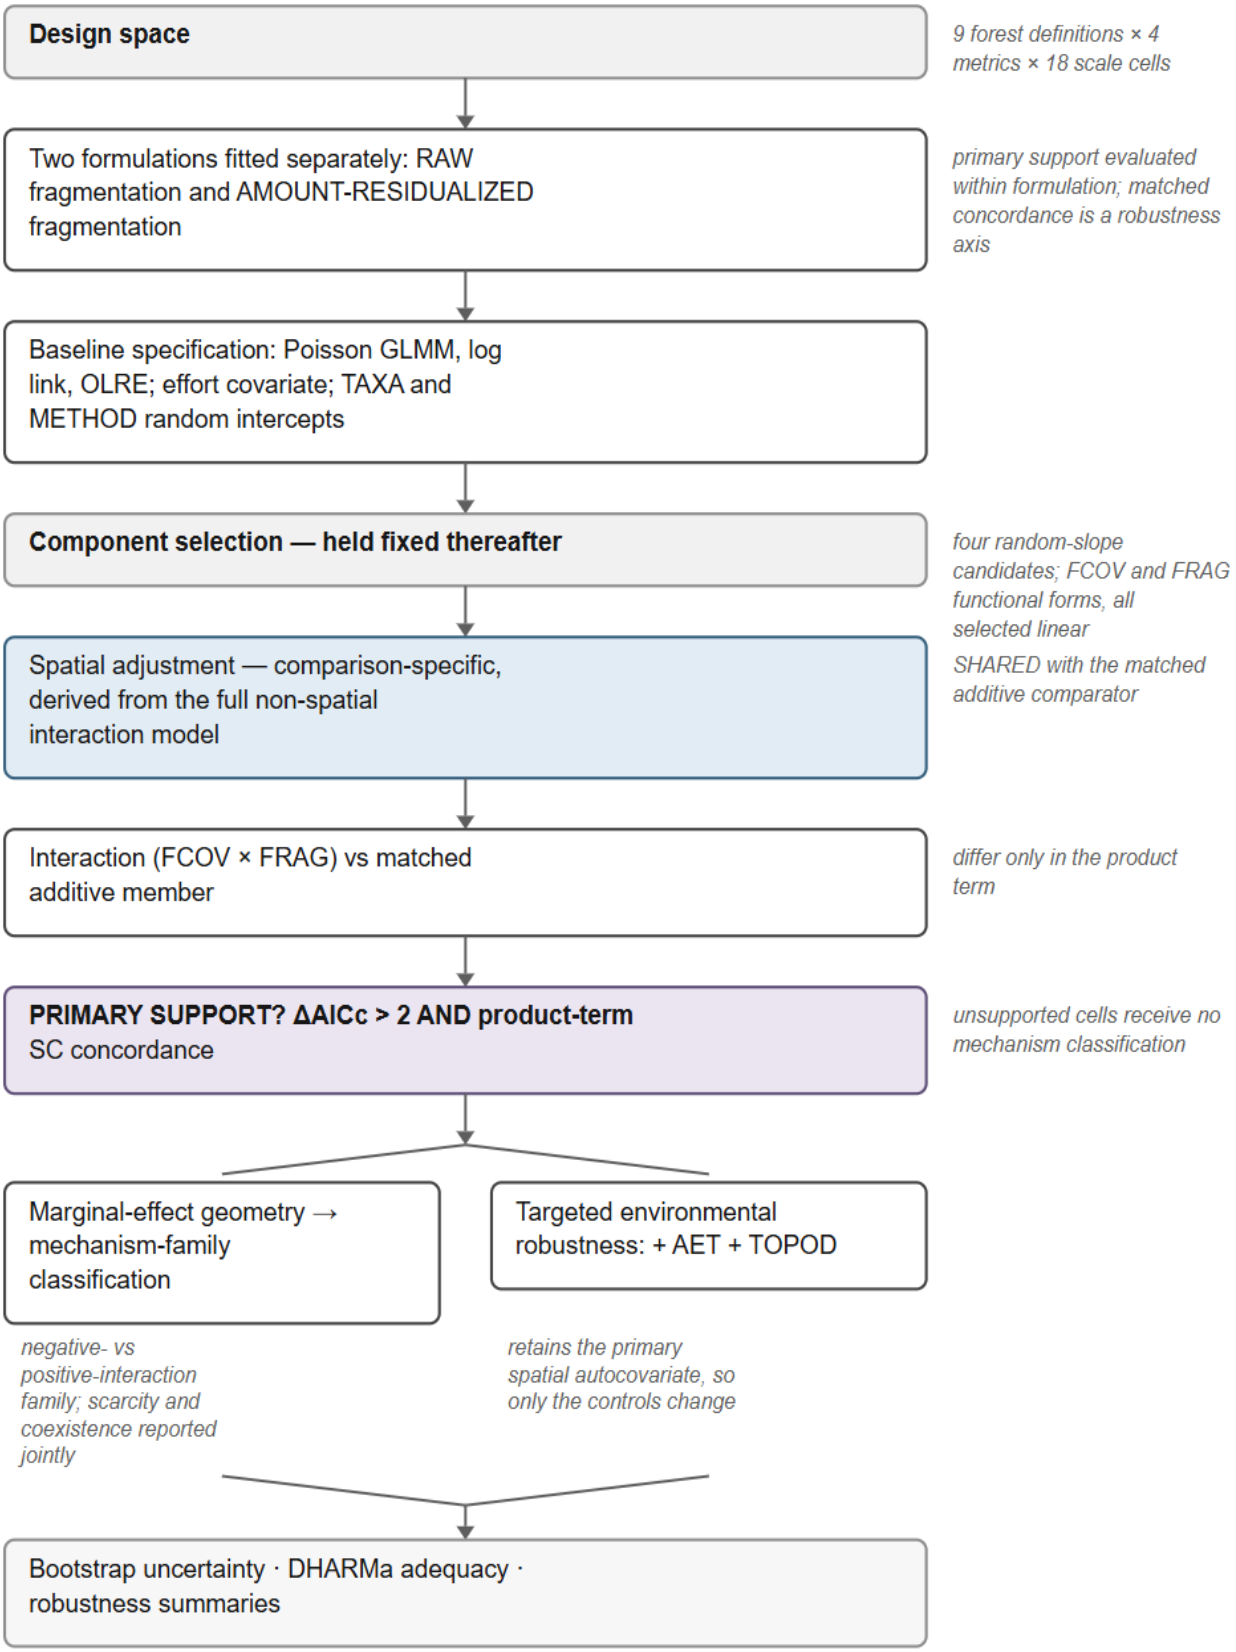

Design space
9 forest definitions × 4 metrics × 18 scale cells
Two formulations fitted separately: RAW fragmentation and AMOUNT-RESIDUALIZED fragmentation
primary support evaluated within formulation; matched concordance is a robustness axis
Baseline specification: Poisson GLMM, log link, OLRE; effort covariate; TAXA and METHOD random intercepts
Component selection — held fixed thereafter
four random-slope candidates; FCOV and FRAG functional forms, all selected linear
Spatial adjustment — comparison-specific, derived from the full non-spatial interaction model
SHARED with the matched additive comparator
Interaction (FCOV × FRAG) vs matched additive member
differ only in the product term
PRIMARY SUPPORT? ΔAICc > 2 AND product-term SC concordance
unsupported cells receive no mechanism classification
Marginal-effect geometry → mechanism-family classification
Targeted environmental robustness: + AET + TOPOD
negative- vs positive-interaction family; scarcity and coexistence reported jointly
retains the primary spatial autocovariate, so only the controls change
Bootstrap uncertainty · DHARMa adequacy · robustness summaries

Figure 1.S4.1. Model-building and inference sequence for the climatic-moderation analysis. Component and functional-form selection precede hypothesis testing and are held fixed within matched comparisons. Tier 1 compares each land-use × climate interaction with its reduced counterpart under a shared spatial autocovariate; only primary-supported interactions, requiring ΔAICc > 2 and product-term structure-coefficient concordance, enter Tier 2. Tier 2 adds AET and TOPOD, recomputes the spatial autocovariate from the corresponding controlled full non-spatial model, and again holds that term fixed between the full and reduced models. Bootstrap, residual and predictive diagnostics evaluate the load-bearing estimates. The sequence is repeated across the eighteen grain × extent combinations and five response strata.

APPENDIX 1.S5. Land-use × climate interaction results and bootstrap validation

This appendix reports the full land-use × climate interaction results that support the main-text Table 1.2 and the load-bearing claims of the Discussion. We compute every quantity from the current pipeline outputs and define interaction support from the support criterion stated in the Methods; we do not infer support from precomputed flags or by matching reported counts. Figure titles describe what each figure plots, not the outcome it shows. The predictor geometry is in Appendix 1.S2, the a priori expectations and support logic in Appendix 1.S3, and the model specification in Appendix 1.S4.

1.S5.1 Interaction support criterion and validation

We classified a land-use × climate interaction as primary-supported at a grain × extent combination when it met both criteria. The interaction had to show a parsimony gain of ΔAICc greater than 2 over the additive model and structure-coefficient concordance, meaning that the interaction coefficient and its structure coefficient shared sign. We treat the likelihood-ratio test ($p < 0.05$) as a secondary consistency check; interactions supported by the likelihood-ratio test alone, or that are structure-coefficient-discordant, are not primary-supported.

This criterion is stricter than the precomputed support flag carried in the pipeline output. Of the 720 tested interaction cells across the five response strata, 145 satisfy the support criterion; a further 56 cells gain parsimony ($\Delta AICc > 2$) or likelihood-ratio support but fail structure-coefficient concordance or the ΔAICc threshold, which we exclude from primary support; the remaining 519 cells are supported by neither. We report the support-criterion results throughout, and the precomputed flag serves only to validate the computation.

1.S5.2 Primary interaction support

The primary support matrix reports, for each response stratum, land-use surrogate, and climate moderator, the number of the 18 grain × extent combinations that meet the support criterion (Table 1.S5.1). Across the five strata, 145 of the 720 cells are primary-supported, and support concentrates strongly in the forest-specialist stratum and in the thermal-variability moderators. Forest-specialist richness shows primary support for the annual-temperature-range interactions at 17 of 18 combinations with cropland and 16 of 18 with non-natural cover, for the daily-temperature-range interactions at 16 (cropland) and 13 (non-natural cover), and for the dry-

season interaction at all 18 combinations with cropland; the peak-heat interaction is supported at 7 of 18 (non-natural cover), whereas cropland × maximum temperature and non-natural-cover × dry-season reach no forest-specialist combination. Forest-facultative support is confined to the non-natural-cover surrogate (7 daily-temperature-range, 6 annual-temperature-range, and 1 peak-heat combinations), and non-forest richness reaches primary support at a single combination. The two aggregate anchors fall between the forest specialist and the weaker guilds (20 and 23 of 144), consistent with dilution of the forest-specialist signal within pooled richness. Zero-support combinations appear explicitly in the matrix, so untested and unsupported cells remain distinguishable.

Table 1.S5.1. Five-stratum primary interaction support — count of the 18 grain × extent combinations meeting the support criterion (ΔAICc > 2 and product-term structure-coefficient concordance), by land-use × climate family and response stratum.

LU × climate family	**total_full**	**total_restr**	**forest-specialist**	**forest-facultative**	**non-forest**
Cropland × DTR	5	0	16	0	0
Cropland × ATR	0	0	17	0	0
Cropland × MAXT	4	0	0	0	0
Cropland × PDRY	7	7	18	0	1
Non-natural cover × DTR	2	1	13	7	0

LU × climate family	total_full	total_restr	forest-specialist	forest-facultative	non-forest
Non-natural cover × ATR	1	14	16	6	0
Non-natural cover × MAXT	1	1	7	1	0
Non-natural cover × PDRY	0	0	0	0	0
All (of 144)	20	23	87	14	1
Thermal (of 72)	8	15	62	13	0

Table 1.S5.2. Forest-specialist primary and controlled interaction support detail (count of 18 grain × extent combinations; controlled = ΔAICc > 2 under AET + TOPOD). The forest-specialist rows of Figure 1.S5.1 show the same cells resolved by grain and extent.

Interaction	Primary (of 18)	Controlled (of 18)
Cropland × daily temperature range	16	16
Non-natural cover × daily temperature range	13	13
Cropland × annual temperature range	17	17
Non-natural cover × annual temperature range	16	16
Non-natural cover × maximum temperature	7	6

Interaction	Primary (of 18)	Controlled (of 18)
Cropland × dry-season precipitation scantness (PDRY)	18	0

Forest-facultative cropland suppressor geometry. Sixteen cropland combinations cleared the parsimony threshold in forest-facultative richness (cropland × daily temperature range 1, cropland × maximum temperature 6, cropland × dry-season 9), but all sixteen were structure-coefficient-discordant and therefore excluded from primary support. The interpretable forest-facultative interactions are non-natural-cover interactions only.

Non-forest precision. Non-forest richness reached primary support at a single combination despite interaction coefficients of magnitude comparable to forest-facultative richness under non-natural cover. The near-absence of support reflects lower precision — non-forest standard errors are approximately 1.8 times larger — so the model-fit gain rarely cleared the parsimony threshold, rather than an absence of interaction geometry.

1.S5.3 Support under additive controls

The controlled analysis refits every fitted interaction with actual evapotranspiration and topographic complexity as additive controls, and reports how many of the primary-supported combinations retain a parsimony gain of $\Delta AICc > 2$ under the controls (controlled support; Appendix 1.S3). We grid-complete the controlled matrix across the same 18 grain × extent combinations as the primary matrix. We record tested cells that lack controlled support as zeros and retain non-computable controlled comparisons as missing, so an absence of support remains distinct from indeterminate controlled status. The thermal-variability interactions are robust to conditioning — the annual-temperature-range interactions retain controlled support at all primary-supported combinations in forest specialists (17 with cropland, 16 with non-natural cover), and the daily-temperature-range interactions at 16 (cropland) and 13 (non-natural cover), with coefficients strengthening rather than attenuating under control (Figure 1.S5.1). By contrast, the cropland × dry-season interaction, primary-supported at all 18 combinations, retains controlled support at none (0 of 18). The interaction loses its parsimony gain under conditioning, visible as a fully supported row carrying no controls-robust marker in Figure 1.S5.1.

The cropland × dry-season interaction is the clearest absorbed case (Appendix 1.S3). It retains its sign and structure-coefficient concordance under the controls (18 of 18) yet loses the ΔAICc gain (0 of 18), so a control absorbs its parsimony support without reversing its direction. This pattern is consistent with the joint-control alternatives anticipated a priori (Appendix 1.S3), including water–energy confounding and terrain-structured cropland placement; because actual evapotranspiration and topographic complexity enter jointly rather than one at a time, the results do not separate productivity from terrain as the absorbing control. The peak-heat interaction retains controlled support at 6 of 7 combinations with a small positive coefficient, consistent with its role in thermal-axis discrimination rather than as a compound-damage interaction, and its support is confined to the coarser response grains in Figure 1.S5.1.

We also refitted the forest-specialist thermal-variability interactions by substituting the Priestley–Taylor coefficient for actual evapotranspiration in the control set. Among the 60 primary-supported combinations with computable ALPHA-controlled comparisons, coefficients agreed in sign with their actual-evapotranspiration counterparts, retained product-term structure-coefficient concordance, and retained controlled support at every combination, with a median coefficient shift of +6.7% (range −0.7% to +13.7%). Two comparisons did not satisfy nested log-likelihood monotonicity, so we treated their coefficient direction, structure-coefficient concordance, and support status as not determined. We interpret the computable results as directional insensitivity to the energy-axis proxy rather than as an independent test of energy confounding, because the Priestley–Taylor coefficient enters the construction of actual evapotranspiration.

1.S5.4 Bootstrap validation

We validate the load-bearing interactions with a parametric bootstrap, computing per-model bootstrap intervals rather than averaging interval bounds across models (which would not yield a valid interval for any quantity), and reporting their per-interaction envelopes. For each interaction we report the primary mean effect size and the 95% interval envelope across the supported primary models (Table 1.S5.3). We also report a count-based aggregate as the number of target models whose bootstrap interval excludes zero (Table 1.S5.4). The target-model denominator equals the primary-supported scale combinations plus their corresponding controlled refits; thus the 16 primary-supported cropland × daily-temperature-range combinations yield 32 target fits, the 17 cropland × annual-temperature-range combinations yield

34, and the 18 cropland × dry-season combinations yield 36. The aggregate discriminates among the interactions. The thermal-variability interactions in forest specialists exclude zero at nearly all target models (daily temperature range with cropland, 31 of 32; annual temperature range with cropland, 33 of 34; and at all target models with non-natural cover, 26 of 26 and 32 of 32). The cropland × dry-season interaction excludes zero at 21 of 36 target models, so the bootstrap shows the same relative weakness in the cropland × dry-season interaction that the controlled analysis identifies as loss of parsimony support. The two diagnostics use the same data and model family, so they corroborate rather than independently confirm one another. The peak-heat interaction (non-natural cover × maximum temperature) excludes zero at most target models (13 of 14), but with a positive coefficient — opposite to the compound-damage direction of the thermal-variability interactions — and in total richness it is surrogate-dependent (negative with cropland, positive with non-natural cover). Because the peak-heat interaction was positive and surrogate-dependent, we interpret it as thermal-axis discrimination linked to geographic context rather than as compound land-use damage or weak support.

Table 1.S5.3. Parametric-bootstrap primary-tier interval envelopes for the load-bearing forest-specialist interactions. Primary mean β averages the SC-concordant primary-supported combinations, and the envelope spans the widest bounds of their per-model 95% intervals; exclude-zero counts appear in Table 1.S5.4. The main-text Table 1.2 instead reports the mean controlled BW and its bootstrap envelope over the same primary-supported cells (Table 1.S5.5).

Interaction	**Primary (of 18)**	**Primary mean β**	**95% CI (primary-model envelope)**
Cropland × daily temperature range	16	-0.073	[-0.142, +0.000]
Non-natural cover × daily temperature range	13	-0.122	[-0.201, -0.028]
Cropland × annual temperature range	17	-0.103	[-0.192, +0.000]

Interaction	Primary (of 18)	Primary mean β	95% CI (primary-model envelope)
Non-natural cover × annual temperature range	16	-0.133	[-0.238, -0.039]
Non-natural cover × maximum temperature	7	+0.073	[+0.012, +0.144]
Cropland × dry-season precipitation scantness (PDRY)	18	-0.104	[-0.196, -0.018]

Table 1.S5.4. Bootstrap validation: target models whose 95% interval excludes zero, by interaction (forest-specialist stratum).

Interaction	Land-use surrogate	Target fits excluding zero
Daily temperature range	Cropland	31 / 32
Daily temperature range	Non-natural cover	26 / 26
Annual temperature range	Cropland	33 / 34
Annual temperature range	Non-natural cover	32 / 32
Maximum temperature	Non-natural cover	13 / 14
Dry-season precipitation scantness (PDRY)	Cropland	21 / 36

From the same bootstrap we summarise the predictive adequacy of the bootstrapped effect-size estimates — the relative root-mean-square error of the bootstrap estimates (0.266 for forest specialists, 0.390 for total richness), the coverage of their 95% intervals (1.000 in both strata), and bootstrap predictive-calibration p-values. The predictive p-values centre near 0.5 in both strata (forest specialists 0.504, total richness 0.500), indicating a well-calibrated bootstrap, and complement the residual-distribution diagnostics in Appendix 1.S4.

These coefficient summaries describe dependent scale combinations rather than estimate a single pooled interaction parameter. The interval envelopes span per-cell bootstrap intervals across the scale matrix and are not confidence intervals for a pooled effect. We therefore use them to document direction, scale coherence, and the provenance of Table 1.2 values, not to rank effect sizes across land-use surrogates or climate moderators. The primary-tier envelopes in S5.2a and the controlled summaries in Table 1.2 refer to different model tiers and must not be combined.

1.S5.5 Provenance of the main-text Table 1.2

Table 1.S5.5 documents the provenance of the six load-bearing forest-specialist rows of the main-text Table 1.2, which carry the chapter's central results. For each, it reports the primary support count (n/18), the controls-robust count as the nested subset of primary-supported cells retaining ΔAICc > 2 under AET + TOPOD (n/n_primary), and the mean controlled BW with its bootstrap envelope over the same primary-supported cells. The mean controlled BW is the mean of the cell-specific controlled bootstrap medians, and the envelope spans their cell-specific 95% bootstrap intervals rather than giving a confidence interval for that mean. Cells enter this summary when they retain sign stability and product-term structure-coefficient concordance under the controls, which for total_full CROP × DTR excludes one of the five controls-robust cells because its controlled structure coefficient is indistinguishable from zero. This makes every forest-specialist Table 1.2 value traceable to the scale combinations that produced it; the full five-stratum primary and controls-robust provenance is given in the main-text Table 1.2 and Table 1.S5.1. Where controls-robust is zero — as for cropland × dry-season in forest specialists (18/18 primary, 0/18 controls-robust) — the primary count documents where the signal was present and the controls-robust count documents that it did not survive conditioning.

Table 1.S5.5. Provenance of the six load-bearing forest-specialist rows of the main-text Table 1.2: per interaction, the primary support count (n/18), the nested controls-robust count (n/n_primary), and the mean controlled BW with its bootstrap envelope over the primary-supported cells. For cropland × dry-season the mean controlled BW of -0.039 is shown although no cell retained controlled support.

Interaction	**Primary (n/18)**	**Controls-robust (n/n_primary)**	**Controlled BW mean [bootstrap envelope]**	**Robustness class**
Cropland × daily temperature range	16/18	16/16	-0.094 [-0.162, -0.019]	robust
Non-natural cover × daily temperature range	13/18	13/13	-0.183 [-0.284, -0.082]	robust
Cropland × annual temperature range	17/18	17/17	-0.111 [-0.192, -0.022]	robust
Non-natural cover × annual temperature range	16/18	16/16	-0.176 [-0.286, -0.070]	robust
Non-natural cover × maximum temperature	7/18	6/7	+0.079 [-0.005, +0.157]	thermal-axis discrimination
Cropland × dry-season precipitation scantness (PDRY)	18/18	0/18	-0.039 [-0.119, +0.041]	absorbed (joint AET + TOPOD)

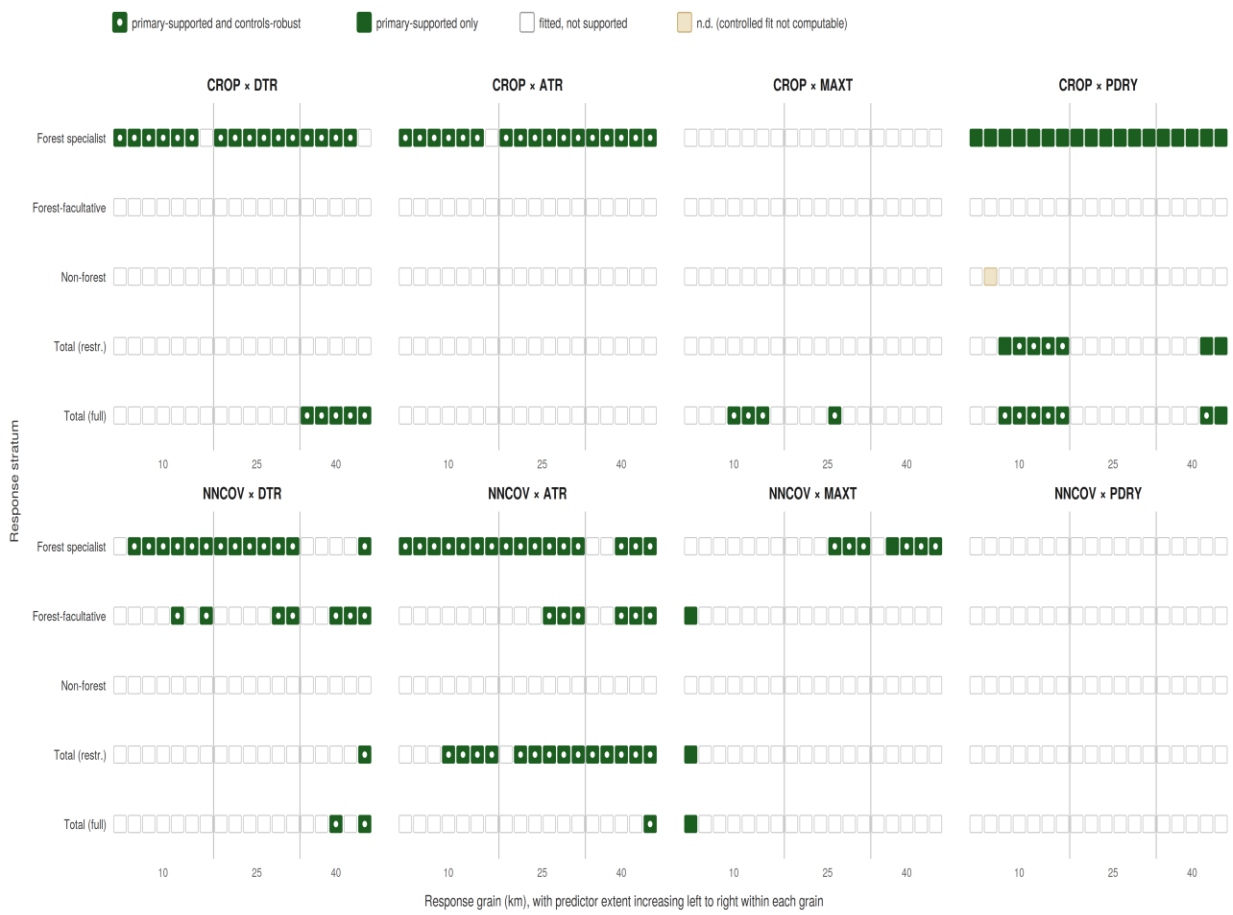


Figure 1.S5.1. Resolved interaction atlas across all five response strata. Each panel shows one land-use × climate family, with response strata as rows and all 18 grain × extent cells as columns, ordered by increasing predictor extent within each response-grain block. Filled cells met the primary support criterion, an inset dot marks cells that also retained the parsimony gain under the joint AET + TOPOD controls, open cells were fitted but not supported, and the shaded cell denotes a controlled comparison that was not computable. Across the design, 145 of 720 cells met the primary criterion and 118 retained controls-robust support.

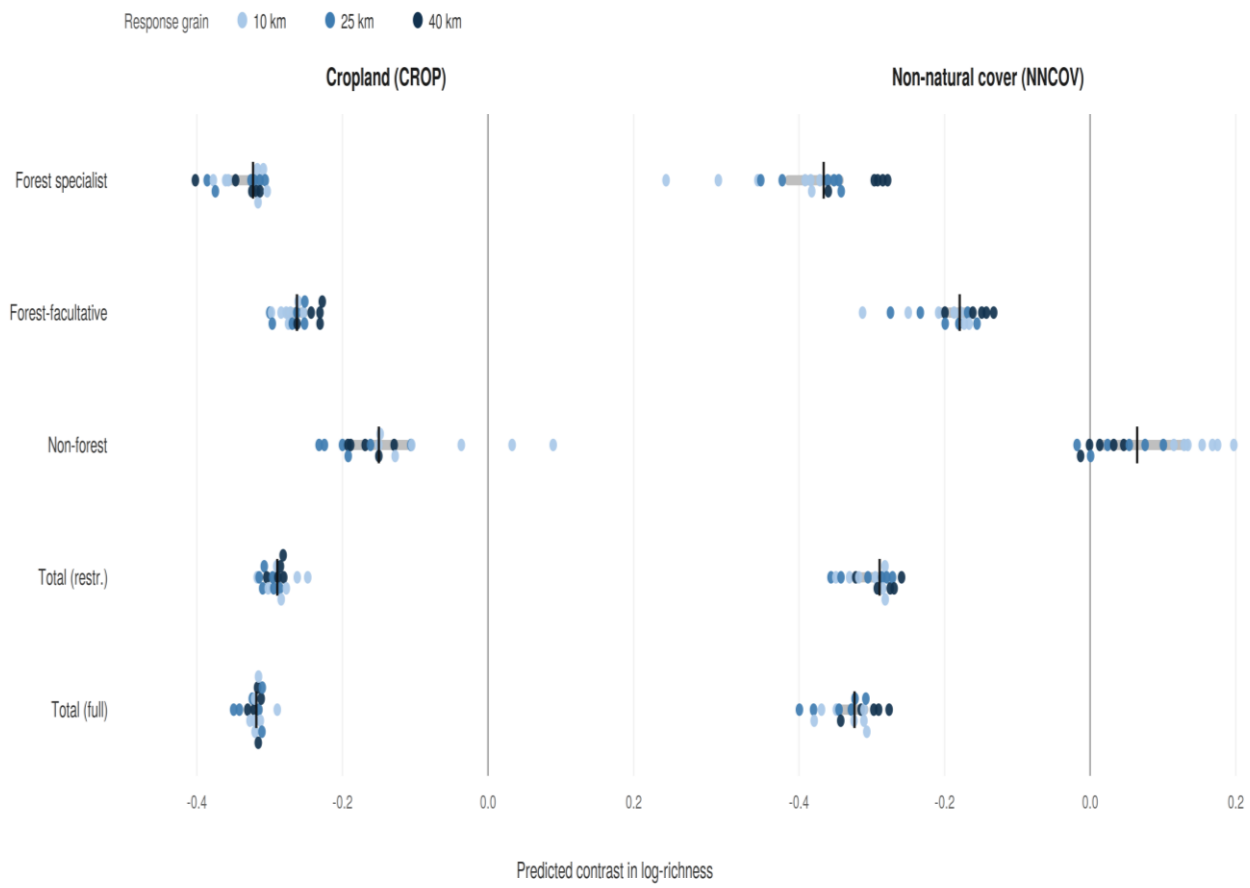


Figure 1.S5.2. Land-use main-effect contrasts resolved by scale. Each point is one grain × extent cell, coloured by response grain, with the interquartile bar and the median across the 18 cells retained for comparison with main-text Figure 1.3. Contrasts are predicted across the shared 10th to 90th percentile land-use interval under a common random-intercept structure with the selected functional form.

APPENDIX 2.S1. Chapter-specific response documentation

This chapter uses the shared Atlantic Forest assemblage compilation, forest-affiliation classification and multiscale response design described in the General Introduction and Appendix 0.S1. Here we document only chapter-specific response construction and diagnostics.

2.S1.1 Stratum precision and effective information

The strata differ substantially in mean richness, and under a Poisson likelihood the precision of a coefficient scales with the mean count. Relative to the restricted total at the 10 km grain, expected standard errors are approximately 1.2 times larger for forest-facultative richness, 2.2 times larger for forest-specialist richness, and 3.0 times larger for non-forest richness. Two consequences follow, and both bear on interpretation. First, the guild-concentration expectation is tested against a precision gradient running opposite to it, so evidence concentrating in the specialist stratum is stronger than raw counts of supported combinations suggest. Second, weak or indeterminate results in the non-forest stratum are consistent both with a true null and with the lowest precision of any stratum, so that stratum is informative about direction but should not be read as a precise zero. We therefore report non-forest outcomes alongside their standard errors rather than as an absence of effect.

APPENDIX 2.S2. Environmental predictors, provenance, and geometry

This appendix documents predictor definitions and provenance, the correlation structure among predictors, and the structure-coefficient classification used to interpret additive coefficients. The fitted factorial combines two land-use surrogates, CROP and NNCOV, with five focal moderator axes — EVICV and EVIENT for vegetational heterogeneity, TOPOD and CLIMD for geophysical heterogeneity, and AET for realised water–energy availability. Figures 2.S2.1 to 2.S2.4 map the native surface of each fitted predictor across the biome, grouped by that structure. We retain GEOD, PET, and ALPHA only outside the fitted interaction factorial. GEOD is a retired relief-based proxy, because terrain ruggedness and climatic heterogeneity together span the abiotic heterogeneity construct. PET and ALPHA are alternative operationalisations of the same water–energy system as AET, from which AET is derived, and re-enter only as targeted proxy-swap robustness checks for energy conditioning.

2.S2.1 Predictor definitions and provenance

Table 2.S2.1. Predictor operationalisation, provenance, and design status. Focal predictors are fitted in the additive and interaction factorial; screen-only predictors are retained in the single-predictor screen only (GEOD retired; PET and ALPHA re-enter as proxy-swap robustness checks). All source layers are native 1 km except non-natural cover, which is 30 m aggregated to the analysis grid, and AET, which is derived. Names in parentheses are the column stems identifying each predictor in the deposited datasets; each is completed with the predictor-extent suffix _{e}mn, where {e} is the extent in kilometres, except climatic diversity, which is climdiv_{e}. PET derives from the ENVIREM current-climate release (~1960–1990) and ALPHA from a soil-water balance averaged over 1950–2000; both are climatological rather than dated observation windows.

Predictor	Source and derivation	Period	Design status
Land-use pressure			

Predictor	Source and derivation	Period	Design status
CROP (cropcov)	Cropland prevalence (%) from the global 1-km consensus land-cover product; cultivated and managed vegetation class aggregated across five source maps (Tuanmu and Jetz 2014).	1992–2006	focal (land use)
NNCOV (natvcov_atlr)	Non-natural cover (%), derived as 100 − natural vegetation cover so that higher values indicate greater transformation. AF-specific 30-m annual land-cover; classes are agriculture, pasture, urban, mining and other anthropogenic (Vancine et al. 2024).	2020	focal (land use)
Vegetational			
EVICV (evicv)	Coefficient of variation of MODIS EVI within the landscape window. First-order greenness texture metric; reflects compositional variability in vegetation greenness. Termed greenness-magnitude variability in the main text (Tuanmu and Jetz 2015).	2000–2005	focal
EVIENT (evient)	Entropy of MODIS EVI within the landscape window. Second-order greenness texture metric; reflects spatial disorderliness of vegetation greenness arrangement. Termed configurational diversity in the main text (Tuanmu and Jetz 2015).	2000–2005	focal
Geophysical			

Predictor	Source and derivation	Period	Design status
TOPOD (rugged)	Terrain ruggedness index (TRI) from GMTED2010-derived global 1-km topographic variables; quantifies local terrain relief and heterogeneity (Amatulli et al. 2018).	Static	focal
GEOD (geolandiv)	Shannon diversity of geomorphological landform classes (10 geomorphon classes: flat, peak, ridge, shoulder, spur, slope, hollow, footslope, valley, pit) derived from the GMTED2010 digital elevation model by morphometric classification (Amatulli et al. 2018).	Static	screen-only (retired)
CLIMD (climdiv)	Shannon diversity of bioclimatic classes within the landscape window. Classes taken from the global environmental stratification of Metzger et al. (2013); Shannon diversity of those strata calculated within each buffer using landscapemetrics.	Static	focal
Water–energy			
PET (ampet)	Annual potential evapotranspiration (mm yr^{-1}) under unlimited moisture supply; indexes atmospheric evaporative demand. ENVIREM product (Title and Bemmels 2018).	~1960–1990	screen-only (proxy-swap)
ALPHA (ptalpha)	Priestley–Taylor coefficient (AET/PET); water-balance efficiency, indexing the fraction of potential evapotranspiration realised given water availability (Trabucco and Zomer 2019).	1950–2000	screen-only (proxy-swap)

Predictor	Source and derivation	Period	Design status
AET (derived)	Realised water–energy flux proxy (mm yr^{-1}), derived as ENVIREM annual PET × Trabucco and Zomer ALPHA before z-transformation. Not a source product; derived from the two baseline layers above (Title and Bemmels 2018; Trabucco and Zomer 2019). ALPHA is itself defined against CGIAR-CSI potential evapotranspiration, so the product approximates rather than reconstructs actual evapotranspiration, which is why AET enters as a water–energy proxy.	Derived	focal (energy)

2.S2.2 Predictor collinearity: correlations and variance inflation across scales

We computed Pearson correlations among all ten screened predictors at every grain and extent, giving 810 predictor-pair by scale combinations. Three pairs dominate the correlation structure, and each is handled by a standing exclusion rather than by variance-inflation control. Realised and potential evapotranspiration and the moisture ratio are definitionally linked, and the strongest of those associations reaches 0.815 on average. The two land-use surrogates correlate at 0.800. Terrain ruggedness and climatic heterogeneity correlate at 0.673. No model in the design contains any of these pairs together.

The correlation structure is stable across the scale matrix. Across the eighteen grain by extent combinations the range of each leading pair is narrow — 0.042 for the strongest energy pair, 0.085 for the land-use surrogates, 0.039 for potential evapotranspiration with the moisture ratio — so these associations are properties of the predictors rather than of any particular scale. Terrain ruggedness with climatic heterogeneity is the exception. Its range of 0.191 is the widest of the leading pairs, and the association weakens monotonically as the predictor window widens, from 0.74 at the narrowest extent to 0.61 at the broadest. The two are therefore most separable where the predictor window is largest.

We therefore paired the correlation audit with a variance-inflation check across fitted and diagnostic predictor configurations (Table 2.S2.2.).

Every configuration the design actually fits remains below 3.3. The values above the conventional threshold arise only in the full diagnostic configuration, which places both land-use surrogates in one model, and that configuration is never fitted because the design admits a single land-use surrogate at a time. These flags therefore describe the geometry of the predictor data rather than the condition of any estimated model. A second diagnostic configuration places terrain ruggedness and climatic heterogeneity together and returns 2.85 and 2.30, so the case for separating those two rests on the interpretability of their partial coefficients and not on variance inflation.

Figure 2.S2.5 summarises predictor geometry in two panels: the correlation matrix among the seven predictors carried into the fitted factorial, at the central scale combination of response grain 25 km and predictor extent 55 km; and the absolute correlation of each moderator with each land-use surrogate across all eighteen grain × extent combinations, ordered by median. Screened predictors are omitted.

Table 2.S2.2. Variance inflation by predictor configuration. Values above the conventional threshold arise only in configurations the design does not fit.

Configuration	**Predictors**	**Highest VIF**	**Fitted in the design**
full_het	Both land-use surrogates with all focal moderators	CROP 6.37; NNCOV 5.39	No — one land-use surrogate per model
lu_veg	Cropland with the vegetational moderators, terrain and energy	CROP 3.22	Yes

Configuration	Predictors	Highest VIF	Fitted in the design
lu_abiotic	Cropland with both geophysical moderators and energy	TOPOD 2.85	No — one geophysical moderator per model
controls	Terrain with the two energy proxies	ALPHA 3.20	Yes
aet_alpha	The two energy proxies with cropland and terrain	ALPHA 4.08	Diagnostic for the proxy-swap geometry

2.S2.3 Relationships among heterogeneity metrics, energy proxies, and terrain

Linear associations were estimated in both directions for every predictor pair at every scale, and each was refitted with a thin-plate spline so that the gain from relaxing linearity could be quantified. This covers 1,620 direction by scale combinations, of which 43 pairs at the representative combination (10 km response grain, 55 km predictor extent) gain more than 0.05 in explained variance under the spline.

The screen separates two superficially similar relationships. Geomorphological landform diversity is close to a nonlinear re-expression of terrain ruggedness. The linear association explains 0.28 of its variance and the spline fit 0.60, the largest nonlinear gain anywhere in the matrix. Terrain ruggedness with climatic heterogeneity shows no nonlinear gain above the reporting threshold at all, so that association is linear and moderate. The two abiotic operationalisations are therefore not equivalent. Landform diversity re-expresses relief in a different functional form, which is why it is screened rather than carried forward, whereas climatic heterogeneity indexes a distinct environmental axis that relief partly generates.

That generative relationship is directional. Relief produces climatic heterogeneity through orographic and lapse-rate effects, and climatic heterogeneity does not produce relief. The two geophysical moderators are consequently not symmetric tests of the persistent-template mechanism. Buffering detected on climatic heterogeneity may reflect terrain acting through it,

whereas buffering detected on terrain ruggedness cannot be attributed to climatic heterogeneity. We fit the two in separate models, read concordant support as one gradient recovered by two operationalisations rather than as two independent confirmations, and weight that support most heavily at the broader extents where the two are least associated.

Realised evapotranspiration sits close to orthogonal to every heterogeneity axis, correlating at -0.02 with the coefficient of variation of greenness, -0.10 with greenness entropy, -0.12 with climatic heterogeneity and 0.15 with terrain ruggedness. Conditioning heterogeneity moderation on realised productivity is therefore a test against a nearly independent axis rather than a partition of shared variance, which is what makes survival under energy conditioning informative.

The three energy proxies differ markedly in their coupling to land use. The moisture ratio correlates at -0.591 with cropland prevalence and -0.592 with non-natural cover, whereas realised evapotranspiration reaches -0.359 and -0.29. Potential evapotranspiration runs in the opposite direction, at +0.557 and +0.616. The proxy carried into the fitted design is thus the least land-use-coupled of the three, and because the swap check substitutes proxies whose coupling to land use runs the other way, an interaction surviving the swap has been tested against opposed coupling geometries rather than against a perturbation of the same one.

Exposure to suppression is asymmetric between the two vegetational metrics. At the representative combination, greenness entropy is moderately coupled to both land-use surrogates, at 0.613 with cropland prevalence and 0.668 with non-natural cover, whereas the coefficient of variation of greenness is weakly coupled at 0.246 and 0.499. A null result therefore carries different weight for the two. For greenness entropy it must be distinguished from suppression, whereas for the coefficient of variation the geometry leaves less scope for suppression to arise.

2.S2.4 Structure-coefficient classification of focal predictors

We classify each focal additive predictor by the relationship among its univariate marginal direction, its uncontrolled structure coefficient, and its partial slope, so that ecological interpretation rests on the structure-coefficient-plus-univariate reference wherever collinearity geometry distorts the partial slope (Table 2.S2.3). The classification covers the five focal predictors; GEOD, PET, and ALPHA are screened out of the fitted factorial and are not

classified here. Both greenness-texture metrics (EVICV, EVIENT) are marginally negative in 100% of cells in all four forest-affiliation layers and positive in the non-forest control, so the forest-layer negativity is a real, control-reversed signal rather than a procedural artefact. The specialist EVICV association is negative on three independent pre-control lines — the univariate marginal slope, the uncontrolled structure coefficient (−0.101) and the primary partial slope (−0.045), both computed over the broad-domain specialist models — but the pooled and facultative EVICV partials are suppression-sensitive, so we treat those as geometry-distorted rather than as ecological signs. Terrain (TOPOD) and realised energy (AET) carry clean, taxon-coherent partial slopes, and AET reverses sign in the non-forest control — the expected signature of a forest-mediated process. Climatic heterogeneity (CLIMD) carries a weak structure coefficient and 0% taxon sign agreement at fine extent in specialists, so its direction is not taxon-coherent where the CROP × CLIMD interaction is strongest.

Table 2.S2.3. Structure-coefficient classification of the five focal predictors: univariate marginal direction, uncontrolled structure coefficient, partial-slope geometry, and classification. From the deposited uncontrolled main-effect structure-coefficient table, columns sc and beta_unctrl, specialist stratum, broad domain, pooled over both land-use pairings. GEOD, PET, and ALPHA are screened out of the factorial and not classified.

Focal predictor	Univariate direction (forest layers)	Uncontrolled SC	Partial-slope geometry	Classification
EVICV	negative, 100% of cells in all four forest layers; positive in non-forest control	−0.101 (specialist)	primary partial −0.045; suppression-sensitive in pooled/facultative	suppression-sensitive; specialist-negative, control-reversed

Focal predictor	Univariate direction (forest layers)	Uncontrolled SC	Partial-slope geometry	Classification
EVIENT	negative, 100% in forest layers; positive in control	concordant with EVICV sign	parallels EVICV; weaker separation from zero	suppression-sensitive; configurational, not magnitude
TOPOD	positive, 100% CI exclusion, 100% taxon agreement, all layers	concordant, strong	clean partial slope	clean; strongest marginal predictor
CLIMD	positive in forest layers; specialist 0% taxon sign agreement at fine extent	weak (|SC| small)	0% CI exclusion at fine extent	SC-weak; direction not taxon-coherent in specialist
AET	positive +0.077 to +0.205, 100% CI exclusion in forest layers; reverses in control	concordant	clean partial slope	clean; control reverses (forest-mediated)

2.S2.5 Nonlinearity among the predictors

Relationships among the predictors are frequently nonlinear (Table 2.S2.4). At the representative combination (10 km response grain, 55 km predictor extent), generalised additive fits outperform linear fits by more than the 0.05 gain threshold in 43 of the 90 ordered predictor–predictor relationships, with a maximum gain of 0.315. This is a property of predictor geometry — the predictors relate nonlinearly to one another — and it reinforces that the linear collinearity diagnostics understate their coupling. The strongest nonlinear relationships involve the energy proxies and the configurational greenness metric (EVIENT), consistent with a curved coupling between realised energy and greenness texture across the sampled landscapes. The two greenness-texture metrics are themselves nonlinearly related (EVICV → EVIENT gain 0.162),

which is why we treat compositional variability (EVICV) and configurational disorder (EVIENT) as distinct axes rather than as interchangeable re-expressions of one greenness signal. We entered AET and TOPOD as linear additive controls, so controls-robust support demonstrates robustness to that stated linear specification but does not exhaust nonlinear predictor coupling; whether the land-use surrogates require polynomial terms in the richness models is resolved separately by model selection (Appendix 2.S4).

Table 2.S2.4. Strongest predictor–predictor nonlinearity at the representative combination (10 km response grain, 55 km predictor extent; top rows exceeding the 0.05 gain threshold, of 90 ordered relationships tested). Gain = GAM R^2 − linear R^2. From the deposited predictor-nonlinearity table. Relationships involving the screen-only predictors GEOD and PET are omitted from the displayed rows.

Relationship (x → y)	**Linear R^2**	**GAM R^2**	**Gain**
ALPHA → EVIENT	0.123	0.311	0.189
ALPHA → NNCOV	0.350	0.526	0.176
AET → NNCOV	0.084	0.253	0.169
EVIENT → AET	0.011	0.176	0.165
EVICV → EVIENT	0.312	0.474	0.162
TOPOD → ALPHA	0.099	0.246	0.147
TOPOD → EVIENT	0.051	0.196	0.145
EVIENT → ALPHA	0.123	0.260	0.138
EVICV → NNCOV	0.249	0.375	0.126

2.S2.6 References

Amatulli G, Domisch S, Tuanmu M-N, et al (2018) A suite of global, cross-scale topographic variables for environmental and biodiversity modelling. Scientific Data 5:180040.

Metzger MJ, Bunce RGH, Jongman RHG, Sayre R, Trabucco A, Zomer R (2013) A high-resolution bioclimate map of the world: a unifying framework for global biodiversity research and monitoring. Global Ecology and Biogeography 22:630–638.

Title PO, Bemmels JB (2018) ENVIREM: an expanded set of bioclimatic and topographic variables increases flexibility and improves performance of ecological niche modeling. Ecography 41:291–307.

Trabucco A, Zomer R (2019) Global High-Resolution Soil-Water Balance. figshare Dataset. https://doi.org/10.6084/m9.figshare.7707605.v3

Tuanmu M-N, Jetz W (2014) A global 1-km consensus land-cover product for biodiversity and ecosystem modelling. Global Ecology and Biogeography 23:1031–1045.

Tuanmu M-N, Jetz W (2015) A global, remote sensing-based characterization of terrestrial habitat heterogeneity. Global Ecology and Biogeography 24:1329–1339.

Vancine MH, Muylaert RL, Niebuhr BB, et al (2024) The Atlantic Forest of South America: spatiotemporal dynamics of the vegetation and implications for conservation. Biological Conservation 291:110499.

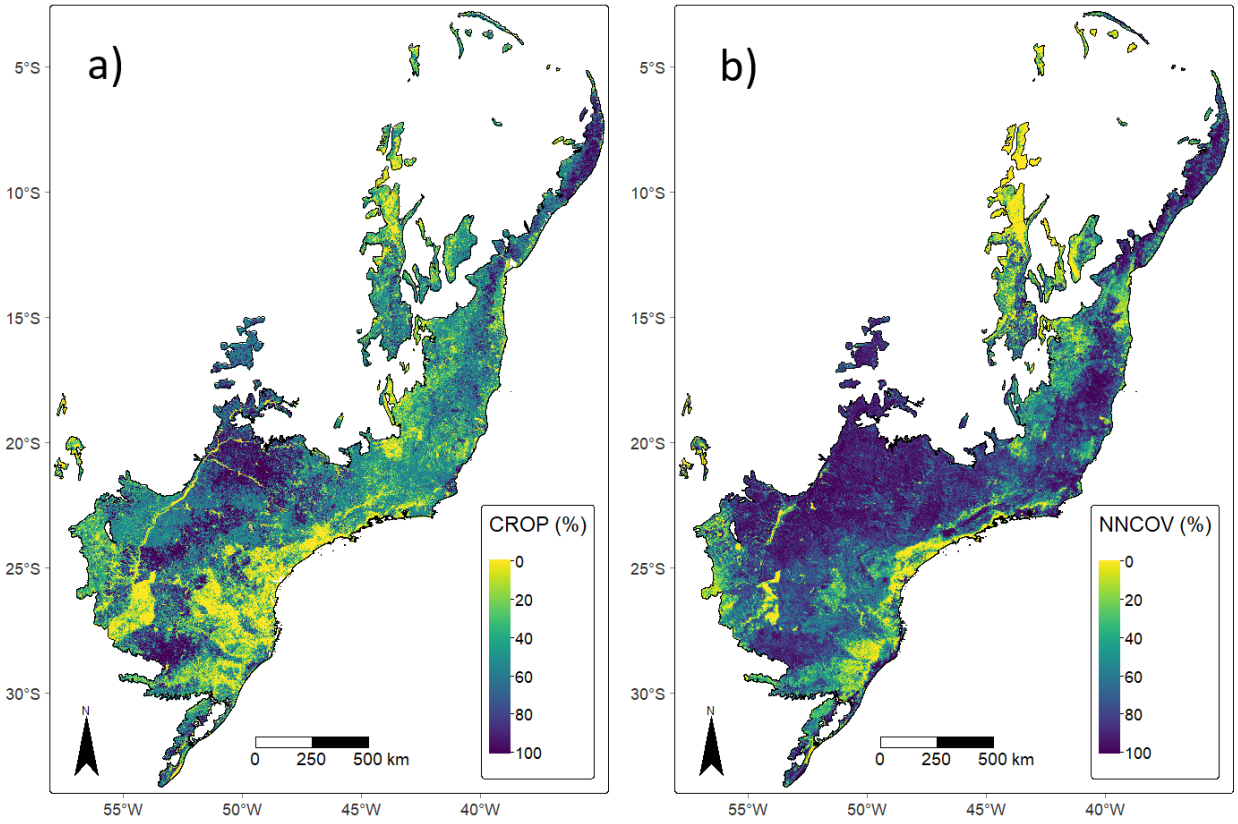


Figure 2.S2.1. Land-use surrogates across the Atlantic Forest. (a) Cropland prevalence (CROP), the percentage of cultivated and managed vegetation per 1-km cell from the global consensus land-cover product. (b) Non-natural cover (NNCOV), the percentage of transformed cover per cell, derived as 100 minus natural vegetation cover from the Atlantic Forest-specific 30-m annual land-cover product resampled to 1 km, so that higher values indicate greater transformation. Both panels show the native predictor surfaces before aggregation to response units, and colour scales run from low (dark) to high (light) values. The two surrogates are related but not interchangeable: cropland forms extensive contiguous tracts in the interior south and west, whereas non-natural cover additionally captures pasture, urban, mining and other anthropogenic classes, giving it broader coverage along the coast and in the north. Full provenance is given in Table 1.S2.2.

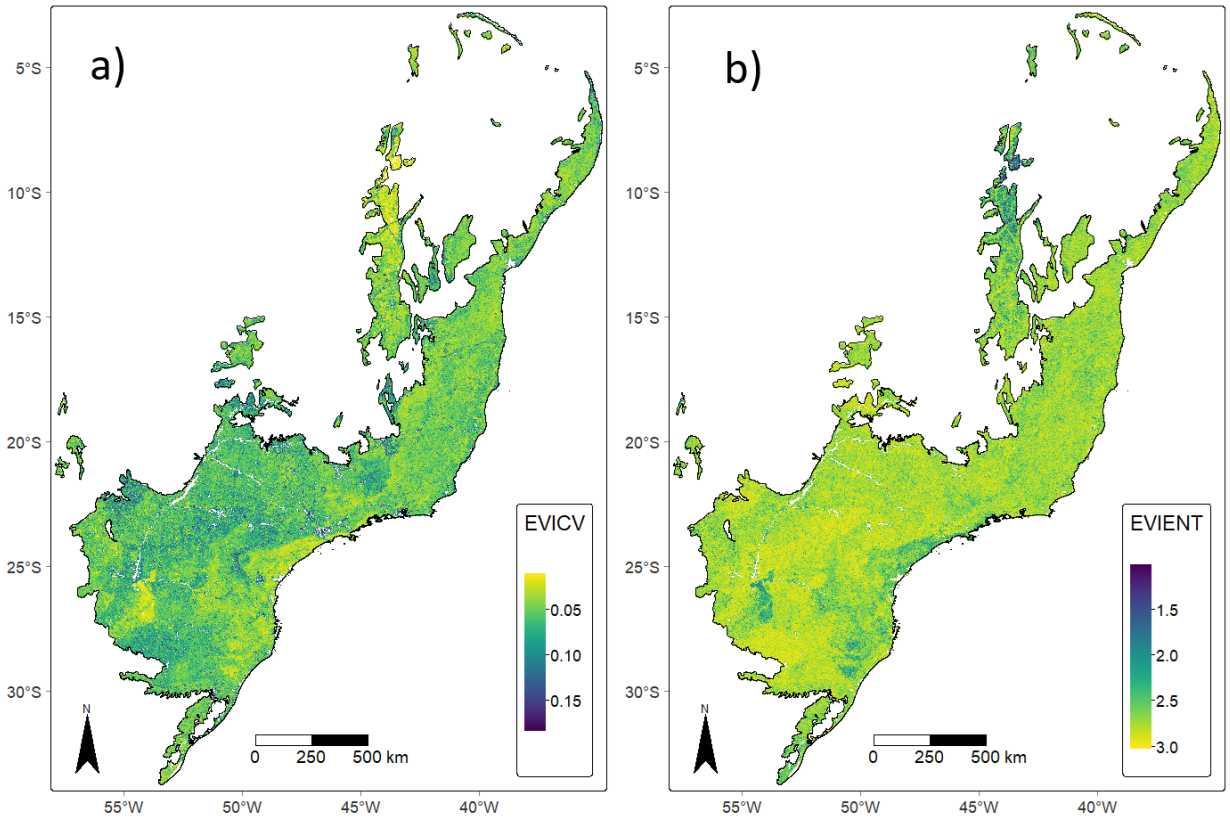


Figure 2.S2.2. Vegetational heterogeneity across the Atlantic Forest. (a) Coefficient of variation of the enhanced vegetation index (EVICV), a first-order texture metric indexing variability in the magnitude of vegetation greenness within the landscape window. (b) Entropy of the enhanced vegetation index (EVIENT), a second-order texture metric indexing the spatial disorderliness of how greenness values are arranged. Both derive from MODIS imagery at 1-km resolution for 2000–2005 and are shown as native surfaces before aggregation to response units. These two axes are land-use-responsive: greenness texture rises through degradation, secondary regrowth, agricultural mosaics and edge proliferation as well as through natural variation in vegetation state, so high values do not distinguish structurally intact heterogeneous forest from transformed mosaics. That ambiguity is why they are separated here from the persistent geophysical axes in Figure 2.S2.3, and why the chapter tests whether the origin of a heterogeneity axis determines the sign of its interaction with land use. Full provenance is given in Table 2.S2.1.

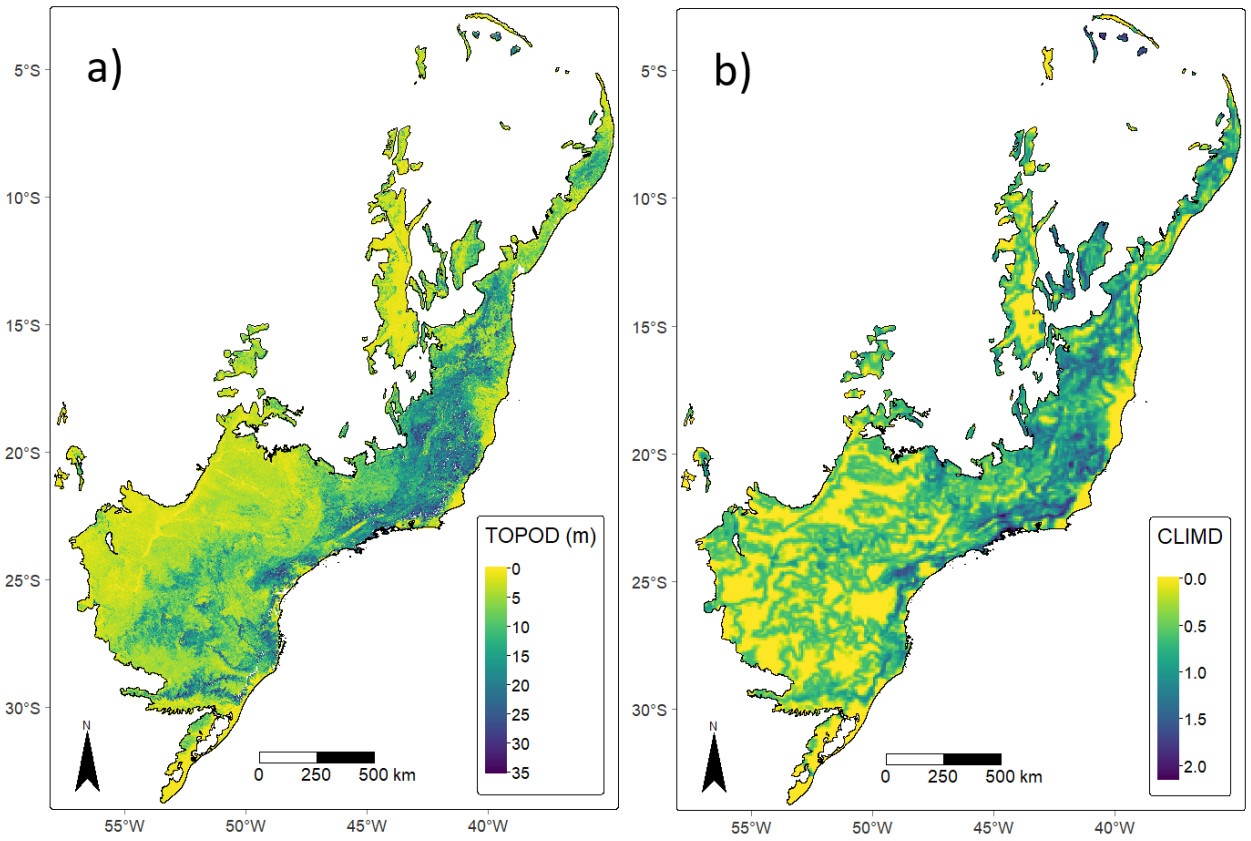


Figure 2.S2.3. Geophysical heterogeneity across the Atlantic Forest. (a) Topographic ruggedness (TOPOD), the terrain ruggedness index derived from GMTED2010-based 1-km topographic variables, indexing local relief. (b) Diversity of bioclimatic classes (CLIMD), the Shannon diversity of global environmental strata within the landscape window. Both are shown as native surfaces before aggregation to response units. Unlike the vegetational axes in Figure 2.S2.2, these describe a persistent physical template rather than a land-use-responsive state: they filter composition, maintain refugia and sustain beta diversity regardless of transformation history, and they are not altered by conversion. Their geographies are related but not identical, because relief generates climatic heterogeneity through orographic and lapse-rate effects while the converse does not hold, so the two are fitted in separate models and concordant support is read as one gradient recovered by two operationalisations rather than as independent confirmation. A third abiotic candidate, geomorphological landform diversity, was screened and retired because these two together span the construct. Full provenance is given in Table 2.S2.3.

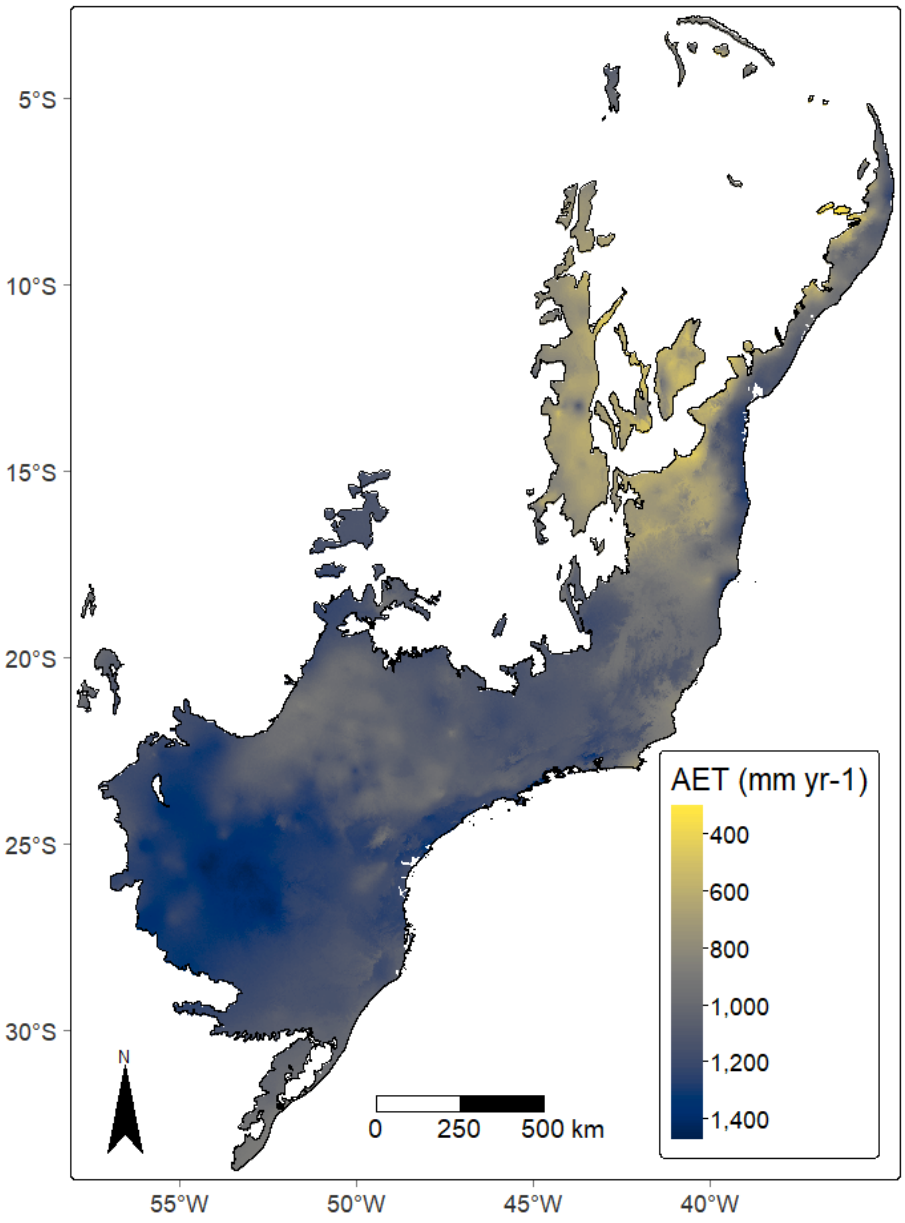


Figure 2.S2.4. Realised water–energy availability across the Atlantic Forest..Actual evapotranspiration (AET), derived before standardisation as the product of annual potential evapotranspiration and the Priestley–Taylor coefficient, indexing the water–energy flux actually realised given moisture availability. The surface is shown before aggregation to response units. AET carries three roles in this chapter: a focal predictor whose own richness association is estimated, a moderation arm in which land-use effects are tested against water–energy context, and a control against which heterogeneity moderations are conditioned, so that a surviving signal cannot be reduced to a productivity gradient. Full provenance is given in Table 2.S2.4

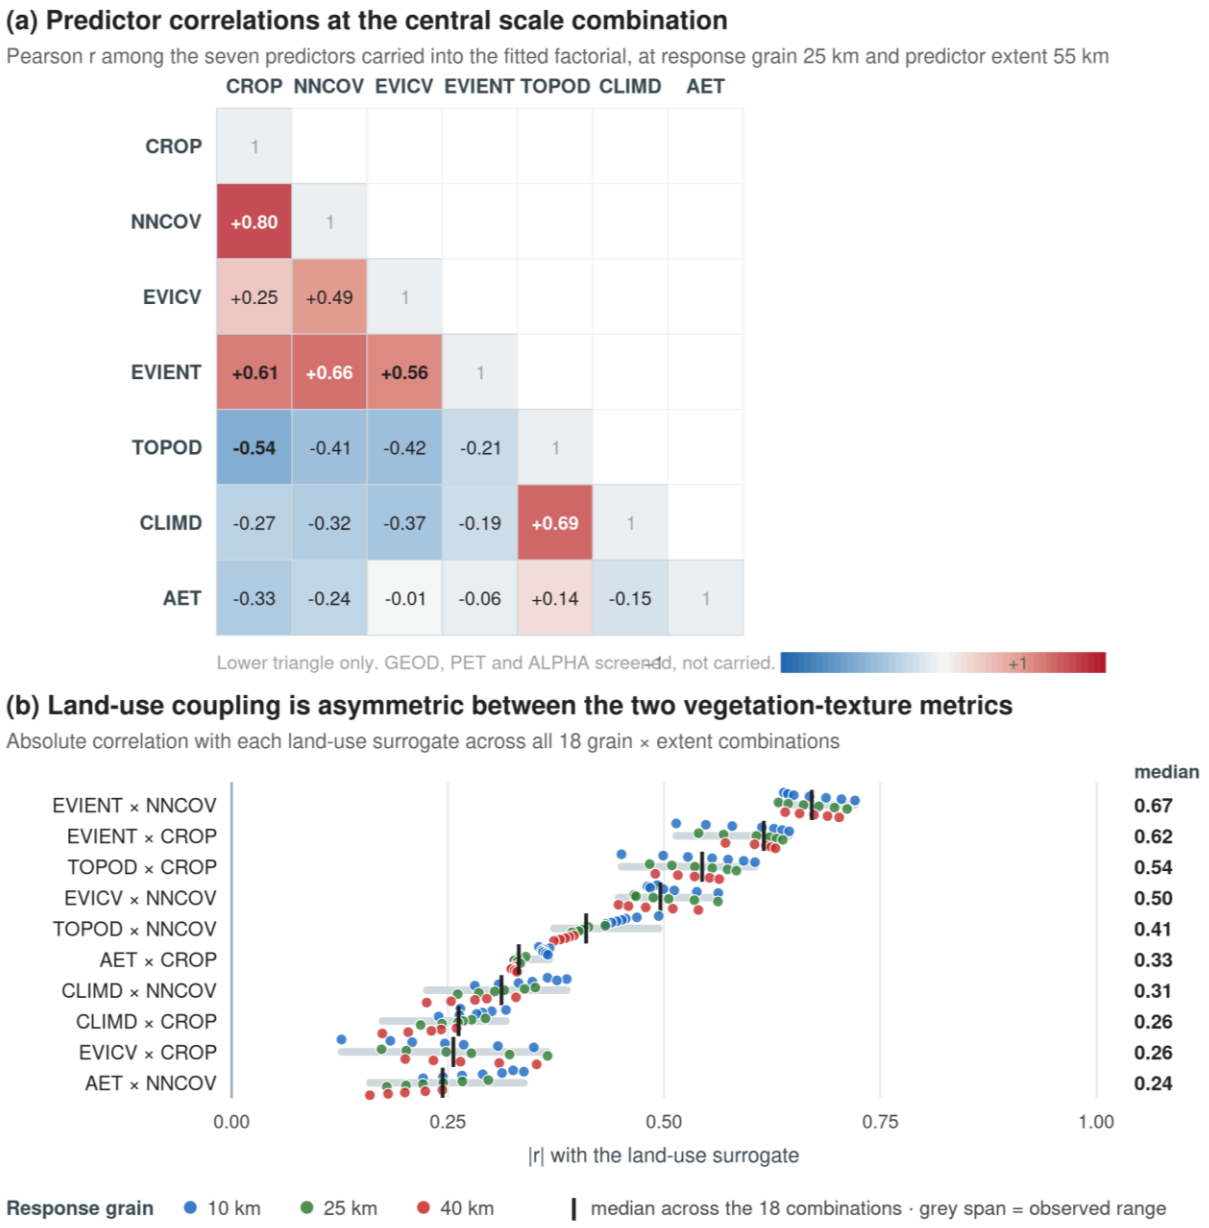


Figure 2.S2.5. Predictor geometry. (a) Pearson correlations among the seven predictors carried into the fitted factorial, at the central scale combination of response grain 25 km and predictor extent 55 km; the lower triangle only is shown, and GEOD, PET and ALPHA are screened out of the factorial and omitted. (b) Absolute correlation of each moderator with each land-use surrogate across all eighteen grain × extent combinations, ordered by median. Configurational diversity (greenness entropy, EVIENT) is moderately coupled to both land-use surrogates (median |r| 0.62 with cropland and 0.67 with non-natural cover) whereas greenness-magnitude

variability (the coefficient of variation of greenness, EVICV) is weakly coupled to cropland (0.26), so suppression risk is asymmetric between the two vegetation-texture metrics and a null result carries more suppression risk for greenness entropy than for the coefficient of variation. The two land-use surrogates are themselves strongly correlated at the central combination (r = +0.80), which is why they are never entered jointly. These are correlations among predictor values, not model coefficients, and not correlations with the response; variance inflation in the fitted models is reported in Section S2.2

APPENDIX 2.S3. Predictor roles and support-classification rules

The main-text tables state the predictor roles and the moderation hypotheses. This appendix expands the ecological reasoning behind those roles and sets out the rules by which we classify and validate interaction support. It does not restate the main-text tables.

2.S3.1 Expected predictor roles

Every predictor carries an expected role fixed before fitting, and each role rests on a stated ecological mechanism. Table 2.1 sets out those roles and the moderation hypotheses in tabular form. What follows expands the reasoning behind them rather than restating the table. The predictors divide into land-use surrogates, two origins of heterogeneity, and a productivity axis, and the central expectation is that heterogeneity of different origins moderates land-use effects in opposite directions.

The two land-use surrogates, cropland prevalence and non-natural cover, index habitat conversion. We expect negative main effects on richness across taxa and scales, filtered through matrix hostility, reduced forest extent and continuity, and resource simplification, and we expect these effects to strengthen with forest dependence across the guild strata. The two surrogates derive from different products and reference years, so agreement between the two land-use surrogates additionally tests robustness to product-specific and temporal idiosyncrasy.

Vegetational heterogeneity enters through two texture metrics computed from the enhanced vegetation index — a first-order compositional measure and a second-order configurational measure. Because vegetation texture is itself responsive to land use, degradation and fragmentation tend to raise apparent textural heterogeneity while removing the continuous forest cover on which forest-dependent species rely. We therefore expect vegetational heterogeneity to compound land-use damage — negative land-use × heterogeneity interactions — with the compositional and configurational metrics indexing proximate and broader dimensions of that texture.

Geophysical heterogeneity enters through terrain ruggedness and climatic heterogeneity. Both are persistent features of the physical template rather than land-use-responsive states, and both can sustain a wider set of persistence pathways under land-use pressure — terrain ruggedness through topographic microrefugia and refuge from conversion on steep ground, climatic

heterogeneity through a spread of climatic conditions and climatic refugia within the landscape. We therefore expect geophysical heterogeneity to buffer land-use damage — positive land-use × heterogeneity interactions — the opposite sign to the vegetational metrics, and the central contrast the chapter tests. Climatic heterogeneity carries a named alternative — an area–heterogeneity tradeoff, in which high climatic heterogeneity within a fixed area reduces the area available per climatic condition and can depress rather than raise richness — which we evaluate against the buffering expectation rather than assume away.

Energy enters through actual evapotranspiration, the realised water–energy flux, retained as the primary productivity axis and expected to associate positively with richness through demographic subsidy. Beyond its own association, we use actual evapotranspiration as a mechanistic control, so heterogeneity moderation that survives conditioning on realised productivity cannot be reduced to a productivity gradient alone.

Three further predictors are screened but not carried into the fitted factorial. Geodiversity is a relief-based metric that is redundant with terrain ruggedness once terrain ruggedness and climatic heterogeneity span the abiotic axes, and it is retired from the fitted and reported design. Potential evapotranspiration and the Priestley–Taylor moisture ratio are alternative operationalisations of the same water–energy system, and actual evapotranspiration is derived from them. Because the three proxies are not independent, we do not use them to define parallel downstream moderation arms; instead, we use them to test whether energy-conditioning conclusions depend on how water–energy structure is operationalised. We report the single-predictor screen for all three so that their exclusion rests on evidence rather than assertion.

2.S3.2 Support and robustness classification

Primary support requires two things at once, at a given grain and extent. The interaction must improve on the additive model by more than two AICc units, and the product term's structure coefficient must carry the same sign as the interaction coefficient itself. Meeting only one condition is not support. The structure coefficient is the correlation between the fitted linear predictor and the product term, and concordance guards against geometrically induced coefficients that carry the opposite sign to the association they purport to represent. We treat the likelihood-ratio test as a secondary consistency check only; an interaction supported by the likelihood-ratio test alone, or one that gains parsimony but fails structure-coefficient

concordance, is not primary-supported. The precomputed support flag in the deposited results records the ΔAICc criterion alone and diverges from the joint rule, so we recompute primary support from the parsimony gain and the concordance test directly rather than reading the flag.

We treat repeated recovery of the same interaction across the grain × extent matrix as scale coherence rather than as independent replication. Scale coherence is an interpretive pattern that raises confidence that a supported interaction reflects a genuine mechanism rather than a scale-specific artefact; it is not an additional per-cell support criterion, and we report scale-restricted support with its scale domain.

We assess robustness using controls matched to the ecological role of the focal moderator. For vegetational heterogeneity, we refit supported interactions with realised water–energy availability and the abiotic template as additive controls, because the enhanced-vegetation-index texture is the most downstream, land-use-responsive signal and is the most exposed to confounding by productivity and terrain. For geophysical heterogeneity, we evaluate survival after energy conditioning, because terrain ruggedness, climatic heterogeneity, and realised productivity can share positive richness gradients. We classify an interaction as controlled-supported only when it retains the parsimony gain of ΔAICc greater than two under the relevant control set and remains sign- and structure-coefficient-concordant under those controls, so that a fit-gain accompanied by geometric failure is never reported as robust.

We also record how surfaces were set aside. Main-text Table 2.2 names three grounds for rejection, of which only structure-coefficient concordance is a computed field: direction, parsimony gain and product-term concordance are stored per cell, and no operational test named directional coherence or non-forest specificity was implemented. Those two labels describe comparisons applied at interpretation rather than gates evaluated during fitting, and we state them as such rather than presenting them as computed criteria. Directional coherence refers to agreement of the interaction sign among eligible forest-specialist cells within a surface; the specificity comparison refers to reversal of that sign in the non-forest stratum on the same predictor. Both are read from the stored per-cell direction and stratum fields, both are post hoc, and neither carries a frozen threshold, so we use them to describe why a surface was not carried forward and not as an affirmative test. The resolved per-cell evidence for every surface, including those set aside, is displayed in Appendix 2.S5. Alongside controlled support we

distinguish three cases. A robust interaction retains the controlled parsimony gain with concordance intact. An absorbed interaction retains its direction and concordance but loses the controlled parsimony gain, indicating that the controls absorb its support without reversing it. An explained interaction reverses sign or loses concordance under the controls, indicating that a control accounts for the apparent moderation.

Because the controls enter jointly rather than one at a time, and because realised productivity and the abiotic template are themselves correlated, the controlled analysis tests whether a moderation survives conditioning rather than decomposing it into productivity and terrain contributions — a joint control cannot rank the two, and we make no such ranking claim. We therefore report survival under conditioning, not a decomposition of the moderation into its putative sources.

Finally, we assess whether the survival of an interaction under energy conditioning depends on which energy proxy conditions it. For each supported interaction we refit the conditioning step with potential evapotranspiration and with the Priestley–Taylor moisture ratio in place of actual evapotranspiration — one proxy at a time, since the three are collinear and cannot enter a single model — and we record whether the interaction retains its support under each swapped conditioner. Proxy-invariant survival indicates that the not-a-productivity-gradient reading is robust to how energy is operationalised; proxy-dependent collapse indicates that the reading is fragile and we report it as such. This is a test of energy-conditioning robustness and not a test of the more-individuals hypothesis — the design carries no abundance mediator, so it cannot isolate a more-individuals pathway, and we do not claim to.

Table 2.S3.1. Criteria used to classify and to set aside interaction surfaces. Only the first three are computed during model fitting and stored per grain × extent cell. Directional coherence and non-forest specificity are comparisons applied at interpretation, read from the stored interaction sign and response stratum; neither was implemented as a test in the analysis and neither carries a frozen threshold, so both describe why a surface was not carried forward rather than constituting an affirmative test.

Criterion	**Definition**	**Stored per cell**	**Status**	**Threshold**
Parsimony gain	Improvement in AICc of the interaction model over the matched additive model with the same random-effect and polynomial structure.	ΔAICc	Computed during fitting	ΔAICc greater than 2
Structure-coefficient concordance	Agreement in sign between the product-term coefficient and the product-term structure coefficient, the correlation between the product term and the fitted values.	Product-term coefficient and structure coefficient	Computed during fitting	Same sign
Primary support	Both of the above met in the same cell. This is the support rule used throughout.	Derived from the two fields above	Computed during fitting	Both met

Criterion	Definition	Stored per cell	Status	Threshold
Directional coherence	Agreement in the sign of the interaction coefficient among the eligible cells of the forest-specialist stratum.	Interaction sign and stratum	Read at interpretation	None fixed
Non-forest specificity	Reversal of the interaction sign in the non-forest stratum relative to the forest-specialist stratum on the same predictor pair.	Interaction sign and stratum	Read at interpretation	None fixed

APPENDIX 2.S4. Model specification, model-building diagnostics, and adequacy

This appendix documents the model formula and random-effect structure, the baseline specification decisions, the polynomial and random-slope (Barr et al. 2013) selection, and residual adequacy. Selection and adequacy summaries follow. All tabulated values are drawn from the deposited code and data archive.

2.S4.1 Model formula and random-effect structure

Richness is modelled as Poisson (Brooks et al. 2017) with a log link, and every observation carries its own random intercept so that overdispersion is absorbed within the likelihood rather than by switching to an alternative count family. Writing LU for a land-use surrogate and M for a moderator, the largest candidate interaction model is

$$\text{resp} \sim \text{EFFsb} + \text{LU} + \text{M} + \text{LU:M} + \text{space}$$
$$+ (1 \mid \text{method}) + (1 + \text{LU} + \text{M} \mid \text{taxa}) + (1 \mid \text{obs})$$

Sampling effort, the two focal terms, their product and a spatial autocovariate enter as fixed effects. Three random components follow — an intercept for sampling method, taxon-level slopes on LU and M shown here in their fullest correlated form, and the observation-level intercept carrying overdispersion. Random-slope structure is not fixed in advance. It is selected per model from the candidate set described in Section S4.4, so the fitted structure varies across models, and distributional adequacy is judged from scaled quantile residuals (Section S4.5).

The spatial term is derived rather than assumed. What this captures is the broad-scale geographic structure the model would otherwise leave in its residuals, so reported coefficients are conditioned on spatial dependence.

In this chapter LU ranges over the two land-use surrogates and M over the five focal predictors — two vegetational heterogeneity metrics, two geophysical metrics, and actual evapotranspiration — across the eighteen grain × extent combinations. Two comparison models accompany each interaction fit. Dropping the product term yields the additive baseline against which support is measured, and adding the control set matched to the focal moderator yields the controlled fit behind the controlled-support assessment. Neither should be confused with the moderator-free additive models behind the main-effect contrasts, which omit M entirely.

2.S4.2 Baseline specification decision audit (effort and year)

We fix the effort covariate and the treatment of survey year a priori and confirm the choices by a two-axis robustness audit. Throughout, we work within an information-theoretic model-selection framework: we compare candidate polynomial terms and random-slope structures by AICc and, where candidates fall within two AICc units of one another, retain the simpler structure (Burnham and Anderson 2002). We compare the retained effort formulation against an alternative unbounded formulation, and the year-free specification against one that adds survey year, and we report the direction agreement and the magnitude ratio of the focal coefficients between the paired specifications, so that the baseline choices rest on demonstrated insensitivity rather than assertion.

The effort audit contrasts the retained site-based effort formulation (EFFsb) against an unbounded alternative across all 18 grain × extent combinations in the total_full and forest-specialist audit strata. Univariate coefficients agree in sign at 96–100% of predictor–scale combinations, and the controlled reported estimates — the AET + TOPOD-conditioned coefficients that carry the reported main-effect inference — agree in sign at 96–100% with magnitude ratios of 0.95–0.98 across strata, so the effort formulation shifts reported magnitudes by at most a few percent without reversing any sign. The year audit contrasts year-retained against year-aggregated responses at the 10-km grain, the only grain permitting yearly disaggregation; sign is stable throughout. Because direction is stable across both audits, the reported interaction inference does not depend on the effort or temporal-aggregation choices, and we adopt the site-based, year-aggregated specification.

2.S4.3 Polynomial-term structure

We allowed polynomial (quadratic) terms for the focal predictors during model building and retained them only where model comparison supported them. Across the 900 fitted interaction models (180 per stratum), the selected structure carries a land-use polynomial in 432 and a moderator polynomial in 207 (Table 2.S4.1) — land-use polynomials in 16–72% of cells across strata and moderator polynomials in 8–39% — not the 0–2% figure carried in earlier drafts. Land-use polynomials are most frequent in forest-facultative richness (129 of 180) and the aggregate strata, and least frequent in non-forest richness (28), whereas moderator polynomials concentrate in non-forest richness (71) and forest specialists (56). The concentration of

polynomial terms on the land-use surrogates is consistent with threshold-like richness responses to land conversion.

Table 2.S4.1. Selected polynomial structure by stratum (counts over the 180 interaction models per stratum). From the deposited interaction-results table.

Stratum	Interaction models	Land-use polynomial	Moderator polynomial
total_full	180	100	42
total_restr	180	85	14
specialist	180	90	56
facultative	180	129	24
non_forest	180	28	71
All	900	432	207

2.S4.4 Random-slope variant selection and convergence

Table 2.S4.2. Selected random-slope structure by response stratum, counted over the 180 selected cells per stratum drawn from the 720 additive fits.

Selected random-slope structure	total_full	total_restr	specialist	facultative	non_forest
Intercept only (v1_none)	42	12	0	76	0
Land-use slope (v2_lu_only)	104	137	80	23	126
Moderator slope (v3_het_only)	22	20	2	76	17

Selected random-slope structure	total_full	total_restr	specialist	facultative	non_forest
Correlated land-use + moderator (v4)	12	11	98	5	37
Total	180	180	180	180	180

Among-taxon response heterogeneity is itself structured by forest dependence. The forest-specialist stratum selects correlated slopes on both terms in 98 of 180 cells, by a wide margin the highest of any stratum, so taxa diverge in their land-use and heterogeneity responses together precisely where forest dependence is greatest. Non-forest richness selects the land-use slope alone in 126 of 180 cells and correlated slopes in only 37, and forest-facultative richness divides between intercept-only and moderator-slope structures with land-use slopes in just 23. The correlated structure is therefore rarely selected outside the forest-specialist stratum yet is the modal structure within it, which is why it is retained in the candidate set throughout.

2.S4.5 Residual diagnostic adequacy

We assess distributional adequacy with scaled quantile residuals (DHARMa; Hartig 2022) computed from the fitted models, testing uniformity (Kolmogorov–Smirnov), dispersion, and outlier frequency at $\alpha = 0.05$ (Table 2.S4.3). Dispersion is adequate in every model of every stratum, so the observation-level random effect (Harrison 2014) captures overdispersion as intended, and outlier frequency is adequate in most models. Uniformity separates the aggregate from the guild-restricted strata: the test is largely adequate in the full total but departs moderately in forest specialists, forest-facultative, and non-forest richness. A representative refit under negative-binomial and zero-inflated alternatives rejected uniformity identically, so the departure is family-invariant and does not support replacing the Poisson observation-level-random-effect baseline. Because DHARMa uniformity tests are amplified by sample size, we judge the departure by the magnitude of the KS statistic rather than the p-value alone: the specialist residuals depart above a well-behaved reference but below the level that accompanies genuine dispersion or zero-inflation problems. We assess residual spatial independence with Moran's I

computed using the *ape* package (Paradis and Schliep 2019), comparing the non-spatial fit with the autocovariate-adjusted fit in every univariate screen cell of every stratum (Table 2.S4.4). The autocovariate reduces Moran's I in 100% of cells in all five strata. After adjustment no cell rejects spatial independence in four strata. Forest specialists are the exception. They carry the highest pre-adjustment autocorrelation and retain 20 of 180 cells at $\alpha = 0.05$ after adjustment, concentrated in the land-use and energy families rather than the heterogeneity moderators, at a mean residual Moran's I of +0.008. We therefore treat residual spatial dependence as substantially but not completely removed in the specialist stratum and read specialist results with that limitation in view. These diagnostics come from the univariate screen on the site-based-effort, temporally aggregated configuration, and the spatial autocovariate is retained in every model regardless of the test, so they are descriptive rather than a selection criterion. We therefore retain the Poisson-OLRE specification and report the guild-stratum uniformity departure as a moderate, unattributed limitation.

One aspect of sampling intensity is not adjusted for: survey duration. Records combining several years of survey enter on the same footing as records from a single period, because the modelling data retain reported effort and the number of consolidated assemblages but not the temporal span behind them.

Table 2.S4.3. Residual-diagnostic adequacy by response stratum (scaled-residual tests at $\alpha = 0.05$). Dispersion is adequate in 100% of models across strata; uniformity departs moderately in the three guild-restricted strata. All models fit without error.

Response stratum	Models	Uniformity (KS)	Dispersion	Outlier
total_full	180	adequate (aggregate)	100%	adequate
total_restr	180	intermediate	100%	adequate
specialist	180	departs (moderate)	100%	adequate
facultative	180	departs (moderate)	100%	adequate

Response stratum	Models	Uniformity (KS)	Dispersion	Outlier
non_forest	180	departs (moderate)	100%	adequate

Table 2.S4.4. Residual spatial autocorrelation before and after the spatial autocovariate, by response stratum and predictor family. Moran's I is computed with the *ape* package and averaged over the univariate screen cells within each family, 180 per stratum across ten screened predictors and eighteen grain × extent combinations. Values are means of full-precision cell estimates rounded once. Counts give the cells rejecting spatial independence at α = 0.05. From the deposited univariate spatial-diagnostic table.

Stratum	Family	Cells	Moran's I before	Moran's I after	p < 0.05 before	p < 0.05 after
total_full	LU	36	+0.021	+0.003	29	0
total_full	HET	90	+0.023	+0.002	75	0
total_full	energy	54	+0.024	+0.003	54	0
total_full	all	180	+0.023	+0.003	158	0
total_restr	LU	36	+0.022	+0.003	24	0
total_restr	HET	90	+0.026	+0.001	67	0
total_restr	energy	54	+0.025	+0.003	49	0
total_restr	all	180	+0.025	+0.002	140	0
specialist	LU	36	+0.083	+0.018	36	14
specialist	HET	90	+0.086	+0.003	90	0

Stratum	Family	Cells	Moran's I before	Moran's I after	p < 0.05 before	p < 0.05 after
specialist	energy	54	+0.097	+0.009	54	6
specialist	all	180	+0.089	+0.008	180	20
facultative	LU	36	+0.023	+0.006	28	0
facultative	HET	90	+0.027	+0.005	82	0
facultative	energy	54	+0.025	+0.007	54	0
facultative	all	180	+0.026	+0.005	164	0
non_forest	LU	36	+0.042	+0.002	36	0
non_forest	HET	90	+0.038	+0.002	90	0
non_forest	energy	54	+0.040	+0.001	54	0
non_forest	all	180	+0.039	+0.002	180	0

2.S4.6 Control-sensitivity audit of the interaction layer

Adding conditioning variables can redistribute variance shared with a focal term and shift its coefficient. We therefore audited the interaction layer for the process-specificity of cancellation between the guild strata rather than interpreting any single controlled coefficient in isolation. The decomposition runs on the uncontrolled additive layer. Each predictor enters several additive models — eight for a moderator, across two land-use pairings and four random-slope variants, and twenty for a land-use surrogate across five moderators and the same four variants — so we resolve the guild coefficient as the median across those models. The spread across models is substantial and we report it as a diagnostic rather than absorbing it: the median across-model range is 0.087 and the maximum is 0.335. Guild coefficients are weighted by each guild's share

of mean restricted-total richness at the response grain, and the cancellation fraction is one minus the absolute summed contribution divided by the summed absolute contributions, computed per cell and summarised as the median across the eighteen grain × extent combinations. Because the statistic is exactly zero whenever all three guild contributions share sign, a zero value records the absence of guild sign opposition rather than a small effect. Using the exact all-cell decomposition (Stage 2a), we quantified, for each land-use × moderator process, the fraction of the guild-level moderation signal that cancels when the guilds are aggregated — the cancellation spine of the chapter.

Cancellation is process-specific and concentrated in the two opposing-guild moderation processes. Greenness-magnitude heterogeneity (EVICV) cancels at 59% and realised energy (AET) at 37%, because forest-dependent and non-forest species respond to these gradients in opposite directions, so aggregate richness destroys most of the signal. By contrast, configurational heterogeneity (EVIENT) cancels at a median of 0% — a median-scale contrast rather than an absolute one, since EVIENT cancellation is zero in twelve of the eighteen cells but the guild contributions oppose in the remaining six, where the non-zero values decline monotonically with predictor extent within each response grain — terrain (TOPOD) at 19%, and climatic heterogeneity (CLIMD) at 12%, and the shared land-use damage terms (CROP, NNCOV) cancel at a median of 0% — guild contributions oppose in only 1 of 18 cropland and 7 of 18 non-natural-cover cells, so the direct cost of conversion survives aggregation essentially intact. The EVIENT 0% versus EVICV 59% contrast reproduces the greenness-magnitude discriminator in the aggregation layer: the same predictor family cancels oppositely because guilds oppose on greenness-magnitude variability in every cell but mostly agree on configurational diversity, opposing in only six of the eighteen combinations.

Table 2.S4.5. Aggregation-cancellation fractions by process (exact all-cell decomposition, Stage 2a): the fraction of the guild-level moderation signal lost when guilds are aggregated to total richness.

Process	Cancellation fraction	Interpretation
CROP (land-use damage)	0%	shared in 17 of 18 cells; survives aggregation
NNCOV (land-use damage)	0%	shared in 11 of 18 cells; survives aggregation
EVICV (greenness magnitude)	59%	opposing guilds; aggregate destroys most signal
AET (realised energy)	37%	opposing guilds; largely cancelled
TOPOD (terrain)	19%	low cancellation
CLIMD (climatic heterogeneity)	12%	low cancellation
EVIENT (configuration)	0%	guilds agree in 12 of 18 cells and oppose in 6

This audit is a diagnostic on the aggregation layer, not an independent confirmation of the guild results: it draws on the same data and model family as the interaction fits. It establishes that the selective blindness of aggregate richness is process-specific — near-total for the opposing-guild vegetational and energy moderation processes, near-zero for shared land-use damage — and it corroborates the greenness-magnitude discriminator reported in the main text.

2.S4.7 References

Barr DJ, Levy R, Scheepers C, Tily HJ (2013) Random effects structure for confirmatory hypothesis testing: keep it maximal. Journal of Memory and Language 68:255–278.

Brooks ME, Kristensen K, van Benthem KJ, et al (2017) glmmTMB balances speed and flexibility among packages for zero-inflated generalized linear mixed modeling. The R Journal 9:378–400.

Burnham KP, Anderson DR (2002) Model Selection and Multimodel Inference: A Practical Information-Theoretic Approach, 2nd edn. Springer-Verlag, New York.

Harrison XA (2014) Using observation-level random effects to model overdispersion in count data in ecology and evolution. PeerJ 2:e616.

Hartig F (2022) DHARMa: residual diagnostics for hierarchical (multi-level / mixed) regression models. R package version 0.4.6.

Paradis E, Schliep K (2019) ape 5.0: an environment for modern phylogenetics and evolutionary analyses in R. Bioinformatics 35:526-528.

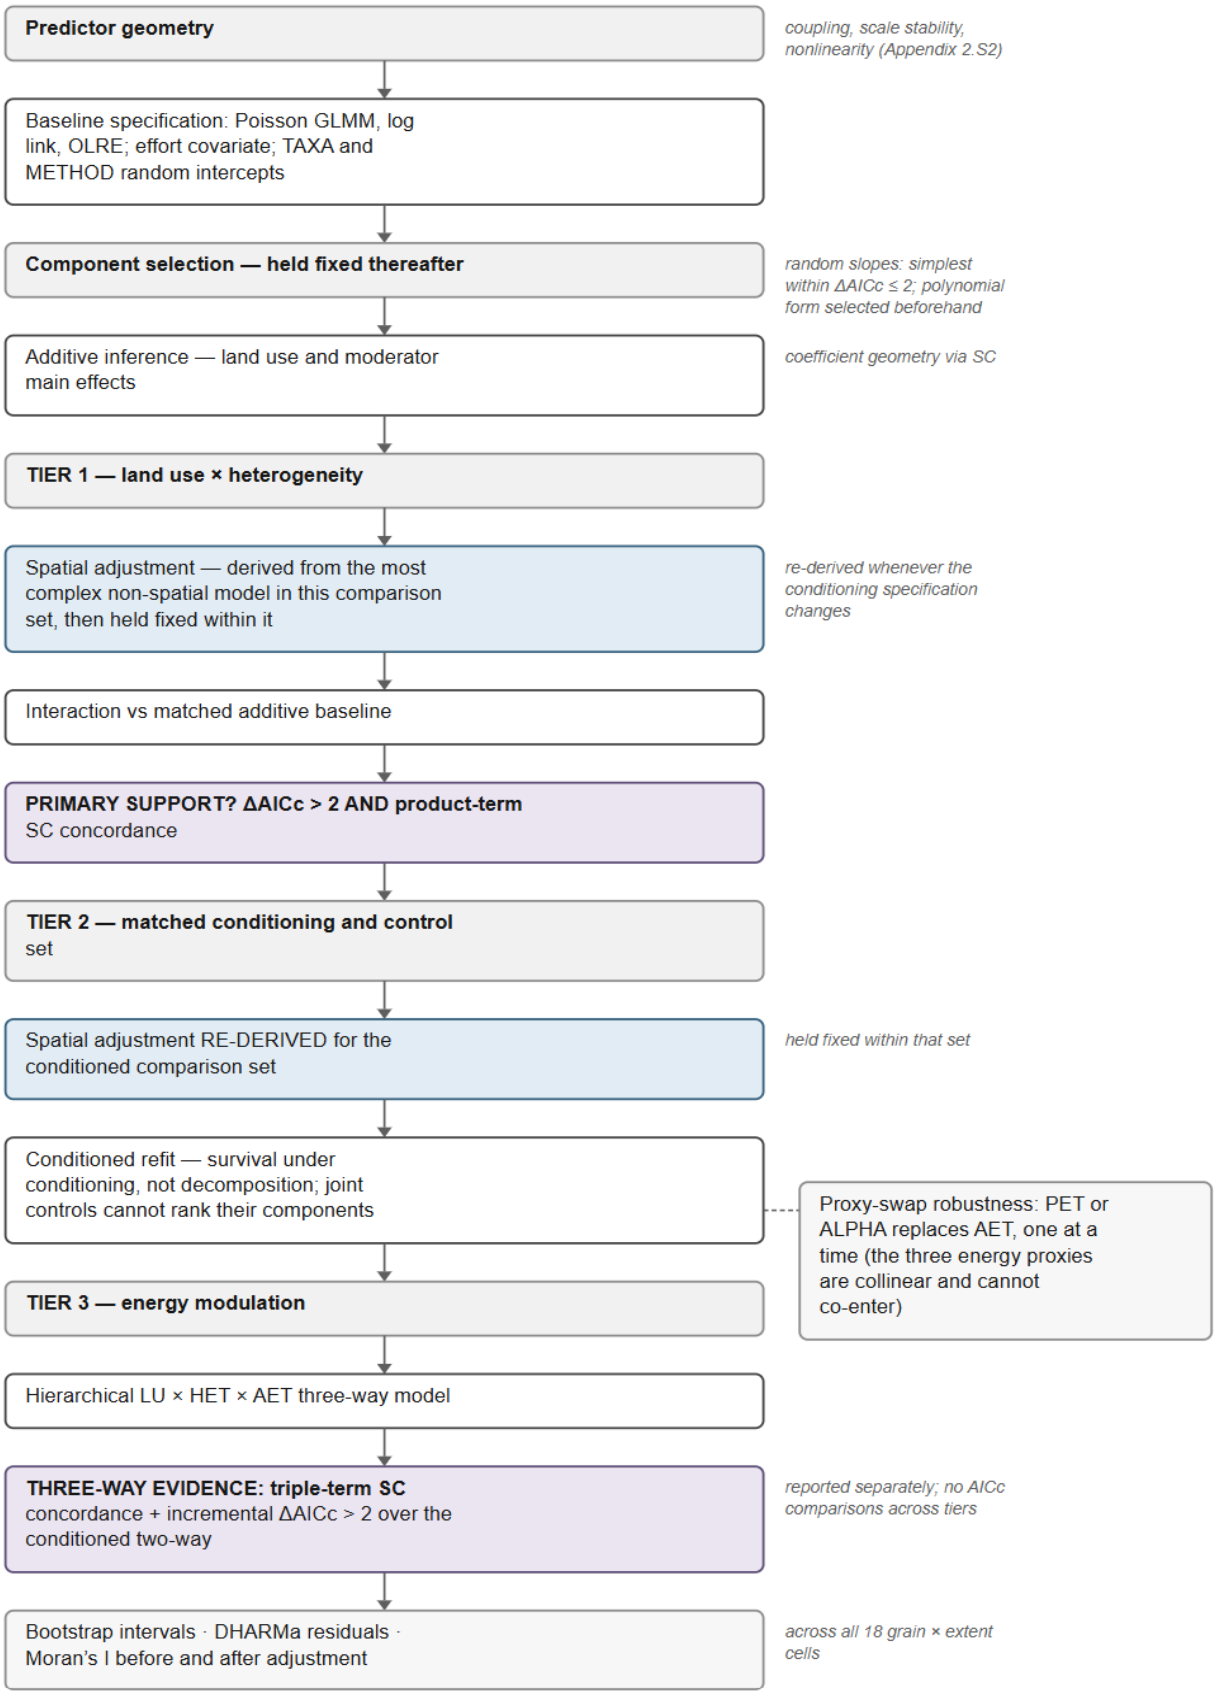

Predictor geometry
coupling, scale stability, nonlinearity (Appendix 2.S2)
Baseline specification: Poisson GLMM, log link, OLRE; effort covariate; TAXA and METHOD random intercepts
Component selection — held fixed thereafter
random slopes: simplest within ΔAICc ≤ 2; polynomial form selected beforehand
Additive inference — land use and moderator main effects
coefficient geometry via SC
TIER 1 — land use × heterogeneity
Spatial adjustment — derived from the most complex non-spatial model in this comparison set, then held fixed within it
re-derived whenever the conditioning specification changes
Interaction vs matched additive baseline
PRIMARY SUPPORT? ΔAICc > 2 AND product-term SC concordance
TIER 2 — matched conditioning and control set
Spatial adjustment RE-DERIVED for the conditioned comparison set
held fixed within that set
Conditioned refit — survival under conditioning, not decomposition; joint controls cannot rank their components
Proxy-swap robustness: PET or ALPHA replaces AET, one at a time (the three energy proxies are collinear and cannot co-enter)
TIER 3 — energy modulation
Hierarchical LU × HET × AET three-way model
THREE-WAY EVIDENCE: triple-term SC concordance + incremental ΔAICc > 2 over the conditioned two-way
reported separately; no AICc comparisons across tiers
Bootstrap intervals · DHARMa residuals · Moran's I before and after adjustment
across all 18 grain × extent cells

Figure 2.S4.1. Model-building and inference sequence for the heterogeneity-moderation analysis. Tier 1 evaluates land-use × heterogeneity support, Tier 2 evaluates robustness under the matched conditioning set, and Tier 3 tests whether actual evapotranspiration modulates a supported two-way interaction. For each nested comparison set, the residual spatial autocovariate is derived from the most complex corresponding non-spatial model and held fixed within that set; it is re-derived when the conditioning specification changes. PET and ALPHA enter only as one-at-a-time proxy swaps in the energy-conditioning robustness analysis, not as parallel three-way moderators. AICc comparisons are made within, not across, tiers.

APPENDIX 2.S5. Land-use × heterogeneity interaction results and bootstrap validation

This appendix reports the interaction support matrix, support under additive controls, the energy-conditioning robustness check, and the bootstrap validation, together with the provenance of the main-text interaction table. All values are drawn from the deposited code and data archive.

2.S5.1 Support criterion and validation

We classify an interaction as primary-supported when it satisfies the joint criterion set out in Appendix 2.S3 — a parsimony gain of ΔAICc greater than two over the additive model together with product-term structure-coefficient concordance — and we recompute this from the parsimony gain and the concordance test rather than reading the precomputed ΔAICc-only flag. We report scale-restricted support with its scale domain and treat repeated recovery across the grain × extent matrix as scale coherence rather than independent replication.

2.S5.2 Primary interaction support across strata

The primary support matrix reports, for each response stratum, the land-use × moderator combinations that meet the support criterion of Appendix 2.S3 — a parsimony gain of ΔAICc > 2 over the additive model together with product-term structure-coefficient concordance, recomputed from the parsimony gain and the concordance test rather than read from the precomputed flag (Table 2.S5.1). Across the five strata, 115 of the 900 grain × extent cells are primary-supported. Support concentrates in forest specialists and in the non-forest control, and the two carry opposite-signed surfaces: the cropland × greenness-magnitude (CROP × EVICV) interaction is supported at 15 of 18 specialist cells, the non-natural-cover × energy (NNCOV × AET) interaction at 12 of 18 specialist cells but 17 of 18 in the non-forest control with reversed sign, and the cropland × climatic-heterogeneity (CROP × CLIMD) interaction at 8 of 18 specialist cells. The table lists every combination reaching at least one supported cell, together with reference rows at zero support; it is not the full 50-row stratum × surrogate × moderator matrix. The complete cell-resolved matrix, in which untested, non-computable and unsupported cells remain distinguishable, is given in the resolved-support atlas (Figure 2.S5.1).

Table 2.S5.1. Five-stratum primary interaction support — count of the 18 grain × extent combinations meeting the support criterion (ΔAICc > 2 and product-term structure-coefficient concordance), by land-use × moderator family. Recomputed from the deposited interaction-results table; the precomputed support flag is ΔAICc-only and is not used. Zero-support tested combinations are shown as 0/18.

Stratum	Land-use × moderator	Primary-supported cells (of 18)
non_forest	NNCOV × AET	17/18
specialist	CROP × EVICV	15/18
non_forest	CROP × AET	13/18
specialist	NNCOV × AET	12/18
specialist	NNCOV × EVICV	9/18
specialist	CROP × CLIMD	8/18
total_full	CROP × EVICV	7/18
facultative	NNCOV × EVICV	5/18
total_restr	NNCOV × EVICV	4/18
non_forest	CROP × TOPOD	4/18
total_full	NNCOV × EVICV	3/18
total_restr	CROP × EVICV	3/18
specialist	NNCOV × EVIENT	3/18
total_full	NNCOV × CLIMD	2/18
facultative	NNCOV × CLIMD	2/18

Stratum	Land-use × moderator	Primary-supported cells (of 18)
total_full	CROP × EVIENT	1/18
total_full	CROP × CLIMD	1/18
total_full	NNCOV × TOPOD	1/18
total_restr	CROP × EVIENT	1/18
total_restr	NNCOV × CLIMD	1/18
specialist	CROP × EVIENT	1/18
facultative	CROP × CLIMD	1/18
non_forest	CROP × CLIMD	1/18
total_full	NNCOV × AET	0/18
total_restr	CROP × CLIMD	0/18
total_restr	NNCOV × AET	0/18
facultative	CROP × EVICV	0/18
facultative	NNCOV × AET	0/18
non_forest	CROP × EVICV	0/18

2.S5.3 Support under additive controls

The controlled analysis refits every supported interaction with actual evapotranspiration (AET) and terrain ruggedness (TOPOD) as additive controls, one and then jointly, and reports how many primary-supported cells retain a parsimony gain of ΔAICc > 2 with structure-coefficient concordance under conditioning (Appendix 2.S3). The forest-specialist CROP × EVICV

interaction is robust to conditioning. All 15 primary-supported cells retain controlled support through the +TOPOD (V1), +AET (V2), and +TOPOD+AET (V3) steps. The interaction retains its negative sign throughout, but its magnitude responds differently to the conditioning axes. Terrain control alone attenuates the negative coefficient in 14 of 15 cells, whereas energy control and joint terrain-and-energy control make it more negative in all 15 cells. This is the load-bearing surface of the chapter, and neither conditioning axis explains it away.

Table 2.S5.2. Forest-specialist controlled interaction support detail (count of grain × extent combinations retaining ΔAICc > 2 with structure-coefficient concordance under each control step). From the deposited control-ladder table (steps V1–V3). Full five-stratum matrices are in the deposited code and data archive.

Interaction (stratum)	Primary (of 18)	V1 +TOPOD	V2 +AET	V3 +TOPOD+AET
CROP × EVICV (specialist)	15	15/15	15/15	15/15

2.S5.4 Energy-conditioning robustness (proxy-swap)

We assess whether the survival of each energy-conditioned interaction depends on which energy proxy conditions it. For each supported interaction we refit the conditioning step with potential evapotranspiration (PET) and with the Priestley–Taylor moisture ratio (ALPHA) in place of actual evapotranspiration, one proxy at a time, and record whether the interaction retains structure-coefficient concordance under each swapped conditioner. This is a test of energy-conditioning robustness and not a test of the more-individuals hypothesis, because the design carries no abundance mediator. The NNCOV × AET surface is directionally insensitive to the proxy: it retains concordance at every specialist cell (12 of 12) and at every non-forest control cell (17 of 17), with the sign split preserved between the two strata. Because PET and ALPHA enter the construction of AET, we read this as directional insensitivity to the energy-axis proxy rather than as an independent test of energy confounding. This section replaces the earlier land-

use × productivity geometry diagnostic, which reported PET and ALPHA as parallel interaction arms.

Table 2.S5.3. Proxy-swap survival for energy-conditioned interactions under potential evapotranspiration and the Priestley–Taylor moisture ratio substituted for actual evapotranspiration. From the deposited energy-robustness table (step E1).

Interaction (stratum)	Cells	SC-concordant under proxy swap
NNCOV × AET (specialist)	12	12/12
NNCOV × AET (non-forest control)	17	17/17

2.S5.5 Bootstrap validation

We validate the load-bearing interactions with a parametric bootstrap (B = 250, seed 2024). The bootstrap ran at several model levels, and the controlled vegetational level returned no usable replicates: all fifty-two target rows at that level completed with zero successful resamples, across every surface and stratum, so no controlled-level interval is available anywhere. Controlled support is therefore evidenced by the robustness ladder rather than by the bootstrap, and the resolved displays distinguish cells that were attempted without replicates from cells that were never attempted. For the levels that did resample, we computed per-model bootstrap intervals rather than averaging interval bounds across models, and we report their per-interaction envelopes. For each interaction we report the mean interaction coefficient across the supported cells, the median, the 95% interval envelope spanning the widest per-model bounds, and the count of cells whose interval excludes zero (Table 2.S5.4). The bootstrap differentiated the load-bearing interactions. The forest-specialist CROP × EVICV two-way excludes zero at all 15 cells with a mean coefficient of −0.084; its energy-buffering triple (CROP × EVICV × AET) excludes zero at 4 of 6 cells with a positive coefficient (+0.071), the opposite sign to the two-way, consistent with energy buffering of the specialist cost. The NNCOV × AET interaction is sign-

split: negative in specialists (−0.082, 10 of 12) and positive in the non-forest control (+0.089, 15 of 17). The CROP × CLIMD interaction excludes zero at all 8 eligible specialist cells (+0.106).

Table 2.S5.4. Parametric-bootstrap interval envelopes for the load-bearing interactions. Mean β and median β average the supported cells; the 95% CI envelope spans the widest per-model bounds; the exclude-zero count follows the main-text Table 2.2 denominator for each surface: L1 for the CROP × EVICV two-way, L3 for the CROP × EVICV × AET triple, primary-level L2c rows for NNCOV × AET, and the eligible conditioned cells for CROP × CLIMD. From the deposited bootstrap-summary table (B = 250, seed 2024).

Interaction (stratum)	Target rows	Mean β	Median β	95% CI envelope	Excl. 0
CROP × EVICV two-way (specialist)	15	-0.084	-0.079	[-0.157, -0.002]	15/15
CROP × EVICV × AET triple (specialist)	6	+0.071	+0.071	[-0.010, +0.141]	4/6
NNCOV × AET (specialist)	12	-0.082	-0.085	[-0.186, +0.005]	10/12
NNCOV × AET (non-forest control)	17	+0.089	+0.091	[-0.014, +0.185]	15/17
CROP × CLIMD (specialist)	8	+0.106	+0.106	[+0.038, +0.176]	8/8

From the same bootstrap we summarise predictive adequacy. Among the 213 of 265 target rows with computable adequacy statistics, interval coverage is 1.000 and the predictive-calibration p-values fall within the 0.05–0.95 band at every one, indicating a well-calibrated bootstrap; the median relative root-mean-square error across all target rows is 0.406. These adequacy statistics complement the residual diagnostics in Appendix 2.S4. These coefficient summaries describe dependent scale combinations rather than a single pooled interaction parameter; the envelopes span per-cell bootstrap intervals across the scale matrix and are not confidence intervals for a

pooled effect. We therefore use them to document direction, scale coherence, and the provenance of Table 2.2 values, not to rank effect sizes across land-use surrogates or moderators.

2.S5.6 Provenance of the main-text interaction table

Table 2.S5.5 documents the provenance of the load-bearing rows of the main-text Table 2.2. For each, it reports the primary support count (n/18), the controls-robust count as the nested subset of primary-supported cells retaining support under conditioning (n/n_primary), the bootstrap mean coefficient with its 95% interval envelope, and the robustness class. This makes each load-bearing Table 2.2 value traceable to the scale combinations that produced it. The remaining Table 2.2 rows — NNCOV × EVICV, NNCOV × EVICV × AET and NNCOV × EVIENT — are reported at primary level only and are not carried as load-bearing results, so they are not traced here.

Table 2.S5.5. Provenance of the load-bearing rows of the main-text Table 2.2: per interaction, the primary support count, the nested controls-robust or bootstrap count, the bootstrap mean coefficient with its 95% envelope, and the robustness class. Bootstrap coefficients from the deposited bootstrap-summary table; support counts from the deposited interaction-results and control-ladder tables.

Interaction (stratum)	Primary (n/18)	Controls-robust (n/n_primary)	Bootstrap β [95% envelope]	Robustness class
CROP × EVICV two-way (specialist)	15/18	15/15	-0.084 [-0.157, -0.002]	robust (headline)
CROP × EVICV × AET triple (specialist)	—	4/6 boot	+0.071 [-0.010, +0.141]	energy buffering
NNCOV × AET (specialist)	12/18	—	-0.082 [-0.186, +0.005]	sign-split (specialist −)

Interaction (stratum)	Primary (n/18)	Controls-robust (n/n_primary)	Bootstrap β [95% envelope]	Robustness class
NNCOV × AET (non-forest control)	17/18	—	+0.089 [-0.014, +0.185]	sign-split (control +)
CROP × CLIMD (specialist)	8/18	—	+0.106 [+0.038, +0.176]	secondary (SC-weak)

The bootstrap coefficients reproduce the main-text Table 2.2 values; the complete per-combination lists are provided in the deposited code and data archive.

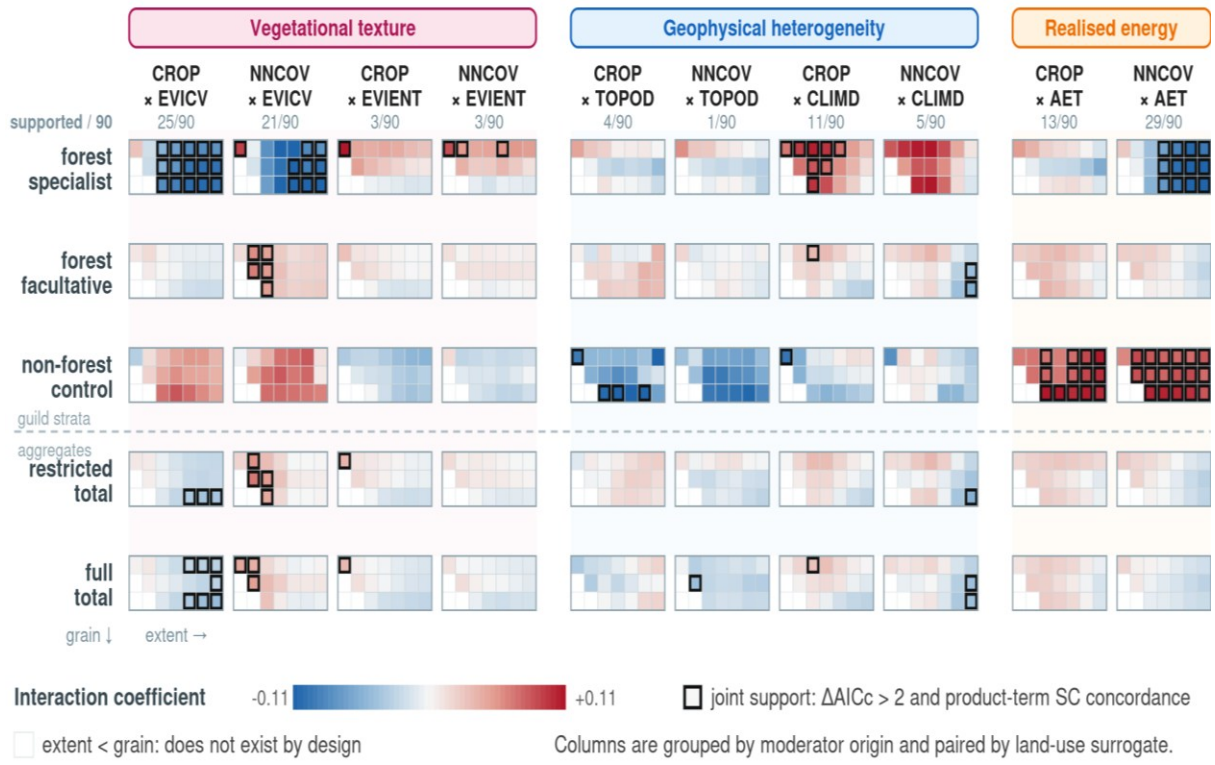


Figure 2.S5.1. Resolved interaction-support atlas. Columns are grouped by moderator origin — vegetational texture, geophysical heterogeneity and realised energy, the three origins of Figure 2.2b — and within each origin the two land-use surrogates are placed side by side on the same moderator, so any adjacent pair differs only in surrogate. Rows run down the forest-dependence gradient, with the two aggregate strata below the dashed rule. Each panel gives the three response grains (rows: 10, 25 and 40 km) by seven predictor extents (columns: 10 to 100 km) for one land-use × moderator surface in one response stratum; combinations where predictor extent is smaller than response grain do not exist by design and are left blank, accounting for three of the twenty-one positions in each panel. Fill gives the interaction coefficient and a heavy border marks cells meeting the joint support rule of ΔAICc greater than two together with product-term structure-coefficient concordance. Counts under each surface are supported cells of 90, that is 18 per stratum across five strata; 115 of the 900 cells are supported overall, no cell was non-computable, and unsupported cells retain their fitted coefficient. Read down the strata within a surface: on cropland × greenness-magnitude variability and on non-natural cover × realised energy the forest-specialist row is blue where the non-forest row is red, and both aggregate rows

are pale on exactly those surfaces. Read across the origin blocks to compare the vegetational and geophysical arms, and across adjacent columns within a block to compare the two land-use surrogates on the same moderator. Coefficients are from conditional Poisson generalised linear mixed models of untransformed species-richness counts, fitted with an observation-level random effect, and are on the log scale.

APPENDIX 3.S1. Diagnostic terminology used in Chapter 3. Thesis-wide terms, including estimand, habitat, habitat amount, configuration and fragmentation per se, are defined in Appendix 0.S2. This appendix defines terminology specific to the diagnostic framework developed here. The definitions are reading aids and do not replace the fuller methodological descriptions in the main text.

Additive model	A regression model entering habitat amount and configuration as separate terms, usually without an interaction; the concern here is interpreting its partial coefficients as independent ecological effects before predictor separability is shown.
Beta weight (BW)	The standardised partial regression coefficient; it measures a predictor's unique contribution to the partial slope.
Coefficient architecture	The joint pattern among marginal associations, partial slopes, fitted-gradient alignment, and unique variance contributions; used here to diagnose attenuation and amplification.
Collinearity	Statistical dependence among predictors; when strong, it prevents a model from attributing response variance to each predictor independently.
Domain stability	Whether a coefficient holds across the ecological gradient where it is interpreted, allowing that different mechanisms may operate in different domains.
Fitted-gradient alignment	The degree to which a predictor tracks the model's fitted response gradient, measured by the structure coefficient.
Geometric separability	Whether the predictor geometry actually supplies predictor variation independent enough for a partial coefficient to read as one predictor's effect rather than a split of a shared gradient; the prior condition in the Figure 3.1 procedure, which predictor-space adequacy evaluates.

Marginal association	A predictor's relationship with the response before other predictors are included, in contrast to its partial coefficient, which is estimated after they are included.
Metric–mechanism correspondence	Whether the fitted configuration metric represents the proposed mechanism rather than a broad proxy for several landscape changes.
Nonlinear coupling	A curved dependence between predictors that linear diagnostics such as correlation and VIF do not register.
Orthogonality	Statistical independence between predictors; the condition under which a partial coefficient reads cleanly as one predictor's effect.
Partial coefficient	A predictor's effect with the other predictors held fixed; it can approximate the estimand only when the predictors are separable in the realised data.
Partial-slope amplification	A coefficient inflated above a predictor's weak standalone association by residual coupling; the Case 2 failure.
Patch number	A count of habitat patches in a sampling unit; the subdivision-based configuration metric used in both cases.
Predictor geometry	The structure of the predictor space: attainable and sampled regions, the coupling between predictors, and any nonlinear or asymmetric dependence among them. Broader than linear collinearity, it is the structure that the realised predictor space instantiates, and predictor-space adequacy asks whether it supplies the same-amount, different-configuration contrasts the estimand requires.
Predictor-space adequacy	Whether the sampled landscapes contain the same-amount, different-configuration comparisons the estimand requires.

Raup-Crick beta diversity	A null-model dissimilarity that reduces the influence of alpha-diversity differences; the alpha-corrected response used as a check in the K26 case.
Realised predictor space	The actual combinations of predictor values present in the sampled data; for this estimand, whether the landscapes include similar habitat amount with differing configuration.
Residualisation	Removing from one predictor the part predictable from another; used here diagnostically, to see how a coefficient behaves as the shared habitat-loss gradient is reduced or reallocated, rather than as a correction.
Response–estimand alignment	Whether the biodiversity response measures the ecological process being interpreted.
Same-amount, different-configuration contrast	The comparison the fragmentation-per-se estimand requires: landscapes of comparable habitat amount but different spatial configuration.
Semi-partial R^2 (R^2sp)	The unique response variance a predictor explains beyond the other predictors.
Shared habitat-loss gradient	The single directed gradient along which habitat loss co-produces habitat amount and configuration, so the two predictors covary by construction rather than varying independently; the structural source of suppressor attenuation and amplification in both cases.
Structure coefficient (SC)	The correlation between a predictor and the model's fitted response; it measures alignment with the fitted gradient and is signed.
Subdivision	The division of remaining habitat into discrete patches during habitat loss; the process patch number indexes. Within a fixed sampling unit the quantities are algebraically tied (mean patch

	size equals habitat area divided by patch number), so a patch-count metric cannot vary fully independently of habitat amount.
Suppressor attenuation	A coefficient pushed toward zero (large SC, near-zero BW) by a coupled dominant predictor despite the subordinate predictor tracking the response; the Case 1 failure.
Total beta diversity	Undecomposed assemblage compositional dissimilarity among sites; one of the K26 response variables.
Triangulation	Testing whether a configuration result converges across responses, metrics, and spatial scales aligned to the proposed mechanism; convergence supports, and divergence undermines, a fragmentation-per-se interpretation.
Variance inflation factor (VIF)	A standard linear-collinearity diagnostic; a low value shows severe linear collinearity is absent but does not establish nonlinear predictor separability.

APPENDIX 5.S1. Chapter-specific response documentation

This chapter uses the shared Atlantic Forest assemblage compilation, forest-affiliation classification and multiscale response design described in the General Introduction and Appendix 0.S1. Here we document only chapter-specific response construction and diagnostics.

5.S1.1 Taxonomic scope and classified taxon pool

Species are classified by forest affiliation from species-level habitat records in the IUCN Red List, coded under the IUCN Habitats Classification Scheme (IUCN 2012), following the level-1 forest-class logic of Donald et al. (2019): species listed only for forest are forest specialists; species listed for forest together with one or more non-forest classes are forest facultatives; species with no forest habitat listed are non-forest. Only the five retained groups carry this coding. Species names were matched to the IUCN habitat records after standardization to the accepted binomial; seven mid–large mammal records resolved only to genus (Cabassous, Coendou, Dasyprocta, Didelphis, Galictis, Leopardus and Mazama) and these were retained with the forest affiliation shared across the congeners recorded in the Atlantic Forest, so that 1,192 of the 1,199 entries are resolved binomials and seven are genus-level. One further bat entry carries a provisional determination but is resolved to species and is counted among the 1,192. The classified pool is therefore counted in taxa rather than species, 1,192 of the 1,199 entries being resolved binomials.

Across the five groups carrying habitat-affiliation coding, 1,199 taxa are classified: 302 forest specialists (25.2 %), 702 forest-facultative taxa (58.5 %) and 195 non-forest taxa (16.3 %). The specialist fraction ranges from 14.1 % in anurans to 38.8 % in small mammals (Table 5.S1.1). Classified-pool totals are 495 for anurans, 428 for passerines, 116 for small mammals, 97 for bats and 63 classified taxa for mid- and large-bodied mammals, comprising 56 species-level and seven genus-level entries; that group's pool is therefore counted in taxa rather than species. For four of the five groups the classified pool matches the curated gamma diversity exactly; mid- and large-bodied mammals are the exception, their classified pool being smaller than the 70 resolved species in the curated gamma diversity. This does not affect the fitted models, because the responses are richness counts over the classified pool rather than species identities. These are pool-level proportions and carry no assemblage-frequency weighting; they are assembly-stage

quantities and are reported as such rather than recomputed here, the aggregated analysis files carrying assemblage richness rather than species identities.

Table 5.S1.1. Classified taxon pools and forest-affiliation composition. Values come from the curated classification; the aggregated analysis files contain assemblage richness rather than taxon-level composition. Percentages are calculated within each classified group and are not weighted by assemblage frequency

Animal group	Classified taxa	Forest specialist n (%)	Forest facultative n (%)	Non-forest n (%)
Anurans	495	70 (14.1)	318 (64.2)	107 (21.6)
Passerines	428	146 (34.1)	227 (53.0)	55 (12.9)
Small mammals	116	45 (38.8)	44 (37.9)	27 (23.3)
Bats	97	30 (30.9)	64 (66.0)	3 (3.1)
Mid–large mammals	63	11 (17.5)	49 (77.8)	3 (4.8)
All five groups	1,199	302 (25.2)	702 (58.5)	195 (16.3)

5.S1.2 Response construction

The chapter fits four response strata over a single restricted universe of the five habitat-classifiable groups: restricted total richness (total_restr) and the three mutually exclusive forest-affiliation guilds — forest-specialist (specialist), forest-facultative (facultative), and non-forest (non_forest) richness. The guilds partition the restricted total exactly:

total_restr = specialist + facultative + non_forest.

For each assemblage, facultative richness is derived as forest-associated richness minus specialist richness, and non-forest richness as restricted total minus forest-associated richness. Forest-associated richness (specialist + facultative) is computed only as the intermediate needed to derive the facultative guild; it is not fitted as a response, and its directly modelled aggregate

estimand is unavailable under the log link rather than statistically redundant. The build enforces non-negativity of every component and the exact partition identity before model fitting; the identity therefore holds by construction in the fitted data.

non_forest is treated as a specificity diagnostic, not a negative control: because its realized richness distribution and effective information content may differ from those of the other strata, its precision is evaluated explicitly (Section S1.3) rather than inferred from support frequency, and weak non_forest support is not interpreted as evidence of a zero effect.

5.S1.3 Stratum precision and effective information

The four strata differ in typical richness and in effective sample size, and these differences bear on the precision with which fragmentation effects can be estimated. Strata with lower mean counts, more zeros, or fewer informative taxon–method observations may yield wider intervals and lower power; the realized precision of each stratum is therefore evaluated empirically from its richness distribution, effective sample size, and fitted-model uncertainty rather than assumed a priori. Consequently, weaker or less frequent support in a low-richness stratum is not evidence of a null effect; it is a precision property that the cross-stratum comparisons (Appendix 5.S3, Appendix 5.S5) interpret alongside sign coherence and matched contrasts rather than as an independent vote. The distributions bear this out. Non-forest richness carries roughly a third zeros (30.3 to 35.7 % across grains) and a mean below three taxa, and forest-specialist richness carries between a fifth and a quarter zeros with a mean below five, whereas restricted-total richness has no zeros and a mean above twenty. The two low-richness strata therefore enter the analysis with materially less information per observation than the aggregate, which is why weaker or less frequent support in either is not read as evidence of a null effect.

Zero prevalence separates the strata more sharply than the means. At the 10 km grain, 35.7% of non-forest records and 24.6% of forest-specialist records are zero, against 2.6% for forest-facultative and none for the restricted total. Zero prevalence declines modestly at coarser grains, to 30.3% and 22.2% respectively at 40 km, because consolidation raises counts. The strata therefore differ not only in mean richness but in how often the response is absent altogether, which bears directly on the precision with which fragmentation effects can be estimated.

APPENDIX 5.S2. Forest definitions, fragmentation metrics, and predictor geometry

Definitions, provenance and design roles come from the analysis design and source documentation. Diagnostic quantities — correlations, variance inflation factors, nonlinear amount–configuration geometry and residualization checks — were recovered from the frozen predictor datasets, which are the surviving source for them. The frozen datasets retained forest cover, the raw and residualized fragmentation metrics and the environmental controls. Amount smooths were refitted only where fitted-shape summaries were unavailable, and the reconstructed residualized predictors were verified against the stored columns (Table 5.S2.4).

5.S2.1 Forest raster streams, provenance, and temporal anchors

Forest cover is represented by three raster streams derived from two products. The Hansen et al. (2013) Global Forest Change product enters at two resampled resolutions, 500 m and 1,000 m, and the Tuanmu and Jetz (2014) global consensus land-cover product (forest classes combined) enters at 1,000 m. Product provenance, native resolution, temporal anchor, and the resampling procedure are reported in Table 5.S2.1, and the three streams are mapped in Figure 5.S2.1, with the environmental controls in Figure 5.S2.2. The two products differ in temporal reference, classification framework, and native resolution; the temporal mismatch between biodiversity sampling and the forest products is treated as a design constraint typical of large-scale landscape studies rather than resolved here.

Table 5.S2.1. Forest raster streams: product, resolution, temporal anchor and provenance, with sources given in the final column

Stream	Product	Native / resampled resolution	Temporal anchor	Provenance / citation
Hansen 500 m	Hansen GFC	30 m / 500 m	percent tree canopy cover for the year 2000	Hansen et al. 2013

Stream	Product	Native / resampled resolution	Temporal anchor	Provenance / citation
Hansen 1000 m	Hansen GFC	30 m / 1,000 m	percent tree canopy cover for the year 2000	Hansen et al. 2013
Tuanmu 1000 m	Tuanmu & Jetz consensus	1,000 m	consensus of source products spanning the 1990s to the mid-2000s	Tuanmu & Jetz 2014

Processing note. Forest streams were built by resampling the continuous canopy-cover surface to the working resolution and applying the canopy threshold to the resampled surface, in that order. The native 30 m Hansen grid is computationally prohibitive across the full multi-scale design and is far finer than the coarsest structure the response grains can resolve, the smallest response unit being 10 km. Order matters and is stated because it defines what counts as forest at each resolution: thresholding after resampling removes forest fragments whose aggregated cover falls below the threshold, whereas thresholding first would have carried them through aggregation as fragmented structure. The resolution contrast between the Hansen 500 m and 1,000 m streams is therefore a contrast in the scale at which cover is averaged before classification, not in the scale at which patches are delineated. The Tuanmu–Jetz consensus product is native to 1 km and is thresholded directly. Patch delineation throughout uses the eight-neighbour rule, so cells meeting only at a corner belong to one patch; NLSI and CLUMPY take no neighbourhood argument and are computed from the adjacency matrix.

5.S2.2 Canopy thresholds and the nine forest definitions

Each stream is binarized into forest and non-forest at three canopy-cover thresholds — 70%, 80%, and 90% — yielding nine forest definitions (three streams × three thresholds). Strict thresholds (90%) emphasize structurally intact interior forest and exclude degraded vegetation; inclusive thresholds (70%) incorporate moderately degraded or recovering forest. The nine definitions are forest definitions, not nine products, and they support three distinct contrasts held fixed within design cell: a product contrast (Hansen versus Tuanmu at matched 1,000 m

resolution), a resolution contrast (Hansen 500 m versus 1,000 m), and a threshold contrast. No single definition is treated as correct; the definitions constitute a robustness dimension.

5.S2.3 Fragmentation metrics and orientation

Forest configuration is quantified with four class-level landscape metrics computed from the binary forest maps: DCLUMP, NLSI, EDGED, and PATCHD (Neel et al. 2004; Wang et al. 2014). Each captures a complementary aspect of configuration — dispersion, shape, edge, and patch density, respectively. All four are oriented so that larger values denote greater fragmentation, so that coefficient signs are comparable across metrics. Metric definitions and orientation are reported in Table 5.S2.2.

Table 5.S2.2. Fragmentation metric definitions and orientation (all oriented so larger = more fragmented).

Metric	Function	Definition	Units	Range and orientation	Transformation applied here
DCLUMP	lsm_c_clumpy	Proportional deviation of the observed proportion of like adjacencies from the proportion expected under a spatially random distribution, computed from the adjacency matrix by the double-count method	none	native CLUMPY spans −1 to +1 (−1 maximally disaggregated, 0 random, +1 maximally aggregated); the fitted predictor is z-standardized and therefore has no fixed bound	z-standardized, then negated: DCLUMP = −z(CLUMPY), so that higher values denote greater disaggregation. Negation reverses direction without changing the standardized scale

Metric	Function	Definition	Units	Range and orientation	Transformation applied here
NLSI	lsm_c_nlsi	$(e_i - \min e_i) / (\max e_i - \min e_i)$: total class edge length relative to its attainable minimum and maximum for the same class area	none	0 to 1; 0 for a single square patch, 1 for a maximally disaggregated checkerboard	z-standardized; already oriented so that higher values denote greater fragmentation
PATCHD	lsm_c_pd (directions = 8)	$n_i / A \times 10{,}000 \times 100$: number of patches of the focal class per unit landscape area	patches per 100 hectares	0 to 10^6 as a package-level bound; the attainable maximum here is set by cell area at 500 m and 1,000 m and is far below that bound. Increasing with patchiness	z-standardized; already oriented so that higher values denote greater fragmentation

Metric	Function	Definition	Units	Range and orientation	Transformation applied here
EDGED	lsm_c_ed (directions = 8, count_boundary = FALSE)	$\Sigma e_{ik} / A \times 10{,}000$: total edge length of the focal class per unit landscape area, with the landscape boundary excluded	metres per hectare	≥ 0; 0 with a single patch and no boundary edge, increasing without limit with patchiness	z-standardized; already oriented so that higher values denote greater fragmentation

5.S2.4 Metric redundancy and amount coupling

A fifth candidate metric, DISSAG, was excluded as highly redundant with NLSI; the reported redundancy range (of order r = 0.92–0.995 across scale combinations) comes from the contemporaneous predictor audit and was not recomputed here. The four retained metrics are not treated as cleanly orthogonal. Because fragmentation metrics can depend nonlinearly and often unimodally on forest cover, standard linear diagnostics may fail to capture amount–configuration coupling that persists after the linear association is removed. The appendix therefore reports retained-metric correlations, variance inflation factors, and the nonlinear amount–configuration geometry across scale combinations (Table 5.S2.3; Figure 5.S2.3), and addresses residual coupling through the residualized formulation (Section S2.5).

The four metrics differed in the shape and strength of their relationship with forest amount and in where residual variation remained along the cover gradient. Residualization reduced the fragmentation term's variance inflation from values as high as 2.61 to 1.03–1.06, confirming that the residualized specifications were well conditioned without making the metrics interchangeable. This distinction matters because the specialist–non-forest sign mirror was pronounced under DCLUMP and NLSI and weaker under PATCHD, but disappeared under EDGED, the only metric for which the non-forest interaction was predominantly negative. These

diagnostics document the predictor estimands but do not identify the cause of the fitted guild contrasts; four metrics in one system cannot do so.

Table 5.S2.3. Retained-metric geometry, in two panels. Panel A: pairwise correlations among the four retained metrics, raw and residualized, with 162 design cells per pair. Residualization changed four of the six pairwise correlations substantially, while DCLUMP–NLSI and EDGED–PATCHD remained comparatively stable. Panel B: variance inflation for forest cover and the fragmentation term. DISSAG was excluded at build and its redundancy with NLSI is sourced to the original rasters, not reconstructed here.

Panel A — pairwise correlations among the four retained metrics, by scale. The DISSAG–NLSI redundancy that motivated DISSAG's exclusion is stated in the text and sourced to the original rasters; it is not tabulated here, DISSAG having been dropped at dataset build.

Metric pair	Median r, raw	Median r, residualized	Change
DCLUMP – NLSI	0.758	0.706	−0.052
EDGED – PATCHD	0.695	0.618	−0.077
NLSI – EDGED	−0.532	−0.030	+0.502
NLSI – PATCHD	−0.258	+0.207	+0.465
DCLUMP – EDGED	−0.294	+0.133	+0.427
DCLUMP – PATCHD	−0.048	+0.312	+0.360

Panel B — variance inflation, by focal predictor and design cell.

Metric	VIF FCOV, raw	VIF FRAG, raw	VIF FCOV, resid.	VIF FRAG, resid.	VIF FRAG, +EVICV, raw
DCLUMP	2.05	1.18	1.91	1.03	1.20
NLSI	2.38	1.51	1.92	1.05	1.55
PATCHD	1.99	1.46	1.93	1.04	1.51
EDGED	3.57	2.61	1.94	1.06	2.89

5.S2.5 GAM residualization and validation

For each forest definition × response-grain × predictor-extent combination, each fragmentation metric is regressed on forest cover using a generalized additive model with a thin-plate regression spline smooth (Wood 2017); the residuals define the amount-adjusted (residualized) metric, denoted resDCLUMP, resNLSI, resEDGED, and resPATCHD. Residualization removes the nonlinear dependence represented by that fitted amount–configuration model and defines an operational estimand — the component of fragmentation not predicted by forest amount under that model — rather than a causal intervention on landscape configuration; whether residual dependence is in fact near zero is an empirical diagnostic, not a property guaranteed by the method. Validation of the residualization step — raw metric-versus-cover curves, GAM fit summaries, residual distributions, residual–cover dependence, and residualization adequacy — is reported in Table 5.S2.4, is visualized in Figure 5.S2.3, and is cross-referenced from the main-text residualization-validation results. The comparison of raw versus residualized fitted effects is a robustness result and is reported in Appendix 5.S5.

Table 5.S2.4. Residualization validation and amount-metric geometry, by metric, over 162 design cells each (nine forest definitions × eighteen grain × extent combinations). Variance removed is the deviance explained by the residualizing smooth; at the metric-summary level, nonlinear excess is the difference between the median deviance explained by the smooth and the corresponding median linear r^2, computed from the underlying unrounded values, and is not the median of cellwise differences. Median peak location is computed over interior-peak cells only, so its denominator differs by metric and is shown. Residual sd low/high is the ratio of residual standard deviation in the lowest to the highest cover quartile. VERIFICATION: refitted smooths practically reproduced the stored residualized predictors after standardization in all 648 cells (minimum correlation 0.999793, median 1.000000, median standardized discrepancy 0.000011). The stored columns were subsequently rescaled — median ratio of stored to refitted standard deviation 1.371, matching the 1.385 expected if the residual is re-standardized to unit variance at a median deviance explained of 0.479 — so relative and standardized residual quantities are interpreted and absolute residual variance is not. The residual reproduction is not uniform across metrics: the fifteen cells with the largest standardized discrepancy are fourteen PATCHD and

one NLSI, concentrated at the 70 % canopy threshold (eight of the fourteen, with five at 80 and one at 90 %) but spread across all three raster streams (seven Hansen 1,000 m, four Hansen 500 m, three Tuanmu 1,000 m), and all retain correlations above 0.9998 with maximum standardized discrepancies below 0.085. The concentration was therefore metric-associated rather than product-associated and did not compromise recovery of the fitted geometry.

Panel A — residualization: variance removed and its nonlinear component.

Metric	Variance removed	Linear r^2	Nonlinear excess	Median edf
DCLUMP	0.192	0.117	0.075	8.05
NLSI	0.377	0.298	0.080	7.85
PATCHD	0.576	0.252	0.323	8.27
EDGED	0.771	0.598	0.174	8.42

Panel B — amount-metric curve geometry and residual distribution along the cover gradient. Peak location is computed over interior-peak cells only, so its denominator differs by metric and is shown.

Metric	Interior peak (n/162)	Median peak, FCOV	Residual sd, low/high cover	Cells with greater low-cover spread
DCLUMP	104/162 (64.2%)	−0.68	1.92	98.8%
NLSI	65/162 (40.1%)	−0.68	2.82	100.0%
PATCHD	162/162 (100%)	+0.67	0.50	1.9%
EDGED	159/162 (98.1%)	+1.60	0.41	9.3%

5.S2.6 Environmental controls and controlled-arm geometry

The targeted environmental-control analysis (Appendix 5.S4.6), which refits each primary-supported interaction, adds actual evapotranspiration (AET), a moisture–energy index of realized water–energy availability, and topographic ruggedness (TOPOD), a terrain-heterogeneity index, to the primary spatially adjusted specification — the models that include effort, forest cover, fragmentation, the hierarchical random effects and the spatial autocovariate but no named environmental covariates, and which the analysis registry records as control-free; vegetation texture (EVICV, the EVI coefficient of variation) is reserved for a separately labelled secondary sensitivity that adds EVICV to AET and TOPOD. Because coefficient change after adding controls can reflect either confounding removal or variance inflation from shared predictor structure, the appendix reports the correlations and variance inflation factors of AET, TOPOD, and EVICV with forest cover and with each fragmentation metric across scales (Table 5.S2.5); these diagnostics condition the interpretation of the controlled-arm coefficient changes reported in Appendix 5.S5.

Table 5.S2.5. Control geometry. Pairwise associations of each environmental control with forest cover and with the fragmentation metrics, and variance inflation within the actual controlled specification (forest cover + fragmentation + AET + TOPOD, and the forest-specialist secondary set adding EVICV). Control variance inflation is stable across metrics within the raw formulation — AET 1.28 (1.19–1.36), TOPOD 1.63 (1.17–2.67), EVICV 1.60 (1.34–2.26) — and what changes with metric and formulation is the fragmentation term itself, reported in Table 5.S2.3 Panel B. In the raw formulation, forest cover exceeds a variance inflation of five in 101 of the 648 raw forest-specialist secondary cells, 77 of them under edge density and the remainder concentrated at broad predictor extents and strict canopy thresholds, reaching 11.9 at the 100 km extent. The two specifications differ materially. Without EVICV, 53 cells exceed a forest-cover inflation of five, all under edge density, with a maximum of 10.8. With EVICV added in the specialist secondary set, 101 of the 648 raw cells exceed five — 77 under edge density, 18 under NLSI and 6 under PATCHD — with a maximum of 11.9, so vegetation texture raises forest-cover inflation enough to carry NLSI and PATCHD cells across the threshold. No residualized cell exceeds five under either specification, the maxima being 2.1 and 3.4. The largest baseline inflation nonetheless arises from the coupling between forest cover and edge density documented in Table 5.S2.4 rather than the control itself. No cell in the residualized formulation exceeds five, and the fragmentation term's inflation there falls to a median of 1.06.

Control	Median |r| with FCOV	Median |r| with metric	Median VIF in raw controlled set	Median VIF in residualized controlled set
AET	0.429	0.274	1.26–1.28	1.26–1.30
TOPOD	0.592	0.317	1.62–1.64	1.62–1.64
EVICV (specialist secondary)	0.566	0.243	1.60 (1.34–2.26)	1.56 (1.34–2.18)

5.S2.7 References

Hansen MC, Potapov PV, Moore R, et al. (2013) High-resolution global maps of 21st-century forest cover change. Science 342:850–853.

Neel MC, McGarigal K, Cushman SA (2004) Behavior of class-level landscape metrics across gradients of class aggregation and area. Landscape Ecology 19:435–455.

Tuanmu M-N, Jetz W (2014) A global 1-km consensus land-cover product for biodiversity and ecosystem modelling. Global Ecology and Biogeography 23:1031–1045.

Wang X, Blanchet FG, Koper N (2014) Measuring habitat fragmentation: an evaluation of landscape pattern metrics. Methods in Ecology and Evolution 5:634–646.

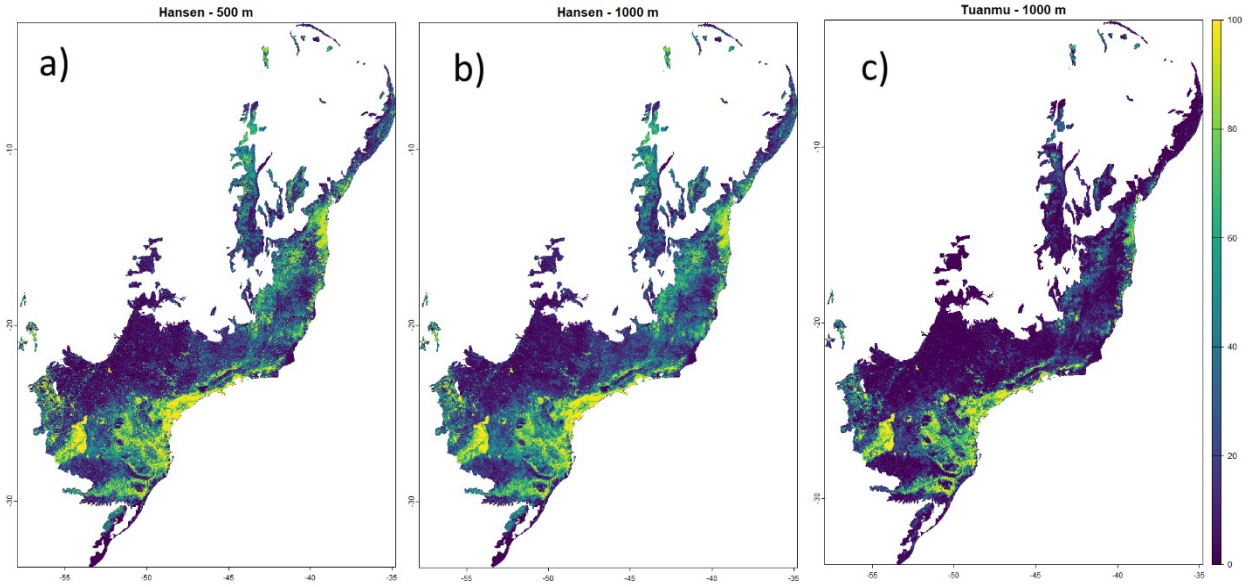
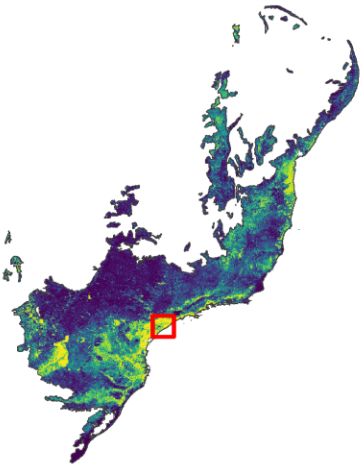
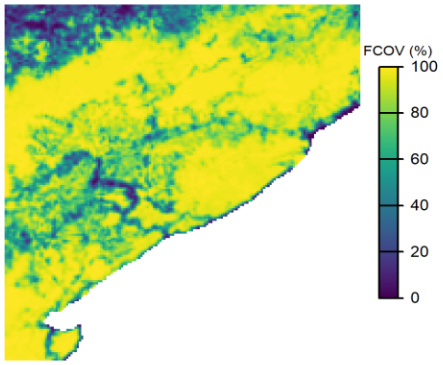


Figure 5.S2.1. Forest-cover streams across the Atlantic Forest. Percentage forest cover from (a) the Hansen Global Forest Change product resampled to 500 m, (b) the same product resampled to 1,000 m, and (c) the Tuanmu and Jetz global consensus land-cover product with forest classes combined, at 1,000 m. All panels share a common 0–100% scale and show continuous cover before binarisation, so the three canopy thresholds that generate the nine forest definitions are applied downstream of what is shown here. The three streams support two distinct contrasts held fixed within design cells: a resolution contrast between panels (a) and (b), which differ only in resampled grain, and a product contrast between panels (b) and (c), which differ additionally in classification framework, native resolution and temporal anchor, so the product contrast changes more than one property at once and is interpreted accordingly. The four configuration metrics are not mapped, because they are computed within predictor extents and have no native surface. Full provenance is given in Table 5.S2.1.

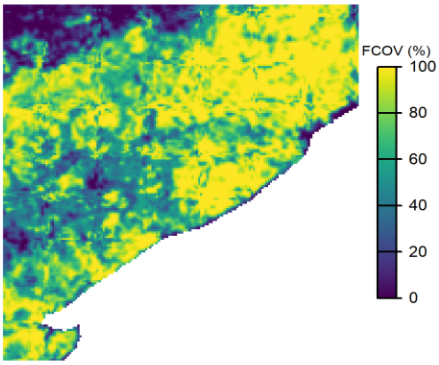
(a) Window location within the Atlantic Forest



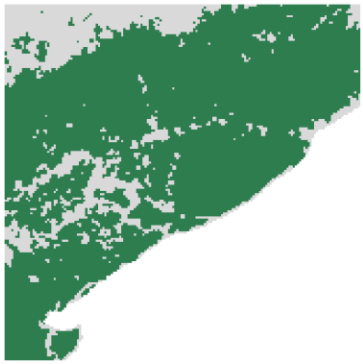
(b) Hansen, continuous



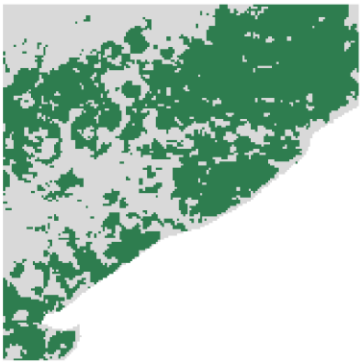
(c) Tuanmu, continuous



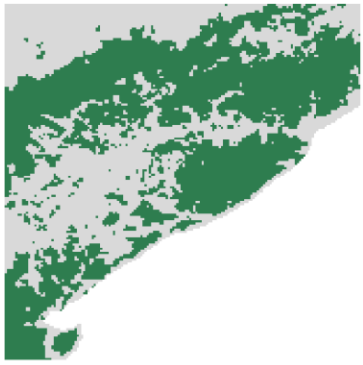
(d) Hansen, 70% threshold



(e) Tuanmu, 70% threshold



(f) Hansen, 90% threshold



(g) Tuanmu, 90% threshold

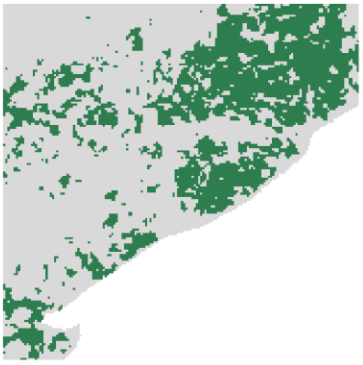


Figure 5.S2.2. Forest definitions at a representative landscape. (a) Location of the window within the Atlantic Forest, marked in red on the Hansen 1,000-m forest-cover surface. The remaining panels show that window under the Hansen Global Forest Change product (b, d, f) and the Tuanmu and Jetz consensus product (c, e, g), both at 1,000 m: continuous percentage forest cover on a common 0–100% scale (b, c), and the same surfaces binarised at the 70% (d, e) and 90% (f, g) canopy thresholds, with forest in green and non-forest in grey. Two operations separate the continuous panels from the rest. Binarisation discards the gradient in cover and commits the analysis to a categorical forest boundary, which is what makes the configuration metrics computable at all. Raising the threshold from 70% to 90% then removes moderately degraded and recovering vegetation, subdividing patches that were contiguous under the inclusive definition and raising edge and patch counts without any change in the underlying landscape. The two products additionally disagree about where boundaries fall, because they differ in classification framework, native resolution and temporal anchor as well as in what counts as canopy: the Tuanmu product fragments this landscape more readily at both thresholds. Both effects act directly on the configuration metrics, which is why forest definition is resolved throughout the analysis rather than fixed, and why canopy threshold is expected to alter the magnitude and support frequency of configuration effects without reversing their predicted direction. The intermediate 80% threshold and the 500-m Hansen stream are omitted here for clarity; all nine definitions enter the analysis. Window: 48.2°W to 46.8°W, 25.2°S to 23.8°S, approximately 140 by 155 km, spanning the Serra do Mar in coastal São Paulo and Paraná. .

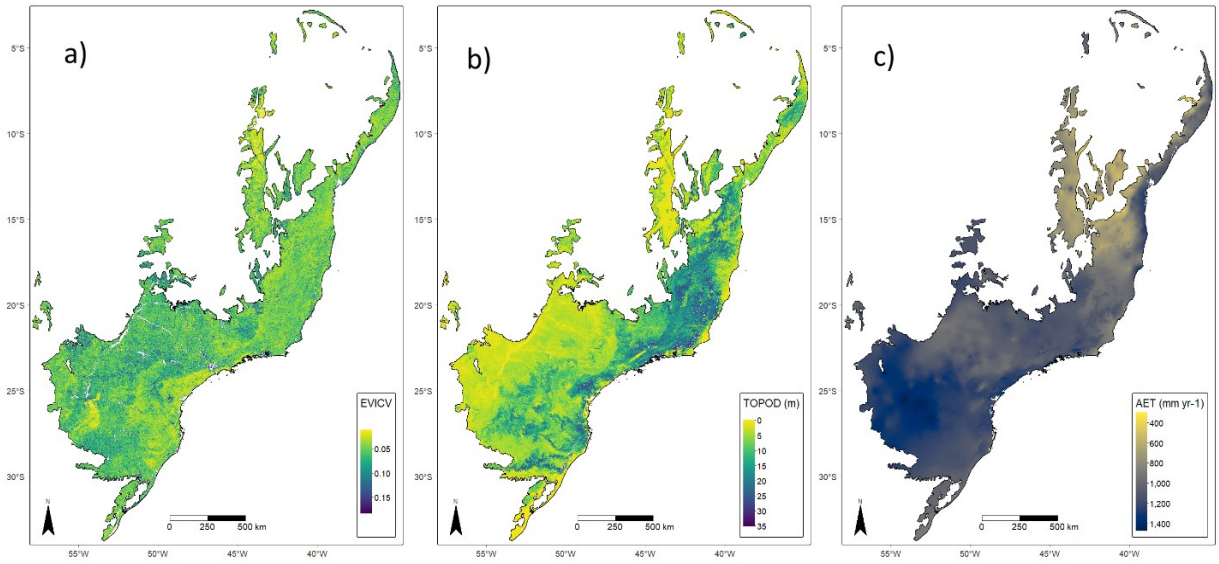


Figure 5.S2.3. Environmental controls across the Atlantic Forest. . (a) Actual evapotranspiration (AET), the realised water–energy flux, derived before standardisation as the product of annual potential evapotranspiration and the Priestley–Taylor coefficient. (b) Topographic ruggedness (TOPOD), the terrain ruggedness index, indexing local relief. (c) Coefficient of variation of the enhanced vegetation index (EVICV), a first-order greenness texture metric from MODIS imagery at 1-km resolution for 2000–2005. All are shown as native surfaces before aggregation to response units. None belongs to the primary factorial. AET and TOPOD enter jointly as a targeted robustness arm applied to primary-supported interactions, refitted while the comparison-specific spatial autocovariate is retained, so that the comparison isolates the addition of the controls. EVICV enters more narrowly still, as a secondary sensitivity for forest-specialist richness alone, added to the joint AET and TOPOD specification, because land-use-responsive vegetation texture is a plausible alternative to configuration for that stratum in particular. Associations of each control with forest cover and with the fragmentation metrics are reported in Table 5.S2.5. Full provenance is given in Table 5.S2.1

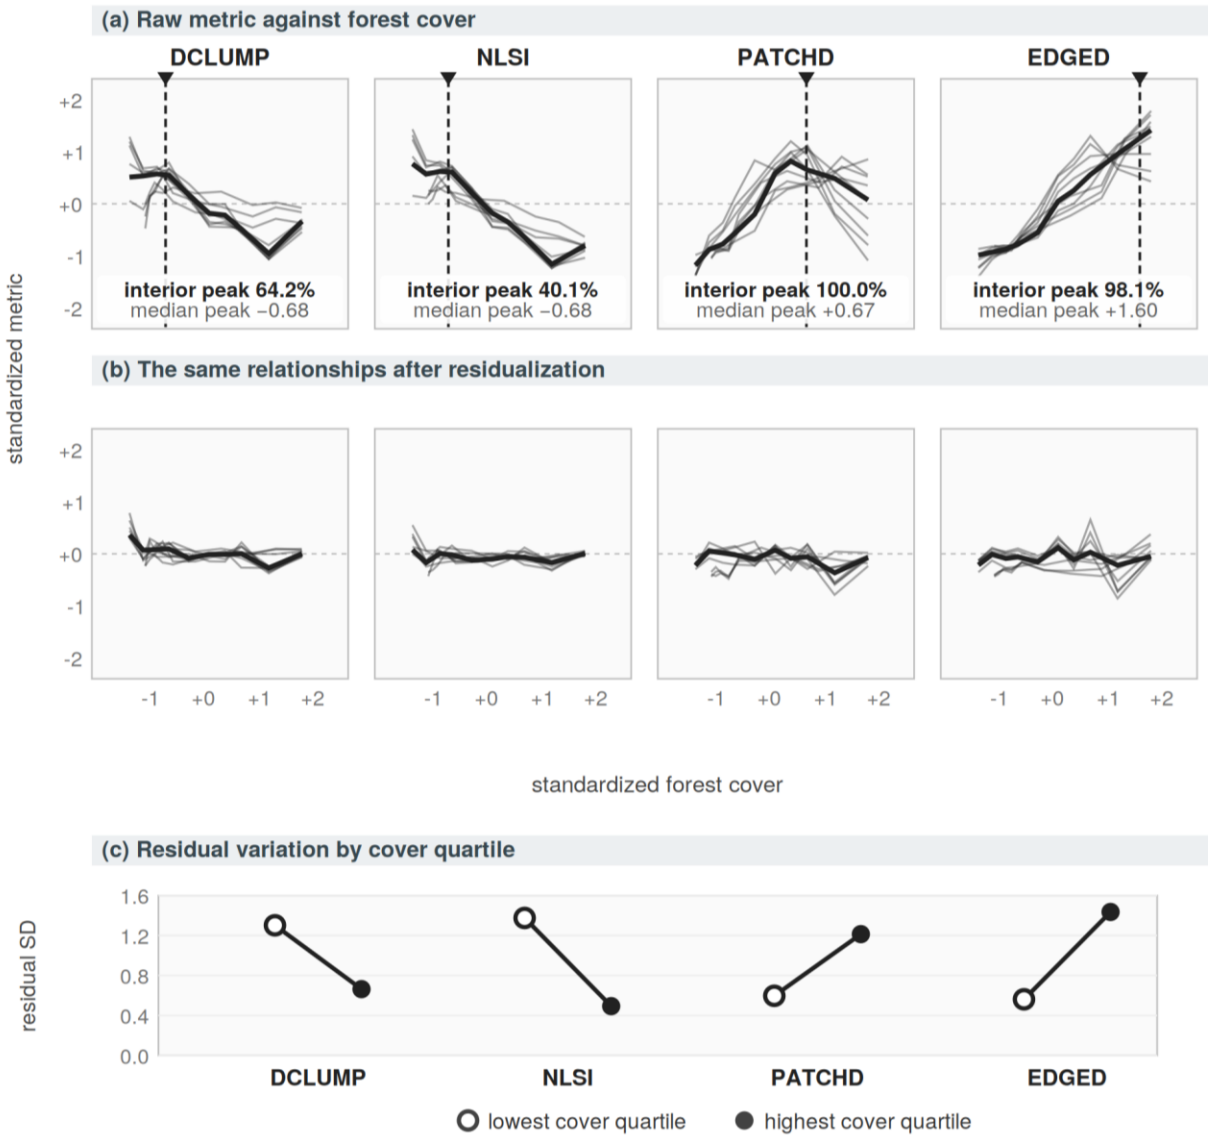


Figure 5.S2.4. Amount–configuration geometry and residualization validation. (a) Fitted relationships between forest cover and each raw fragmentation metric. Curves show the nine forest definitions at the fixed 25 km response grain and 55 km predictor extent; the interior-peak prevalence and median peak-location annotations summarize all 162 forest-definition × scale cells per metric; the median peak location is computed over interior-peak cells only, so its denominator differs by metric. Unimodality was nearly universal for EDGED and PATCHD and less frequent for DCLUMP and NLSI. (b) The same nine curves after residualization, on a common vertical scale, so that the flattening of the amount–metric association is visible directly rather than only as a summary statistic. (c) For each metric, paired points joined by a segment give the standard deviation of the residualized metric within the lowest and the highest quartile

of forest cover, so that the direction of the difference is read from the slope of the segment; DCLUMP and NLSI retained more residual variation at low cover, whereas EDGED and PATCHD retained more at high cover. Quantities are those tabulated in Tables 5.S2.3 and 5.S2.4. Variance inflation before and after residualization is reported in Table 5.S2.5 and is not repeated here.

APPENDIX 5.S3. Mechanism classification, eligibility, and evidential rules

The main-text Table 5.1 states the competing frameworks and their predicted interaction signatures. This appendix documents the operational rules used to map estimated effects to mechanism families, to classify eligibility and support, to compare across the robustness space, and to read guild decomposition — the evidential reference for the chapter rather than a second prediction table. Predictor definitions and geometry are documented in Appendix 5.S2; model specification and adequacy in Appendix 5.S4; empirical support in Appendix 5.S5. The design universe was code-anchored before full interpretation. The formal reporting, support and denominator rules were consolidated after initial directional summaries had been viewed and before R1 construction and manuscript drafting. They are fixed for current reporting and no rule may change because a result is inconvenient, but they are not presented as preregistered, prospectively frozen, or result-blind.

5.S3.1 From interaction sign to mechanism family

The chapter classifies the mechanism family from the sign of the forest-cover × fragmentation interaction and the geometry of the resulting marginal fragmentation effect across the forest-cover gradient (main-text Figure 5.2). The configuration-independent habitat-amount null is treated as a limiting theoretical null rather than an affirmatively testable outcome. Under the β_{frag} / $\beta_{frag \times amount}$ distinction maintained throughout this chapter, configuration irrelevance would require both the average fragmentation effect and the interaction to be near zero, with the marginal-effect curve inside a near-zero band across the evaluation domain; a near-zero interaction accompanied by a clearly non-zero average effect is gradient invariance, not configuration irrelevance. Absent a frozen equivalence test (bounds, scale, domain, and decision rule), a non-supported, near-zero, or imprecise estimate is classified as inconclusive, not as affirmative evidence for the null. A positive interaction is consistent with the scarcity and coexistence family, under which a negative fragmentation effect is strongest at low cover and weakens toward high cover; whether that effect crosses zero at sufficiently high cover would in principle discriminate coexistence (crossing) from scarcity (no reversal), but no operational crossing rule was implemented in this chapter: separating them requires a zero crossing supported within the observed cover domain, and point-estimate crossings without interval support are not sufficient. Positive-interaction cells are therefore reported as a joint positive-

interaction family and are not subdivided. A negative interaction is consistent with extinction filtering, under which the negative fragmentation effect instead strengthens as forest cover increases (Table 5.S3.1). This mapping is a classification of consistency, not a demonstration of demographic process: with richness-based, cross-sectional data the sign and geometry identify the most consistent mechanism family, and guild specificity and robustness calibrate credibility rather than providing independent causal proof.

Table 5.S3.1. Operational classification: how an estimated interaction is assigned to a mechanism class. This is an operational decision table, not a second prediction table. Classes are labelled operationally — negative-interaction family, positive-interaction family — with their directional correspondence to extinction filtering and to scarcity and coexistence given in the classification column rather than in the class name, and never by demonstrated process, since richness-based cross-sectional data identify the most consistent family and do not establish demographic mechanism. The SC-discordant diagnostic subset and the unsupported class are separate: only model-supported but structure-coefficient-discordant cells belong to the 298.

Operational class	Interaction estimate	Marginal-effect geometry	Classification
Equivalence-compatible null	within a declared equivalence region — NOT AVAILABLE: no equivalence region was declared, so this class cannot be reached in this chapter	flat, near zero	compatible with the amount-only limiting null

Operational class	Interaction estimate	Marginal-effect geometry	Classification
Negative-interaction family	negative, primary-supported	negative effect strengthens toward high cover	filtering-direction-consistent. Table 5.1's full filtering signature also requires the absence of a sign reversal, which this cell-level rule does not test; 86.8 % of primary-supported negative-interaction specialist cells cross zero within the observed cover domain. Mechanism attribution is therefore made at chapter level from interaction direction, high-cover intensification, guild specificity and robustness jointly, not from any single cell's classification.
Positive-interaction family	positive, primary-supported	negative effect weakens toward high cover	scarcity/coexistence-direction-consistent, left joint. Separating the two would require a supported zero crossing of the marginal-effect curve within the observed cover domain; no crossing rule was operationalized in this chapter, and point-estimate crossings without interval support are not sufficient. Cells in this class are therefore not subdivided.

Operational class	Interaction estimate	Marginal-effect geometry	Classification
SC-discordant diagnostic subset	model-supported but structure-coefficient discordant	indeterminate	no mechanism classification. Exactly the 298 cells of this subset, and no others.
Unsupported or inconclusive	not model-supported; may also be imprecise or near zero	indeterminate	no mechanism classification. Not part of the 298-cell SC-discordant subset.

5.S3.2 Eligibility and support classification

For every interaction cell the chapter retains the following fields separately and does not collapse them into a single flag:

- direction_int — the sign of the interaction coefficient;
- interval_supported_int — whether the 95% confidence interval excludes zero;
- model_supported_int — whether the interaction earns a parsimony gain of ΔAICc greater than two over the matched additive model;
- sc_concordant_int — whether the interaction coefficient and its product-term structure coefficient share sign;
- primary_supported_int — the conjunction model_supported_int AND sc_concordant_int.

Primary support is recomputed from the parsimony gain and the concordance test rather than read from any precomputed model-support-only flag (Table 5.S3.2). The primary gate selects downstream work; it is not the sole truth, and the interval, model, structure-coefficient, sign, and magnitude fields remain separately visible. Structure-coefficient discordance is treated as a

suppression or geometry warning, not as proof of invalidity: every eligible cell remains in R1 regardless of concordance status, and model-supported but structure-coefficient-discordant cells form a diagnostic subset of 298 cells. Because the composition of that subset was known before any computational sampling rule was fixed, the prospective option has expired: any additional analysis of it must either include all 298 deterministically, use an explicitly post hoc subset with its selection rule reported, or be omitted. Additive structure-coefficient concordance is reported as a diagnostic only and is not an additive support gate, asymmetrically with the interaction gate by design. Raw–residualized concordance is a matched robustness axis, not a support criterion; primary support is evaluated within each formulation.

Table 5.S3.2. Support and eligibility fields and the primary-support rule.

Field	Definition
direction_int	sign of the interaction coefficient
interval_supported_int	95% interval excludes zero; see the availability note below
model_supported_int	ΔAICc > 2 vs the matched additive model
sc_concordant_int	interaction coefficient and product-term structure coefficient share sign
primary_supported_int	model_supported_int AND sc_concordant_int

One availability qualification applies to the first of those fields. The stage-06 outputs store the interaction coefficient and its structure coefficient but no standard error, so interval_supported_int is not computable for the interaction term from the model-fitting outputs. It is available for forest-specialist richness from the parametric bootstrap, which yields percentile intervals for every term in each resampled cell; for the remaining three strata the interaction is reported by direction, model support and structure-coefficient concordance without an interval. Interval support is available throughout for the additive family, whose coefficient table stores standard errors. The primary-support rule does not depend on the interval field, so this gap constrains what can be reported rather than what can be classified; it is stated wherever interval support is cited.

5.S3.3 Robustness comparisons and denominators

Support is reported over an explicit accounting sequence — designed, attempted, eligible, and supported cells — so that untested, non-computable, and unsupported cells remain distinguishable rather than absorbed into a single count. Every comparison across the robustness space is matched within the design cell (Table 5.S3.3), holding fixed the fragmentation metric and formulation, forest stream, canopy threshold, response grain, and predictor extent. The eighteen grain × extent combinations are nested, overlapping windows on the same assemblages; repeated recovery across them is reported as scale coherence with the scale profile shown, not as independent replication, and interval envelopes formed as unions of per-cell intervals are consistency statements rather than confidence intervals. Product, resolution, threshold, scale, and formulation comparisons are reported as robustness summaries with their paired denominators, never as independent votes.

Table 5.S3.3. Design dimensions that condition the mapping from observed interaction signatures to ecological interpretation. These are cross-cutting: each conditions the mapping from an observed signature to an ecological mechanism rather than serving as a rival explanation for one framework.

Dimension	Predicted or analytical role	Interpretive limitation
Forest-definition threshold	Stricter canopy definitions isolate structurally intact interior forest. Each ecological framework predicts stronger effects under stricter definitions — coexistence and scarcity stronger negative effects overall, extinction filtering stronger effects at high cover — while the limiting null predicts none. Thresholds are predicted to alter magnitude and detectability, not to reverse the direction each framework predicts	Reclassification also changes amount–configuration geometry, so magnitude and detectability shift together and cannot be attributed separately
Habitat-affiliation guild	Localizes which assemblage component carries a response and tests whether an aggregate signal is diluted	Guilds differ in counts, zero prevalence and composition; restricted total richness is their arithmetic sum, not an independent replicate

Dimension	Predicted or analytical role	Interpretive limitation
Fragmentation metric and formulation	Robustness axis for the conditional pattern	Metric choice alters the recovered guild contrast, so fragmentation is not a single invariant exposure; four metrics cannot establish a general rule about metric classes
Response grain and predictor extent	Sensitivity axis; effect magnitude may vary with the extent at which configuration is measured	The eighteen combinations are nested, overlapping windows on the same assemblages; repeated recovery is scale coherence, not independent replication
Cross-sectional richness	Tests consistency with a mechanism family from interaction sign and marginal-effect geometry	Cannot establish demographic history, and cannot separate species replacement from ordered loss; a beta-diversity decomposition would be required

5.S3.4 Guild decomposition and specificity

Because restricted total richness is the arithmetic sum of the three guilds, aggregate-versus-component comparisons are interpreted as linked decompositions of the same assemblages, not as independent ecological replications; they localize which component carries a response and test whether an aggregate signal is diluted. On matched cells the chapter reports paired specialist-versus-non-forest, specialist-versus-facultative, and facultative-versus-non-forest contrasts in magnitude, sign coherence, support frequency, and product or formulation concordance, each with uncertainty. Under the documented interpretation rule: a response concentrated among specialists supports concentration of fragmentation costs among forest-dependent taxa, whereas comparable and comparably supported responses in specialist and non_forest weaken a dependency-graded reading and elevate explanations shared across guilds; such a pattern does not by itself demonstrate artefact or exclude extinction filtering. No magnitude-ratio threshold is

imposed, and a higher frequency of supported cells is not interpreted as a larger effect. Any empirically observed precision difference involving non_forest (Appendix 5.S1.3) is carried into the corresponding matched contrasts and their interpretation.

APPENDIX 5.S4. Model specification, model building, and adequacy

Residual adequacy is assessed for forest-specialist richness (Section S4.5, Table 5.S4.3).

5.S4.1 Model formulas and the four response estimands

Each richness response is fitted with a log-link Poisson generalized linear mixed model in glmmTMB (Brooks et al. 2017) in which an observation-level random intercept carries overdispersion (Harrison 2014; O'Hara and Kotze 2010), so that overdispersion is handled inside the likelihood rather than by changing the count family. Writing FCOV for forest cover and FRAG for a fragmentation metric, the analysis comprises three model families. The additive raw family models richness as a function of forest cover and fragmentation (control-free primary), with random intercepts for TAXA and METHOD and the observation-level intercept. The additive residualized family replaces the raw metric with its residualized counterpart. The interaction family adds the forest-cover-by-fragmentation product term. Each interaction family is specified in raw and residualized formulations. Environmental controls (AET, TOPOD; and EVICV in sensitivity) enter only in the separate control-comparison arms (Section S4.6), never in the primary models.

The same factorial is fitted separately to each of the four response estimands — total_restr, specialist, facultative, and non_forest — which are parallel aggregate and guild-specific estimands, not independent replicates, since restricted total richness is the arithmetic sum of the three guilds. The factorial spans the nine forest definitions, the four fragmentation metrics, and the eighteen grain × extent combinations, in raw and residualized formulations.

Writing resp for a response estimand, FCOV for forest cover, FRAG for a fragmentation metric, resFRAG for its residualized counterpart, EFFsb for the standardized sampling-effort covariate, and space for the spatial autocovariate, the specifications are:

- additive raw: resp ~ EFFsb + FCOV + FRAG + space + RS_TAXA + (1 | METHOD) + (1 | obs)
- additive residualized: resp ~ EFFsb + FCOV + resFRAG + space + RS_TAXA + (1 | METHOD) + (1 | obs)

- interaction raw: resp ~ EFFsb + FCOV + FRAG + FCOV:FRAG + space + RS_TAXA + (1 | METHOD) + (1 | obs)
- interaction residualized: resp ~ EFFsb + FCOV + resFRAG + FCOV:resFRAG + space + RS_TAXA + (1 | METHOD) + (1 | obs)
- targeted environmental-control analysis: the matched interaction model with + AET + TOPOD added and nothing else changed
- EVICV secondary sensitivity: the control arm with + EVICV added to AET and TOPOD

Sampling effort entered every model as EFFsb, a standardized fixed-effect covariate, rather than as an offset. All continuous predictors — forest cover, fragmentation, the environmental controls and sampling effort — were standardized before fitting; richness remained on its original count scale. An observation-level random intercept carried overdispersion, and TAXA and METHOD random intercepts carried the sampling structure throughout. Polynomial forms in forest cover replace the linear term where selected (Section S4.3).

5.S4.2 Baseline specification audits

No alternative effort specification and no year-retained model were fitted. Richness was aggregated across years to retain adequate data density within response cells.

5.S4.3 Polynomial-form selection

Polynomial forms in forest cover were considered for every cell under a minimum-AICc rule with the simpler form preferred among candidates within two AICc units of the minimum. No cell selected a polynomial form. All 5,184 cells across the four response strata and both formulations retained the linear structure in both forest cover and fragmentation (poly_fcov and poly_frag false in 100 % of cells; selected structure FCOV+frag throughout; Table 5.S4.1). The fitted richness responses therefore did not warrant polynomial terms in either forest cover or fragmentation under the stated parsimony rule. This is a statement about the richness models and carries no implication for the amount–configuration geometry among the predictors themselves, which is nonlinear in general and is handled separately by the GAM residualization of Appendix 5.S2.5 and the stage-02 diagnostics. Because every selected richness model was linear in forest cover and fragmentation, the marginal fragmentation effect was linear in forest cover in every

cell; no cell was excluded from derivative estimation and no mixed-functional-form caveat is needed.

Table 5.S4.1. Polynomial-form selection by response and formulation. No cell selected a polynomial form in either forest cover or fragmentation. All rows were selected under the same rule — minimum AICc, with the simpler form preferred among candidates within two AICc units — so that rule is stated here rather than repeated in every row; note that a linear selection is therefore consistent with a polynomial candidate lying inside the two-unit margin, and the candidate-level differences are unavailable.

Response	Formulation	Cells	Selected form	% of cells
specialist	raw	648	linear (FCOV+frag)	100%
specialist	residualized	648	linear (FCOV+frag)	100%
facultative	raw	648	linear (FCOV+frag)	100%
facultative	residualized	648	linear (FCOV+frag)	100%
non_forest	raw	648	linear (FCOV+frag)	100%
non_forest	residualized	648	linear (FCOV+frag)	100%
total_restr	raw	648	linear (FCOV+frag)	100%
total_restr	residualized	648	linear (FCOV+frag)	100%

(Displayed as a summary; the complete machine-readable selection record carries the design keys so individual cells are locatable.)

5.S4.4 Random-slope selection and provenance

The random-slope structure is selected per model among four candidates — no random slope, a forest-cover slope, a fragmentation slope, and correlated forest-cover and fragmentation slopes — with random intercepts for TAXA and METHOD throughout and the interaction-term random slope fixed out on estimability grounds. The four candidate structures and their selected frequencies are reported in Table 5.S4.2; the complete candidate-level AICc record is unavailable. The selected structure, its AICc, the no-slope AICc and their difference were retained, whereas values for the other two candidates are unavailable and are not reported. The linkage between stages 04 and 06 was audited using the complete design key. All 5,184 cells resolve uniquely to one selected random-slope structure and one selected polynomial structure, the polynomial structures agree in every cell, and the random-slope structures agree in 5,180, the four exceptions being cells where the selected correlated-slope structure failed to converge and the fitter recorded the no-random-slope structure instead. A forest-cover random slope was selected in 45.1 to 87.3 % of cells and was the majority choice in seven of the eight stratum × formulation combinations; correlated forest-cover and fragmentation slopes were the plurality choice only for specialist richness in the raw formulation (54.6 %). The selected structures therefore support among-taxon heterogeneity in forest-cover responses. A fragmentation-only slope was rarely selected (0 to 15.4 %), and was never selected under the residualized formulation for specialist or facultative richness; the no-slope structure appeared only in restricted-total and non-forest richness (7.4 to 26.2 %), never in specialist or facultative richness. The selected-structure source is the model-selection output, never a crossed-structure winners table.

Table 5.S4.2. Random-slope candidate selection across the four-candidate set, by response stratum and formulation. Counts are cells out of 648 per stratum × formulation.

Response	Formulation	none	FCOV slope	FRAG slope	FCOV +FRAG	Modal choice
specialist	raw	0	292 (45.1%)	2 (0.3%)	354 (54.6%)	FCOV+FRAG

Response	Formulation	none	FCOV slope	FRAG slope	FCOV +FRAG	Modal choice
specialist	residualized	0	501 (77.3%)	0	147 (22.7%)	FCOV
facultative	raw	0	566 (87.3%)	22 (3.4%)	60 (9.3%)	FCOV
facultative	residualized	0	500 (77.2%)	0	148 (22.8%)	FCOV
non_forest	raw	48 (7.4%)	408 (63%)	81 (12.5%)	111 (17.1%)	FCOV
non_forest	residualized	49 (7.6%)	471 (72.7%)	11 (1.7%)	117 (18.1%)	FCOV
total_restr	raw	122 (18.8%)	382 (59%)	100 (15.4%)	44 (6.8%)	FCOV
total_restr	residualized	170 (26.2%)	352 (54.3%)	26 (4%)	100 (15.4%)	FCOV

(Selected structures only. The model-structure stage retained the winning structure and not the losing-candidate values, so the complete four-candidate AICc record is unavailable and is not reported.)

5.S4.5 Spatial adjustment and residual adequacy

Spatial dependence is represented by a derived autocovariate. For each comparison, Pearson residuals from the corresponding full non-spatial model are regressed on longitude and latitude with a two-dimensional thin-plate smooth, gam(residuals ~ s(long, lat)) (Wood 2017), under a Gaussian family with the default basis dimension. The fitted surface is z-standardized and entered as the fixed-effect covariate space in both the interaction model and its matched additive comparison, so the pair differs only in the interaction term and their AICc difference is

attributable to that term alone. Two additive fits therefore exist in this chapter and serve different purposes. The additive coefficients reported in the main text, Table 5.2, Figure 5.3 and Table 5.S5.3 come from the additive family fitted at stage 04, whereas the parsimony comparison that produces every primary-support classification uses the additive member refitted alongside each interaction model at stage 06. The two runs differ in every design cell, by a median of 0.002 and a ninetieth-percentile 0.010 standardized units in the fragmentation coefficient, and they converge on slightly different sets: four cells lose the stage-06 comparison and four lose the stage-04 fit, overlapping in only two. Reported additive coefficients and the model-support gate are therefore not the same fitted models and should not be reconciled cell by cell.

Table 5.S4.3. Simulation-based residual diagnostics, forest-specialist richness. Scaled quantile residuals were computed with DHARMa (Hartig 2022) from the model objects retained by stage 06, with 250 simulations per model and no refitting. The audit set is the design reference forest definition (Hansen 1,000 m, 80 % threshold) crossed with all four fragmentation metrics, both formulations and all eighteen grain × extent combinations — 144 cells, deterministic and defined by design position rather than sampled. Note that glmmTMB does not implement conditional simulation, so all residuals are simulated unconditionally on the random effects. Results are descriptive: 144 non-independent views of one dataset, not independent tests, and no multiplicity correction is applied because none would be meaningful. These diagnostics assess distributional adequacy and play no part in model selection.

Test	Cells	Result	Reading
Dispersion	144	median statistic 0.716; no cell rejects	No over- or underdispersion detected. The statistic sits below 1 throughout — 0.483–0.885 raw, 0.466–0.901 residualized, median 0.716 — but never significantly.

Test	Cells	Result	Reading
Uniformity (KS)	144	rejects in every cell	Universal but small: residuals occupy the central two quartiles in 45.6% of cases against 50 expected, and the outer 10% of the scale in 12.3% against 10. At ~1,020 observations per cell a KS test detects departures of this size reliably, so the p-value indexes sample size as much as severity. The direction is mild tail-heaviness, not a central pile.
Outliers	144	95.8% of cells show no excess	No systematic outlier problem; the six exceptions are scattered across metrics and scales.
Zero inflation	144	no cell rejects	No excess zeros, despite 22–25% of forest-specialist observations being zero. This was the plausible distributional failure for this stratum and it does not occur.

Residual diagnostics indicated mild distributional departures but no systematic dispersion, outlier or zero-inflation problem. Dispersion and zero inflation were not detected in any audited cell, while outlier excess was absent in 95.8 % of them, and zero inflation in particular was the plausible failure for a stratum in which between a fifth and a quarter of observations are zero. Uniformity rejects everywhere, but the magnitude of departure is small and its direction is a slight excess of extreme residuals rather than the central concentration that unmodelled grouping structure would produce. The departure does not concentrate in any metric, and neither formulation sits systematically closer to nominal dispersion than the other. The audit covers forest-specialist richness at one forest definition, so it does not establish adequacy for the other three strata or the other eight definitions. It uses unconditional simulation because glmmTMB implements no conditional alternative, which inflates simulated variance relative to the data where random-effect variance is substantial and is a plausible contributor to the residual underdispersion. A uniformity departure of this size cannot be resolved into a specific

misspecification without a targeted alternative-model comparison, which this chapter does not undertake.

5.S4.6 Targeted environmental-control robustness

Environmental controls were evaluated for interaction surfaces meeting the primary-support criterion. This analysis was not an additional primary-support screen across the full eligible factorial; it asked whether the interactions selected for ecological interpretation remained directionally and quantitatively stable after explicit environmental adjustment. For each primary-supported interaction the matched model was refitted with AET and TOPOD added jointly, retaining the response data, raw or residualized fragmentation formulation, selected polynomial structure, random-effects architecture and remaining model components. Additional leave-one-control-out specifications included AET alone and TOPOD alone, allowing dependence on control composition to be assessed. The primary spatially adjusted model — which includes effort, forest cover, fragmentation, the hierarchical random effects, an observation-level random effect and the comparison-specific spatial autocovariate shared by the matched models, but no explicit environmental covariates — remained the inferential reference throughout. EVICV entered only as a separately labelled secondary sensitivity for forest-specialist richness, where it was added to AET and TOPOD; this restricted scope reflects that stratum's direct alignment with the forest-dependency hypothesis and does not imply EVICV coverage for the other response strata. The controls analysis reports convergence, coefficient direction, proportional coefficient change and sign agreement relative to the matched primary model (Table 5.S4.4). Because it was applied to primary-supported interactions, it evaluates the stability of detected surfaces rather than changes in support classification among cells not selected for interpretation.

Table 5.S4.4. Targeted environmental-control accounting. Every primary-supported interaction was refitted with AET and TOPOD; EVICV was added for forest-specialist richness only.

Response	Cells refitted	Sign reversals	Median coefficient change
specialist	411	0	−1.24%
specialist, +EVICV	411	0	≈ +2%

Response	Cells refitted	Sign reversals	Median coefficient change
facultative	274	0	+1.33%
non_forest	201	0	+0.75%
total_restr	362	0	−4.11%

5.S4.7 References

Brooks ME, Kristensen K, van Benthem KJ, et al. (2017) glmmTMB balances speed and flexibility among packages for zero-inflated generalized linear mixed models. R Journal 9:378-400.

Harrison XA (2014) Using observation-level random effects to model overdispersion in count data in ecology and evolution. PeerJ 2:e616.

Hartig F (2022) DHARMa: Residual diagnostics for hierarchical (multi-level/mixed) regression models. R package.

O'Hara RB, Kotze DJ (2010) Do not log-transform count data. Methods in Ecology and Evolution 1:118-122.

Wood SN (2017) Generalized Additive Models: An Introduction with R, 2nd edn. CRC Press.

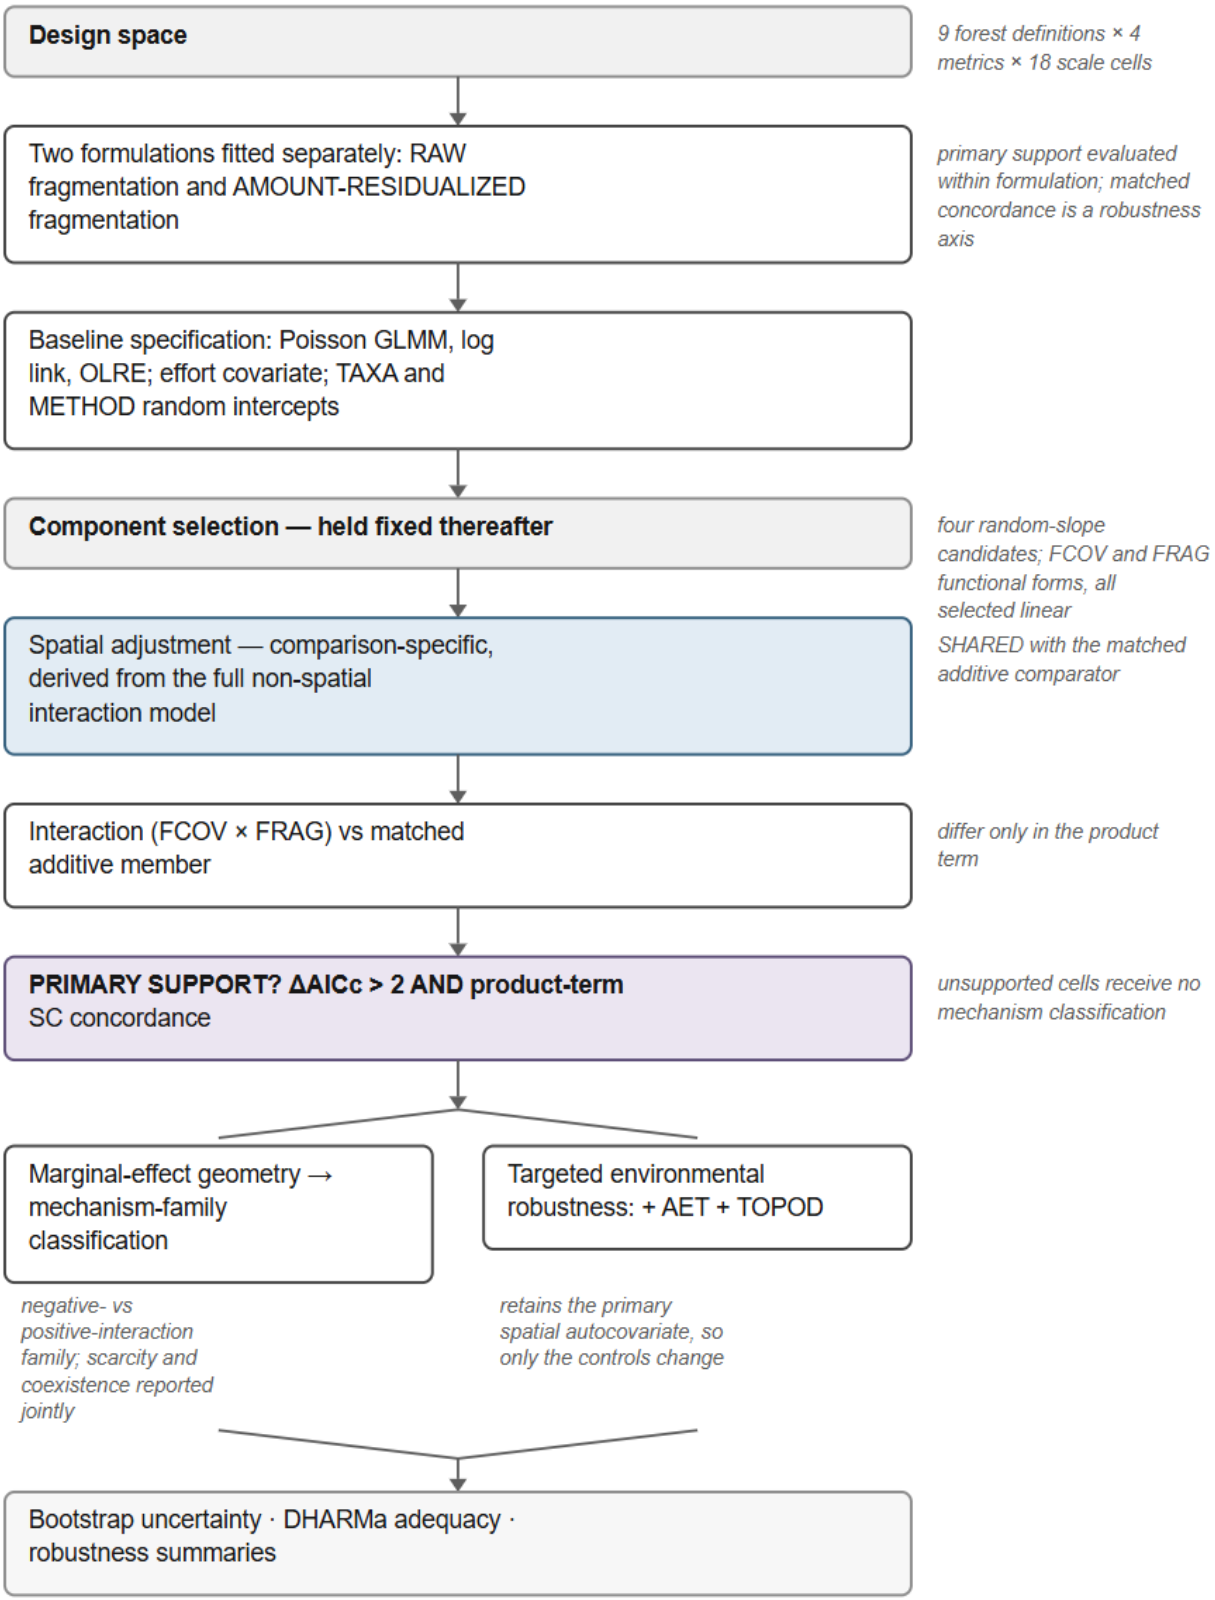

Design space
9 forest definitions × 4 metrics × 18 scale cells
Two formulations fitted separately: RAW fragmentation and AMOUNT-RESIDUALIZED fragmentation
primary support evaluated within formulation; matched concordance is a robustness axis
Baseline specification: Poisson GLMM, log link, OLRE; effort covariate; TAXA and METHOD random intercepts
Component selection — held fixed thereafter
four random-slope candidates; FCOV and FRAG functional forms, all selected linear
Spatial adjustment — comparison-specific, derived from the full non-spatial interaction model
SHARED with the matched additive comparator
Interaction (FCOV × FRAG) vs matched additive member
differ only in the product term
PRIMARY SUPPORT? ΔAICc > 2 AND product-term SC concordance
unsupported cells receive no mechanism classification
Marginal-effect geometry → mechanism-family classification
Targeted environmental robustness: + AET + TOPOD
negative- vs positive-interaction family; scarcity and coexistence reported jointly
retains the primary spatial autocovariate, so only the controls change
Bootstrap uncertainty · DHARMa adequacy · robustness summaries

Figure 5.S4.1. Model-building and inference sequence for the amount-conditioned fragmentation analysis. Raw and amount-residualized fragmentation formulations are fitted separately and primary support is evaluated within formulation; their matched concordance is treated as a robustness property. For each primary comparison, the residual spatial autocovariate is derived from the full non-spatial interaction model and shared with its matched additive comparator. Primary-supported interactions are then evaluated through two complementary paths: marginal-effect geometry and mechanism-family classification, and targeted environmental robustness after adding AET and TOPOD while retaining the primary spatial autocovariate. Unsupported cells receive no mechanism classification.

APPENDIX 5.S5. Fragmentation × amount results and robustness validation

5.S5.1 Eligibility and primary-support accounting

The full design comprised 1,296 cells per response stratum — nine forest definitions × four fragmentation metrics × eighteen grain × extent combinations × two formulations — and 5,184 cells across the four strata. Every designed cell was attempted and all 5,184 interaction models converged; their matched additive comparison converged in 5,180 cells; the remaining four lacked a converged matched additive comparison. Those four are ineligible for one reason: the interaction model converged while the matched additive comparison did not, leaving no baseline for the parsimony test. They are classified as missing paired member rather than as failed fits, and all four fall in the residualized formulation of the PATCHD metric at broad predictor extents (three facultative, one restricted total). Eligibility is therefore 648 of 648 raw cells in every stratum, and 648, 647 or 645 residualized cells depending on stratum. Primary support was recomputed as model support and structure-coefficient concordance jointly, evaluated within each formulation and never pooled across them; the recomputation reproduced the stored classification exactly in all eight stratum × formulation cells (Table 5.S5.1, Table 5.S5.2), and the parsimony threshold was ΔAICc greater than two throughout. Across the design, 298 model-supported cells failed structure-coefficient concordance and entered the diagnostic subset described in Appendix 5.S3.2, concentrated in the residualized formulation of the specialist (144) and non-forest (44) strata.

Table 5.S5.1. Design, attempt, eligibility, and failure accounting, by stratum and formulation.

Stratum	Formulation	Designed	Attempted	Interaction-model convergence failures	Non-computable	Excluded (reason)	Eligible
specialist	raw	648	648	0	0	—	648
specialist	residualized	648	648	0	0	—	648
facultative	raw	648	648	0	0	—	648

Stratum	Formulation	Designed	Attempted	Interaction-model convergence failures	Non-computable	Excluded (reason)	Eligible
facultative	residualized	648	648	0	0	missing paired member (3)	645
non_forest	raw	648	648	0	0	—	648
non_forest	residualized	648	648	0	0	—	648
total_restr	raw	648	648	0	0	—	648
total_restr	residualized	648	648	0	0	missing paired member (1)	647

Table 5.S5.2. Primary interaction support by stratum × metric, evaluated within each formulation (not pooled). Designed maximum = 18 grain × extent × 9 forest definitions = 162 per stratum × metric × formulation; entries are supported/eligible counts (eligible excludes non-computable cells).

Stratum	Formulation	DCLUMP (supp/elig)	NLSI	PATCHD	EDGED
specialist	raw	70/162	14/162	61/162	69/162
specialist	residualized	90/162	48/162	41/162	18/162
facultative	raw	16/162	27/162	52/162	89/162
facultative	residualized	54/162	21/162	12/159	3/162
non_forest	raw	6/162	43/162	40/162	33/162
non_forest	residualized	5/162	10/162	27/162	37/162

Stratum	Formulation	DCLUMP (supp/elig)	NLSI	PATCHD	EDGED
total_restr	raw	36/162	14/162	48/162	90/162
total_restr	residualized	78/162	48/162	25/161	23/162

5.S5.2 Additive effects across raw and residualized formulations

Forest-amount effects exceeded fragmentation effects in specialist, facultative and restricted-total richness under both formulations: median forest-cover coefficients are reported by fragmentation metric, because neither the fragmentation coefficient nor its magnitude relative to forest amount is invariant to metric choice. Across the 24 specialist, facultative and restricted-total metric × formulation combinations, forest amount exceeded fragmentation in every case, with median per-cell absolute amount-to-fragmentation ratios from 2.42 to 8.50. Metric-specific ranges were 3.33 to 7.80 for DCLUMP, 2.42 to 8.50 for NLSI, 3.73 to 7.77 for PATCHD and 3.23 to 5.13 for EDGED. Non-forest ratios ranged from 1.18 to 2.18, indicating that non-forest richness responded to configuration about as strongly as to amount. Non-forest richness inverted both terms, but its amount interval excluded zero in only 23.5 % of raw and 9.3 % of residualized cells, against 86.7 to 100 % elsewhere, so the inversion is directional rather than precisely estimated. Residualization raised the proportion of negative fragmentation coefficients in specialist, facultative and restricted-total richness while lowering the proportion whose interval excluded zero — facultative from 68.5 to 84.1 % negative against 34.3 to 20.7 % interval-supported, restricted total from 71.3 to 82.6 against 40.1 to 26.9 — residualization increased directional consistency in these strata, while interval support rose for specialist richness from 40.4 to 53.2 % and fell for facultative and restricted-total richness; the cause of that divergence is not established here. Additive structure-coefficient concordance is reported as a diagnostic and is not a support gate. Metric-marginal coefficients and ratios are not used to characterize the magnitude relationship; the complete metric-resolved values are in Table 5.S5.3, and the per-cell additive fragmentation coefficient across the whole design is shown in Figure 5.S5.3.

5.S5.3 Interaction support and marginal-effect geometries

Under Table 5.S3.1, cells divide into three classes: primary-supported with a negative interaction, assigned to the negative-interaction family and directionally consistent with

extinction filtering; primary-supported with a positive interaction, assigned to the positive-interaction family and directionally consistent with scarcity and coexistence; and cells that are not primary-supported, which receive no mechanism classification; only the 298 model-supported but structure-coefficient-discordant cells form the diagnostic subset of Appendix 5.S3.2. In the raw formulation, forest-specialist richness assigns 212 cells to the negative-interaction family against 2 to the positive, a ratio above one hundred to one; 174 restricted-total cells fall in the negative-interaction family against 14 in the positive, and 154 facultative cells against 30. Non-forest richness inverts, assigning 37 cells to the negative-interaction family against 85 to the positive. The residualized formulation gives the same ordering with smaller supported counts — specialist 182 against 15, restricted-total 166 against 8, facultative 82 against 8, non-forest 38 against 41. Between 67.0 and 87.8 % of eligible cells in every stratum are not primary-supported and receive no mechanism assignment, and support frequency is not read as evidence of effect size in itself. One further limitation is inherent to the operational table: primary-supported positive cells cannot be split between scarcity and coexistence cell by cell, because that distinction requires a zero crossing of the marginal-effect curve within the observed cover domain, which is evaluated over the marginal-effect geometry rather than from the interaction coefficient alone. Those cells are therefore classified as positive-interaction family without further subdivision. The fixed-cell marginal-effect curves are shown in Figure 5.S5.1, and the per-cell interaction coefficient across the whole design in Figure 5.S5.4.

5.S5.4 Product, resolution, threshold, scale, and formulation robustness

Robustness is reported over the nine forest definitions, with the product contrast (Hansen 1,000 m versus Tuanmu 1,000 m at matched resolution), the resolution contrast (Hansen 500 m versus Hansen 1,000 m within one product) and the threshold contrast held separate and never conflated; over the eighteen grain × extent combinations; and over the raw and residualized formulations; the additive fragmentation coefficient and the interaction are displayed for every fitted forest-specialist cell in Figures 5.S5.3 and 5.S5.4. Every comparison is matched within design cell, holding fixed the metric, forest stream, canopy threshold, response grain and predictor extent not being contrasted. Direction showed moderate-to-high concordance throughout and magnitude was less stable: matched sign concordance ran 68 to 88 % for the product contrast, 73 to 83 % for the resolution contrast, 61 to 84 % between the 70 and 90 % thresholds. Comparing median coefficients, Hansen at 1,000 m yielded specialist interactions

about 1.8 times as large in absolute magnitude as Tuanmu at matched resolution in the raw formulation and about 3.2 times as large in the residualized formulation, and the coarser Hansen resolution yielded systematically more negative coefficients than the finer one. Scale support is reported as counts across the matrix with the scale profile shown (Table 5.S5.4); the eighteen combinations are nested, overlapping windows on the same assemblages, so repeated recovery across them is scale coherence rather than independent replication, and interval envelopes formed as unions of per-cell intervals are consistency statements rather than confidence intervals.

Because the eighteen scale combinations, nine forest definitions, four metrics and two formulations are design factors rather than replicates, each capable of shifting direction and magnitude, we also profile the robustness space by factor. Coefficient summaries use the complete 5,184-cell interaction universe; support counts use the 5,180 eligible cells, the difference being four cells whose matched additive comparison failed to converge. The design is balanced in the coefficient universe — every stratum × stream × threshold × metric × formulation × scale combination appears exactly once — so marginal medians are computed over identical sets of every other factor. No variance partition, interval or test is reported: these are one dataset viewed 1,296 ways per stratum. Sensitivity is summarised by the median absolute deviation of level medians from the stratum median.

All thirty one-factor specialist marginal medians were negative, ranging from −0.109 to −0.005. At the joint stream × threshold × metric × formulation level, 63 of 72 specialist profiles were negative, and 83.3 % of the 1,296 individual specialist cells were negative. The few one-factor departures for restricted-total and facultative richness were confined to the three smallest scale combinations — 10 km grain with 10 or 25 km extent, and 25 km grain with 25 km extent — where only 0 of 36 or 2 of 36 eligible cells were supported, indicating weak or unresolved effects rather than robust reversal.

The nine positive specialist joint profiles clustered in three architectures: raw NLSI under Hansen 500 m, and residualized EDGED or PATCHD under Tuanmu 1,000 m, each at all three canopy thresholds. Threshold therefore did not explain the departures. Their coefficients were small, from +0.001 to +0.024, against a design-wide specialist median of −0.057 and strongest profiles reaching −0.203, indicating localized attenuation or weak reversal rather than a broad

directional change. DCLUMP produced no positive specialist joint profile. This clustering is descriptive and was not tested as a formal factor interaction.

The amount-dependent geometry of the four metrics, the variance each residualization removed, and the distribution of residual variation along the cover gradient are reported in Appendix 5.S2.4 and S2.5; those diagnostics condition the interpretation of the metric-dependent guild contrasts recorded here. Two observations are carried there rather than developed here: the departing metric differs between formulations, and between-metric displacement for forest specialists increases rather than decreases under residualization. Neither is explained by the present analysis, and with four metrics in one system neither can be.

Fragmentation metric produced the largest displacement of factor-level medians in all eight stratum × formulation profiles, with scale second in six. These summaries are descriptive and rest on one balanced factorial; the design factors differ in type and in number of levels, and level counts are reported alongside rather than used to normalize. Residualization increased metric-associated displacement for specialists from 0.025 to 0.056, reduced it for facultative richness from 0.022 to 0.014, and left restricted-total and non-forest displacement nearly unchanged. Matched raw–residualized median differences remained small across strata, from 0.002 to 0.008, although individual-cell signs differed more often than those central summaries suggest.

Design factors also move the difference between guilds, and which factor does so is not the same as within a single stratum. The specialist-versus-non-forest contrast varies most across metric levels (range of level medians 0.101) and scale levels (0.095); the specialist-versus-facultative and specialist-versus-restricted-total contrasts vary most across scale (0.078 and 0.075). These ranges are computed over paired differences and are not directly comparable with the within-stratum dispersion measures above, which summarise a different quantity. Metric choice most strongly altered the specialist–non-forest contrast. On matched cells the paired difference is −0.124 under NLSI, −0.120 under DCLUMP and −0.059 under PATCHD, but only −0.023 under EDGED, with sign agreement between the two strata rising from 23.5 % under NLSI to 62.3 % under EDGED. Edge density is also the only metric under which the non-forest median interaction is negative (−0.034). The specialist–non-forest sign mirror therefore disappeared under that metric. The metric-level paired difference was smaller than under the other metrics because the raw and residualized formulation-specific contrasts partly offset; it should not be

read as uniform attenuation within both formulations. Matched guild contrasts are reported in Table 5.S5.5, the ordered guild gradient in Figure 5.S5.5, and the corresponding predictor geometry in Appendix 5.S2.4.

Within the scale matrix the gradient is carried by predictor extent rather than response grain. The Spearman correlation between extent and the interaction coefficient is negative in twenty-two of twenty-four stratum × formulation × grain combinations, strongest for forest specialists (−0.46 to −0.22) and weakening systematically as grain coarsens. The trend is directional but not monotone: coefficients strengthen with extent to roughly 70 to 85 km and turn back at 100 km. Grain contributes little once extent is held common: restricted to the extents shared by all three grains (40 to 100 km), specialist medians are −0.063, −0.074 and −0.060 at the 10, 25 and 40 km grains. The strongest single combination is a 25 km grain with a 70 km extent for specialist richness (median −0.103, 18 of 36 cells supported) against −0.005 and 5 of 36 at the 10 km grain with a 10 km extent. The five combinations in which predictor extent equals response grain are a distinct measurement situation, because the landscape window and the assemblage unit share a footprint; they are retained as valid design cells under the extent-greater-than-or-equal-to-grain rule, and are identified here so that readers can locate them rather than treating them as ordinary landscape windows. The full design-factor profile is shown in Figure 5.S5.2.

Canopy threshold is tested on two quantities, since Table 5.1 predicts stronger effects under stricter thresholds — a claim about magnitude — while the underlying rationale, that stricter thresholds isolate structurally intact interior forest, is a claim about detectability. On magnitude, raising the threshold from 70 to 90 % in the raw formulation made median interactions progressively more negative for forest-specialist richness (−0.053, −0.059, −0.063), facultative richness (−0.006, −0.019, −0.023) and restricted-total richness (−0.009, −0.019, −0.026), and attenuated the predominantly positive non-forest interaction toward zero (+0.029, +0.032, +0.004). The same direction held for facultative and restricted-total richness after residualization, but the residualized specialist coefficient became less negative at the strictest threshold (−0.073, −0.072, −0.048). Absolute changes across the sweep are formulation-specific and must not be pooled: in the raw formulation non-forest richness moves most (0.025) and forest specialists least (0.010), with restricted-total and facultative richness intermediate (0.017 each), whereas in the residualized formulation the specialist change is the largest of the eight profiles (0.025) and runs in the opposite direction. Relative changes were larger for restricted-total and facultative

richness, but those percentages exceed 100 % because both coefficients are close to zero at the inclusive threshold, so absolute changes are the primary comparison.

On detectability the prediction is more clearly supported than on magnitude. Primary support — an operational detectability measure combining the parsimony gain with structure-coefficient concordance, not a general index of statistical power — rises with canopy threshold in five of the six specialist, facultative and restricted-total formulation profiles, monotonically in four of them: 25 to 34 % for raw restricted-total richness, 22 to 33 % residualized, 26 to 31 % for raw facultative richness, 9 to 20 % residualized, and 28 to 35 % for raw specialist richness, which peaks at 36 % under the intermediate threshold rather than rising monotonically. Restricted-total richness is an aggregate containing all three guilds and is not itself a forest-associated guild; that it tracks the specialist and facultative profiles rather than the non-forest one indicates that the forest-associated components carry the threshold signal within the aggregate. Non-forest richness moves the other way, falling from 21 to 13 % raw and from 21 to 5 % residualized, which is coherent with the same delineation logic: a stricter definition of forest sharpens the configuration signal for species that depend on forest and degrades it for species that do not. The single exception is forest-specialist richness under the residualized formulation, where support falls from 37 to 24 %. That exception is concordant across both quantities, appearing in the coefficient medians and the support frequencies alike. Those summaries derive from the same fitted cells, and the coefficient sign enters the support decision through the structure-coefficient test, so their agreement shows internal concordance rather than independent corroboration and is not evidence against sampling variation. It is the one result in the threshold sweep that the prediction does not accommodate.

Table 5.S5.3. Additive forest-amount and fragmentation coefficients by response stratum, fragmentation metric and formulation. Each row summarises the 162 design cells at that stratum, metric and formulation, obtained by marginalizing over nine forest definitions and eighteen grain × extent combinations. Coefficients are medians across those cells. The ratio is the median of the per-cell absolute ratios |βFCOV| / |βFRAG|, not the ratio of the two marginal medians, and the median is used because per-cell ratios have a heavy upper tail. The interval columns give the

percentage of cells whose 95 % interval excludes zero. No coefficient or ratio is pooled across fragmentation metrics, because neither quantity is invariant to metric choice.

Stratum	Metric	Formulation	Median βFCOV	Median βFRAG	Median |βFCOV| / |βFRAG|	FCOV CI excl. 0	FRAG CI excl. 0
Forest specialist	DCLUMP	raw	+0.310	−0.094	3.33	100.0%	48.1%
Forest specialist	DCLUMP	residualized	+0.356	−0.062	5.61	100.0%	53.1%
Forest specialist	NLSI	raw	+0.283	−0.124	2.42	98.8%	47.5%
Forest specialist	NLSI	residualized	+0.357	−0.052	6.90	100.0%	40.7%
Forest specialist	PATCHD	raw	+0.371	−0.020	7.77	100.0%	22.2%
Forest specialist	PATCHD	residualized	+0.345	−0.062	5.14	100.0%	57.4%
Forest specialist	EDGED	raw	+0.398	−0.065	5.13	100.0%	43.8%
Forest specialist	EDGED	residualized	+0.339	−0.075	4.73	100.0%	61.7%
Forest facultative	DCLUMP	raw	+0.118	−0.019	6.33	88.3%	25.9%
Forest facultative	DCLUMP	residualized	+0.122	−0.018	6.56	100.0%	16.0%

Stratum	Metric	Formulation	Median βFCOV	Median βFRAG	Median |βFCOV| / |βFRAG|	FCOV CI excl. 0	FRAG CI excl. 0
Forest facultative	NLSI	raw	+0.112	−0.025	4.16	71.6%	26.5%
Forest facultative	NLSI	residualized	+0.123	−0.013	6.92	100.0%	20.4%
Forest facultative	PATCHD	raw	+0.129	+0.012	3.73	95.1%	46.3%
Forest facultative	PATCHD	residualized	+0.121	−0.025	4.72	100.0%	21.6%
Forest facultative	EDGED	raw	+0.158	−0.040	3.23	92.0%	38.3%
Forest facultative	EDGED	residualized	+0.119	−0.031	4.07	100.0%	24.7%
Non-forest	DCLUMP	raw	−0.061	+0.025	1.82	13.6%	12.3%
Non-forest	DCLUMP	residualized	−0.065	+0.036	1.75	9.3%	20.4%
Non-forest	NLSI	raw	−0.059	+0.017	1.53	26.5%	17.9%
Non-forest	NLSI	residualized	−0.067	+0.029	2.18	9.3%	22.2%
Non-forest	PATCHD	raw	−0.109	+0.102	1.19	30.2%	43.2%
Non-forest	PATCHD	residualized	−0.069	+0.040	1.39	9.3%	14.8%
Non-forest	EDGED	raw	−0.082	+0.013	1.18	23.5%	55.6%
Non-forest	EDGED	residualized	−0.073	−0.006	1.59	9.3%	32.7%
Restricted total	DCLUMP	raw	+0.144	−0.023	6.11	100.0%	33.3%

Stratum	Metric	Formulation	Median βFCOV	Median βFRAG	Median |βFCOV| / |βFRAG|	FCOV CI excl. 0	FRAG CI excl. 0
Restricted total	DCLUMP	residualized	+0.155	−0.019	7.80	100.0%	18.5%
Restricted total	NLSI	raw	+0.141	−0.025	4.82	100.0%	40.1%
Restricted total	NLSI	residualized	+0.155	−0.017	8.50	100.0%	22.8%
Restricted total	PATCHD	raw	+0.155	+0.006	4.67	100.0%	46.3%
Restricted total	PATCHD	residualized	+0.152	−0.019	6.43	100.0%	26.5%
Restricted total	EDGED	raw	+0.183	−0.038	4.12	100.0%	40.7%
Restricted total	EDGED	residualized	+0.150	−0.033	4.67	100.0%	39.5%

Table 5.S5.4. Robustness across product, resolution, threshold, scale, and formulation, matched within cell.

Contrast	Matched dimension held fixed	Concordance (n/N)	Notes
product (Hansen1000 vs Tuanmu1000)	metric, formulation, threshold, grain, extent	68.1–88.4% of 214–216 matched pairs	product contrast at matched 1,000 m resolution

Contrast	Matched dimension held fixed	Concordance (n/N)	Notes
resolution (Hansen500 vs Hansen1000)	metric, formulation, threshold, grain, extent	72.7–83.3% of 213–216 matched pairs	resolution contrast within the Hansen product
threshold 70 vs 80	metric, formulation, stream, grain, extent	80.1–87.0% of 213–216 matched pairs	adjacent canopy thresholds
threshold 80 vs 90	metric, formulation, stream, grain, extent	71.3–92.6% of 215–216 matched pairs	adjacent canopy thresholds
threshold 70 vs 90	metric, formulation, stream, grain, extent	61.1–83.8% of 214–216 matched pairs	extreme canopy thresholds
formulation (raw vs residualized)	metric, forest definition, grain, extent	64.0–88.0% of 645–648 matched pairs	robustness axis, not a support criterion

5.S5.5 Aggregate dilution and guild specificity

Restricted total richness is compared with its three guild components as linked decompositions of the same assemblages rather than as independent ecological replications, since it is their arithmetic sum. On matched cells the specialist and non-forest interaction coefficients shared sign in only 40.6 % of raw and 35.2 % of residualized comparisons, with median paired differences of −0.107 and −0.087 — the two strata move in opposite directions along the cover gradient. Specialist coefficients were more negative than restricted-total coefficients on matched cells, consistent with dilution of a response concentrated among forest-dependent taxa when guilds are aggregated. The non-forest precision caveat of Appendix 5.S1.3 is carried through: that stratum's amount effect is interval-supported in under a quarter of cells, so the specificity

argument rests on the interaction-sign contrast across all cells and not on the amount-effect reversal. No magnitude-ratio threshold is applied. The mirror pattern is consistent with the guild specificity expected under the filtering interpretation, while remaining compatible with composition-dependent sampling alternatives, but it does not by itself exclude artefacts acting through the composition of what is detected — detectability substitution within a survey, or edge- or road-biased survey placement, would produce a similar contrast without ecology, and are addressed in the main-text limitations.

Seventeen primary-supported specialist cells returned positive interaction coefficients, two raw and fifteen residualized, and all seventeen occurred at the 10 km response grain. They fall into three architectures. Three arose where the predictor extent equalled the response grain — raw and residualized DCLUMP at Hansen 1,000 m and 70 % (+0.038, +0.040) and raw NLSI at Hansen 500 m and 70 % (+0.058). Eight arose under residualized edge density at Tuanmu 1,000 m, at 80 and 90 % thresholds and extents of 40 to 100 km (+0.044 to +0.065); in those cells the non-forest interaction was negative in seven of eight, so specialist and non-forest responses oppose one another and the scale account does not cover them. The remaining six arose under residualized PATCHD at predictor extents of 10 to 40 km: Hansen 1,000 m at 70 % with a 10 km extent (+0.103), Hansen 500 m at 70 % with 10, 25 and 40 km extents (+0.149, +0.078, +0.116), Hansen 500 m at 80 % with a 10 km extent (+0.120), and Tuanmu 1,000 m at 80 % with a 40 km extent (+0.049). In the same design cells the other residualized configuration metrics ran predominantly neutral or negative, with DCLUMP between −0.071 and +0.040, NLSI between −0.113 and +0.004, and edge density between −0.019 and +0.054. At Hansen 500 m, 70 %, 10 km grain and 25 km extent the bootstrap intervals were DCLUMP −0.091 to −0.012, NLSI −0.134 to −0.003 and PATCHD +0.021 to +0.144; at the 40 km extent they were −0.119 to −0.018, −0.154 to −0.027 and +0.041 to +0.193. None of the corresponding raw PATCHD interactions was primary-supported, the raw coefficients ranging from +0.045 to −0.011. All six retained sign under the AET and TOPOD control arm and under the EVICV sensitivity step.

Table 5.S5.5. Guild-specificity matched contrasts.

Contrast	Magnitude difference	Sign coherence	Support frequency	Uncertainty
specialist vs non_forest (raw)	median paired difference -0.1068	40.6% same sign	36 of 648 both primary-supported	IQR -0.1500 to -0.0423
specialist vs non_forest (residualized)	median paired difference -0.0865	35.2% same sign	25 of 648 both primary-supported	IQR -0.1651 to -0.0075
specialist vs facultative (raw)	median paired difference -0.0563	63.7% same sign	75 of 648 both primary-supported	IQR -0.0874 to -0.0204
specialist vs facultative (residualized)	median paired difference -0.0443	74.8% same sign	28 of 630 both primary-supported	IQR -0.1032 to -0.0035
facultative vs non_forest (raw)	median paired difference -0.039	52.2% same sign	31 of 648 both primary-supported	IQR -0.0963 to -0.0004
facultative vs non_forest (residualized)	median paired difference -0.0462	44% same sign	9 of 645 both primary-supported	IQR -0.0820 to 0.0090
specialist vs total_restr (raw)	median paired difference -0.0469	69.8% same sign	76 of 648 both primary-supported	IQR -0.0794 to -0.0154
specialist vs total_restr (residualized)	median paired difference -0.0389	79.8% same sign	70 of 639 both primary-supported	IQR -0.0938 to 0.0080

Contrast	Magnitude difference	Sign coherence	Support frequency	Uncertainty
facultative vs total_restr (raw)	median paired difference 0.0011	75.2% same sign	96 of 648 both primary-supported	IQR -0.0218 to 0.0267
facultative vs total_restr (residualized)	median paired difference 0.0089	77.7% same sign	39 of 636 both primary-supported	IQR -0.0112 to 0.0228
non_forest vs total_restr (raw)	median paired difference 0.0413	65.7% same sign	38 of 648 both primary-supported	IQR 0.0120 to 0.0703
non_forest vs total_restr (residualized)	median paired difference 0.0413	47.6% same sign	25 of 639 both primary-supported	IQR 0.0032 to 0.0762

5.S5.6 Environmental controls and EVICV sensitivity

The environmental-control robustness analysis comprised the 1,248 interactions meeting the primary support criterion: 362 restricted-total, 411 specialist, 274 facultative and 201 non-forest cells. Adding AET and TOPOD produced no interaction-sign reversals in any response stratum. Median coefficient changes relative to the matched primary spatially adjusted models were −4.11 % for restricted-total richness, −1.24 % for specialist richness, +1.33 % for facultative richness and +0.75 % for non-forest richness, with the specialist interquartile range running −12.8 to +8.0 % and 203 of 411 specialist cells shifting by more than 10 %, so direction is insensitive to control specification while per-cell magnitude is not. The AET-only and TOPOD-only specifications likewise retained interaction direction throughout, indicating that the result did not depend on one member of the joint control set; convergence held except in two cells across the four strata. The secondary EVICV sensitivity covers the 411 primary-supported forest-specialist cells only, producing no sign reversals and a median interaction-coefficient change of approximately +2 %. This supports directional stability of the specialist interaction under additional adjustment for vegetation texture; it does not establish that EVICV is geometrically independent of fragmentation. These controlled fits are interpreted as a robustness layer on the interaction surfaces selected under the primary criterion (Table 5.S5.6). They do not replace the primary support classification and do not estimate control-induced support transitions in cells that were not selected for interpretation.

Table 5.S5.6. Environmental-control sensitivity, targeted coverage: the 1,248 primary-supported interaction cells.

Stratum	Cells refitted with AET+TOPOD	Sign reversals	Median coefficient change	EVICV secondary layer
specialist	411	0	−1.24%	411/411 sign retained, median ≈ +2%
facultative	274	0	+1.33%	—
non-forest	201	0	+0.75%	—

Stratum	Cells refitted with AET+TOPOD	Sign reversals	Median coefficient change	EVICV secondary layer
restricted total	362	0	−4.11%	—

5.S5.7 Bootstrap validation and main-text provenance

The parametric bootstrap is an adequacy check on inference already obtained, not the inferential basis. Coverage is forest-specialist richness at three levels: control-free (411 cells), with actual evapotranspiration added (410) and with topographic ruggedness added (410), at 250 resamples each. No both-controls rung exists, because the stage-07 slimming step retained only the single-control model objects. The non-forest bootstrap was abandoned at 70 of 201 cells on compute grounds; the completed subset is first in job order rather than a random sample and is skewed across metrics, so no percentage is computed over it and it is not reported. Restricted total and facultative richness were not bootstrapped. At the control-free level, interval-based and likelihood-based support agreed in all 411 cells, every intended bootstrap cell completed, and the median relative root-mean-square error was 0.271. Coefficient recovery therefore showed no obvious instability under the resampling design. The bootstrap assesses recovery only; broader distributional adequacy was evaluated separately for forest-specialist richness using simulation-based residual diagnostics (Appendix 5.S4.5; Table 5.S4.3). Structure-coefficient intervals excluded zero in 63.1 % of raw and 27.4 % of residualized cells while the coefficient interval excluded zero in all of them, a pattern showing that inferential support for the partial interaction term was stronger than its simple association with the response; the structure coefficient is not a variance partition and no decomposition of explained variance is implied. Every numerical row in the main-text Results traces to its source cells in the frozen results master, with eligibility denominators and support fields explicit (Table 5.S5.7, Table 5.S5.8, Table 5.S5.9, Table 5.S5.10, Table 5.S5.11). The partial coefficient of determination produced by the bootstrap stage is not reported: it divides a variance of log-scale linear predictors by a variance of raw counts, is small by construction and can go negative.

Table 5.S5.7. Bootstrap validation and main-text provenance. All entries are descriptive medians taken across cells: the coefficient column is the median of per-cell observed coefficients, and the bounds are the medians of the per-cell bootstrap percentile limits. They are not a pooled estimate and not a confidence interval for any population quantity, since the cells are overlapping views of one dataset.

Main-text row	Source cell(s)	Bootstrap interval	Eligibility denominator	Support fields
specialist L1 (control-free), raw	214 primary-supported cells	median -0.1157; median per-cell bounds -0.1809, -0.0539	214 of 214 bootstrapped	100% interval excludes zero; SC 63.1%; rel. RMSE 0.2731
specialist L2a (+AET), raw	214 primary-supported cells	median -0.1102; median per-cell bounds -0.1738, -0.048	214 of 214 bootstrapped	93% interval excludes zero; SC 66.4%; rel. RMSE 0.2861
specialist L2b (+TOPOD), raw	214 primary-supported cells	median -0.1204; median per-cell bounds -0.1845, -0.0565	214 of 214 bootstrapped	95.8% interval excludes zero; SC 66.4%; rel. RMSE 0.2675
specialist L1 (control-free), residualized	197 primary-supported cells	median -0.1157; median per-cell bounds -0.1766, -0.0566	197 of 197 bootstrapped	100% interval excludes zero; SC 27.4%; rel. RMSE 0.2693
specialist L2a (+AET), residualized	196 primary-supported cells	median -0.1107; median per-cell bounds -0.174, -0.0482	196 of 197 bootstrapped	95.9% interval excludes zero; SC 27.6%; rel. RMSE 0.2902

Main-text row	Source cell(s)	Bootstrap interval	Eligibility denominator	Support fields
specialist L2b (+TOPOD), residualized	196 primary-supported cells	median -0.1286; median per-cell bounds -0.1915, -0.0694	196 of 197 bootstrapped	96.4% interval excludes zero; SC 29.1%; rel. RMSE 0.2479

Post hoc descriptive characterization of the bootstrap marginal-effect curves. The material below was assembled after the results were inspected and is descriptive. No operational crossing rule was implemented, and the endpoint classifications are not a mechanism classifier.

Table 5.S5.8. Direction of bootstrap interval support for the forest-specialist marginal effect at the lower cover percentiles, among primary-supported cells (214 raw, 197 residualized). Intervals are pointwise percentile intervals evaluated separately at each percentile.

Formulation	Percentile	Interval > 0	Interval < 0	Overlaps zero
Raw	5th	100	3	111
Raw	25th	78	14	122
Residualized	5th	40	8	149
Residualized	25th	16	11	170

Table 5.S5.9. Cells whose interval was positive at the 5th cover percentile and negative at the 95th, by stream, metric and formulation, over primary-supported cells. Denominators are support-selected, so rates are conditional on primary support and are not design-wide frequencies. Cells with single-digit denominators are shown as counts only.

Formulation	Stream	Metric	Opposite-signed / n
Raw	Hansen	EDGED	51 / 69
Raw	Hansen	PATCHD	19 / 52
Raw	Hansen	DCLUMP	5 / 43
Raw	Hansen	NLSI	1 / 2
Raw	Tuanmu	PATCHD	3 / 9
Raw	Tuanmu	DCLUMP	0 / 27
Raw	Tuanmu	NLSI	0 / 12
Residualized	Hansen	PATCHD	23 / 40
Residualized	Hansen	DCLUMP	6 / 61
Residualized	Hansen	NLSI	2 / 30
Residualized	Hansen	EDGED	0 / 10
Residualized	Tuanmu	NLSI	7 / 18
Residualized	Tuanmu	DCLUMP	0 / 29
Residualized	Tuanmu	EDGED	0 / 8
Residualized	Tuanmu	PATCHD	0 / 1

No endpoint pattern was observed among the primary-supported Tuanmu DCLUMP cells in either formulation. No equivalence test was defined, so these absences are not evidence of a zero effect. All raw primary-supported EDGED cells derive from the Hansen stream, so the apparent stream contrast partly reflects metric composition rather than an independent product effect.

Table 5.S5.10. Threshold, response-grain and predictor-extent profiles of the opposite-signed endpoint pattern, over primary-supported specialist cells. Primary-support selection changes the metric and stream composition of every denominator, so these are selected-cell descriptive profiles and not threshold, grain or extent effects.

Factor	Level	Raw	Residualized
Canopy threshold	70 %	34 / 61	21 / 80
Canopy threshold	80 %	29 / 77	12 / 65
Canopy threshold	90 %	16 / 76	5 / 52
Response grain	10 km	27 / 74	16 / 71
Response grain	25 km	29 / 81	13 / 78
Response grain	40 km	23 / 59	9 / 48
Predictor extent	10 km	1 / 5	0 / 5
Predictor extent	25 km	7 / 14	3 / 12
Predictor extent	40 km	11 / 26	3 / 28
Predictor extent	55 km	10 / 38	2 / 34
Predictor extent	70 km	13 / 46	10 / 41
Predictor extent	85 km	18 / 46	12 / 42
Predictor extent	100 km	19 / 39	8 / 35

Table 5.S5.11. Paired within-cell ratio of the residualized to the raw coefficient magnitude, across the 648 forest-specialist design cells matched between formulations. Raw and residualized predictors were standardized separately and the stored residualized columns were re-standardized to unit variance (Appendix 5.S2), so these are formulation-specific standardized contrasts and NOT fractions of an effect removed.

Metric	n	Fragmentation, median [IQR]	Interaction, median [IQR]	Interaction ratio < 1
DCLUMP	162	0.99 [0.79–1.26]	1.36 [1.16–1.74]	8.0%
NLSI	162	1.16 [0.90–1.55]	4.37 [2.38–11.96]	3.1%
PATCHD	162	1.26 [0.39–3.46]	0.65 [0.30–1.69]	64.8%
EDGED	162	0.83 [0.45–1.82]	0.20 [0.09–0.43]	88.9%

Residualization did not attenuate coefficients uniformly in forest-specialist cells. The metric-resolved paired ratios differed between the additive fragmentation and interaction terms. For the additive fragmentation coefficient, the median ratio was near or above one under DCLUMP, NLSI and PATCHD and below one under edge density. For the interaction coefficient, the median ratio exceeded one under DCLUMP and NLSI and fell below one under PATCHD and edge density. No pooled attenuation summary is reported because the ratio is not invariant to metric choice. Because the formulations are non-symmetric and were standardized separately, attenuation where observed is a formulation-specific contrast, not an estimate of the fraction of an effect removed and not evidence of ecological absence.

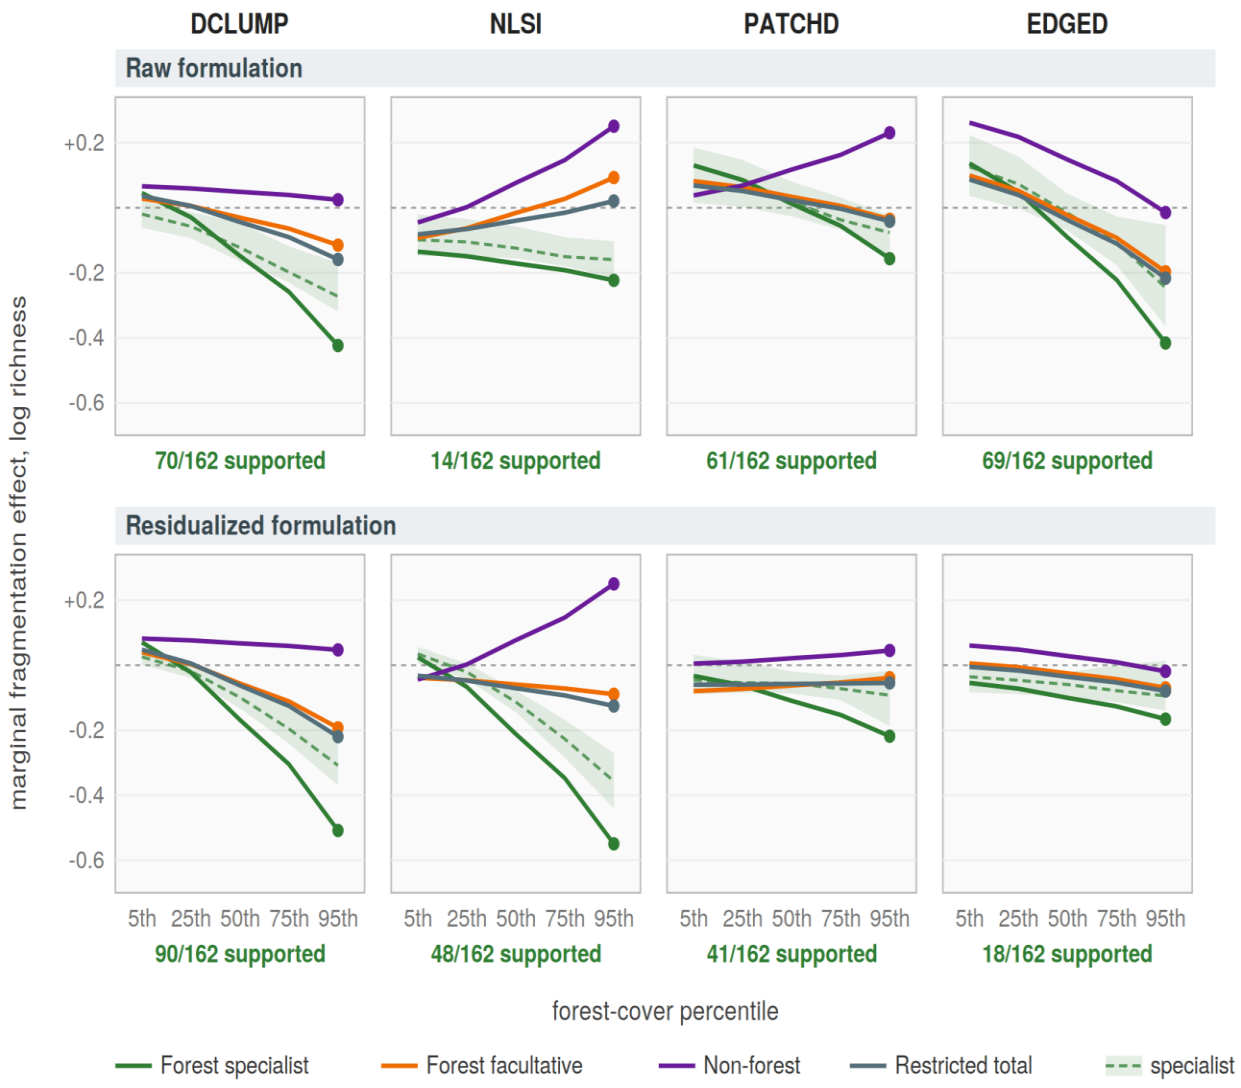


Figure 5.S5.1. Marginal fragmentation effect across the forest-cover gradient at a fixed display cell, by response stratum and metric. . Curves are shown for restricted-total, forest-specialist, forest-facultative and non-forest richness under each of the four fragmentation metrics, at the Hansen 1,000 m forest definition with an 80 % canopy threshold, a 25 km response grain and a 55 km predictor extent. The display cell was fixed on design grounds before the curves were inspected: it is the central valid grain–extent combination under the Hansen 1,000 m reference stream and the intermediate 80 % canopy threshold, and was not chosen for its estimated result. Percentile intervals are shown for forest-specialist richness only, the parametric bootstrap having been run on all 411 primary-supported forest-specialist cells across all four metrics, including 62 NLSI cells. Intervals are therefore absent from a specialist panel only where this particular display cell did not reach primary support under that metric and formulation, not because the

metric was left unbootstrapped. The other three strata are shown as point estimates throughout, since no bootstrap was run for them. A fifth, dashed series shows the median forest-specialist marginal effect with its interquartile range across all 162 design cells of each metric and formulation, so that the single display cell can be read against the design-wide specialist distribution; it is a summary across cells and not a fifth response stratum.

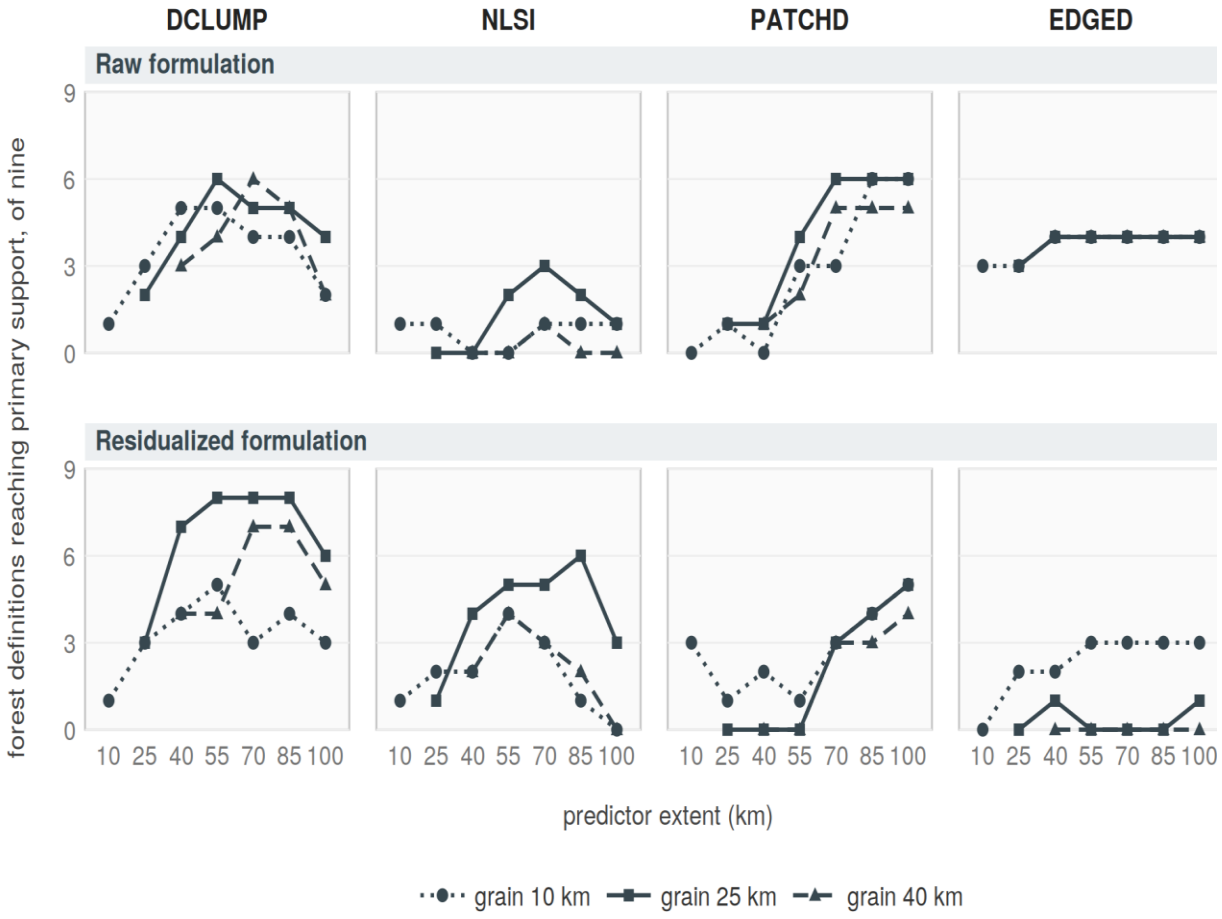


Figure 5.S5.2. Primary-support frequency across predictor extent, by metric and formulation. The figure is a support and accounting display and carries no coefficient trajectories, which belong to main-text Figure 5.4. Columns are fragmentation metrics and rows are formulations, never pooled. Within each panel the vertical axis gives the number of the nine forest definitions reaching primary support at that grain and extent, so the denominator is nine in every case and is stated here rather than repeated in the figure. Response grain is distinguished by both line type and symbol shape. Combinations in which predictor extent is smaller than response grain do not exist by design and are absent rather than plotted as zero. Primary support requires both a parsimony gain over the matched additive model and structure-coefficient concordance, so it is an operational gate rather than an interval criterion. Because forest definition, metric, scale and formulation are design factors rather than replicates, the panels describe how support is distributed across the design and do not constitute independent tests.

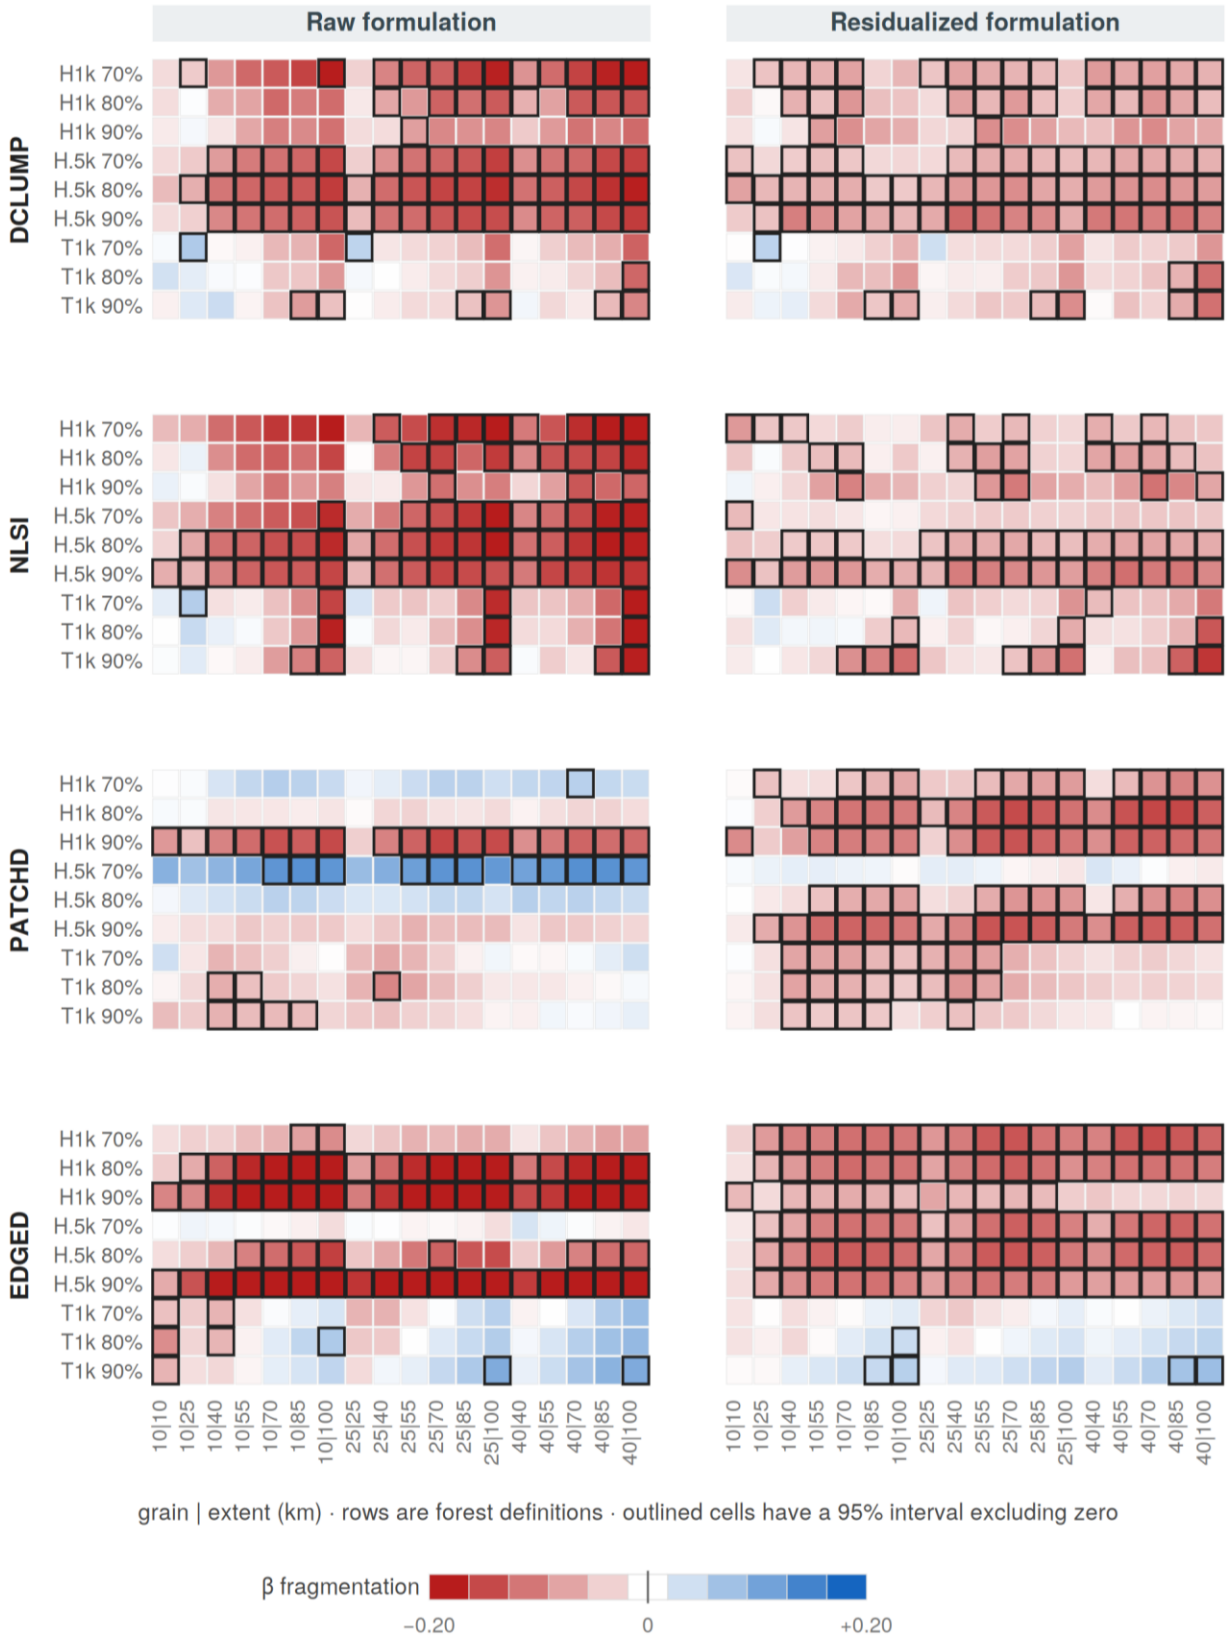

Raw formulation
Residualized formulation
DCLUMP
NLSI
PATCHD
EDGED
H1k 70%
H1k 80%
H1k 90%
H.5k 70%
H.5k 80%
H.5k 90%
T1k 70%
T1k 80%
T1k 90%
10|10
10|25
10|40
10|55
10|70
10|85
10|100
25|25
25|40
25|55
25|70
25|85
25|100
40|40
40|55
40|70
40|85
40|100
grain | extent (km) · rows are forest definitions · outlined cells have a 95% interval excluding zero
β fragmentation
−0.20
0
+0.20

Figure 5.S5.3. Additive fragmentation coefficient in every fitted forest-specialist design cell. The plate shows βFRAG from the additive family, in which forest amount and fragmentation enter without an interaction term, so the quantity displayed is fragmentation per se after amount control — the estimand at issue in the fragmentation debate. Rows of blocks are fragmentation metrics and columns are formulations, never pooled. Within each block, the nine forest definitions run down the vertical axis and the eighteen grain × extent combinations along the horizontal, so every one of the 1,296 fitted forest-specialist cells appears exactly once. Fill encodes the coefficient and an outline marks cells whose 95 % interval excludes zero; unlike the interaction, the additive family stores standard errors for every cell, so this is an interval criterion rather than an operational support gate. The coefficient is negative in 84.4 % of cells and the interval excludes zero in 46.8 %. Because forest definition, metric, scale and formulation are design factors rather than replicates, the plate displays the distribution of the estimate across the design and does not constitute independent tests.

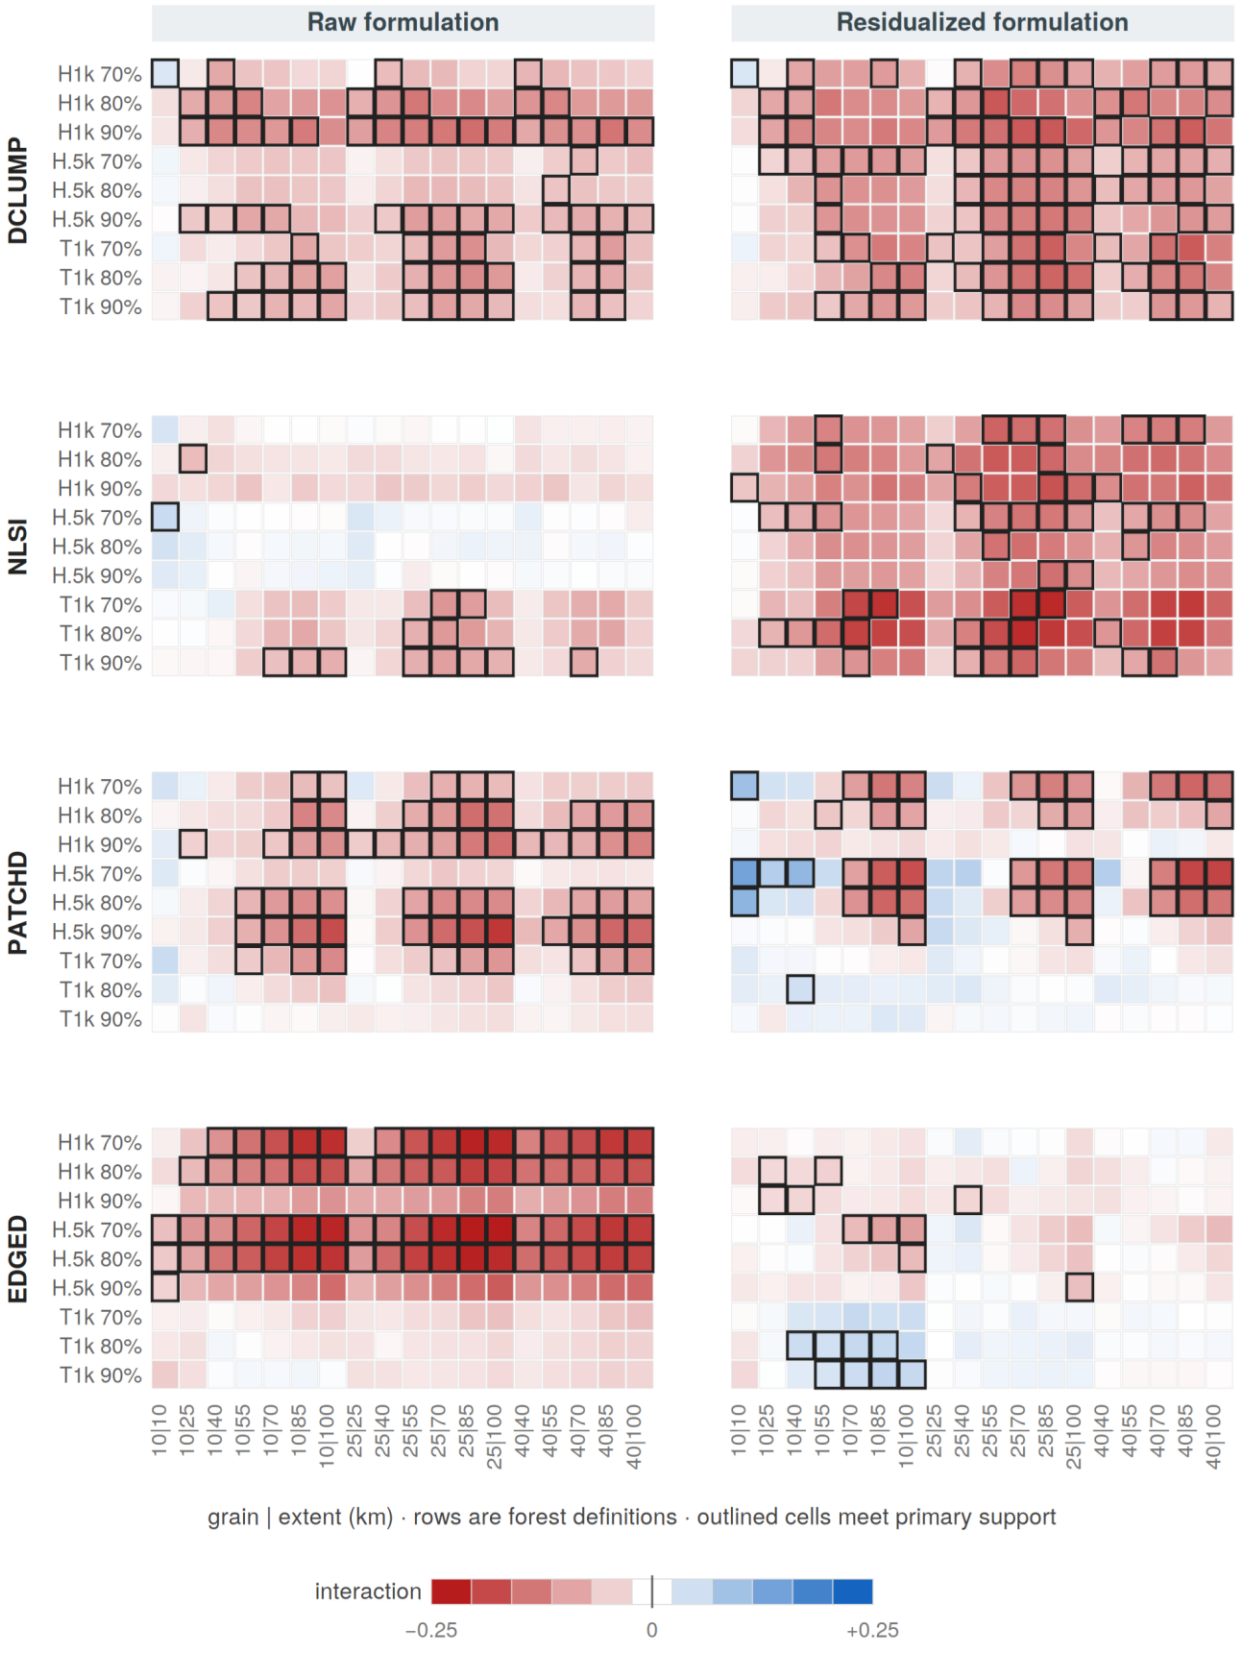

Raw formulation
Residualized formulation
DCLUMP
NLSI
PATCHD
EDGED
H1k 70%
H1k 80%
H1k 90%
H.5k 70%
H.5k 80%
H.5k 90%
T1k 70%
T1k 80%
T1k 90%
10|10
10|25
10|40
10|55
10|70
10|85
10|100
25|25
25|40
25|55
25|70
25|85
25|100
40|40
40|55
40|70
40|85
40|100
grain | extent (km) · rows are forest definitions · outlined cells meet primary support
interaction
−0.25
0
+0.25

Figure 5.S5.4. Interaction coefficient in every fitted forest-specialist design cell. The plate is the companion to Figure 5.S5.3 and uses the same layout, showing the forest-cover × fragmentation interaction rather than the additive fragmentation coefficient. Rows of blocks are fragmentation metrics and columns are formulations. Within each block, the nine forest definitions run down the vertical axis and the eighteen grain × extent combinations along the horizontal, so every one of the 1,296 fitted forest-specialist cells appears exactly once. Fill encodes the coefficient and an outline marks cells meeting the primary-support criterion, which combines a parsimony gain over the matched additive model with structure-coefficient concordance. That is an operational gate rather than an interval, because the stored outputs carry no standard error for the interaction term; interval support for the interaction is available only for forest-specialist richness through the parametric bootstrap, and is reported in Table 5.S5.7; pointwise intervals for the derived marginal fragmentation effect are shown in Figure 5.S5.1. The same design-factor caution applies as in Figure 5.S5.3.

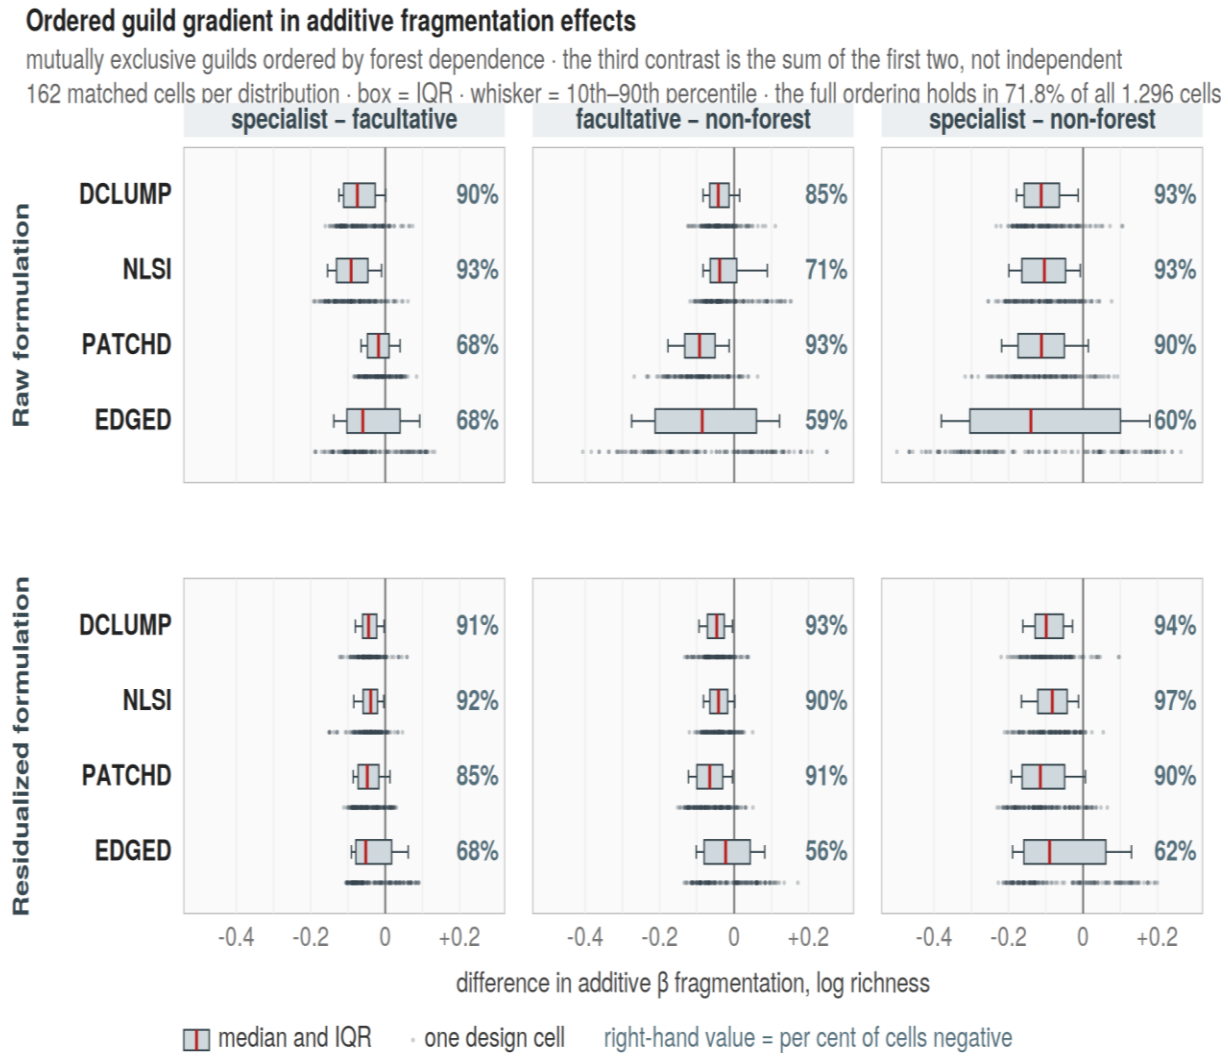


Figure 5.S5.5. Ordered guild gradient in the additive fragmentation coefficient, by metric and formulation. Each distribution summarises 162 matched design cells, one for every combination of forest definition and valid grain-by-extent pair, in which the same landscape, forest definition, threshold, response grain and predictor extent are held constant and only the response stratum differs. The box spans the interquartile range, whiskers reach the 10th and 90th percentiles, the red line marks the median, and individual points show single design cells; the value at the right of each row is the percentage of cells in which the difference is negative. Raw and residualized formulations occupy separate rows because their predictors were standardized separately and are not on a common unit. The three contrasts are algebraically dependent: within every cell the specialist-minus-non-forest difference is the exact sum of the two adjacent contrasts, so the adjacent steps establish the ordering and the endpoint contrast displays its extent, and they are

not three independent comparisons. Median profiles follow the specialist, facultative, non-forest ordering in all eight metric by formulation combinations; cellwise the full ordering holds in 926 of the 1,292 matched cells in which every stratum's additive model both converged and returned a positive-definite Hessian, 71.7 %, ranging from 54.9 % under residualized edge density to 87.7 % under residualized DCLUMP. Four residualized PATCHD cells are excluded because their additive fit did not converge and their Hessian is not positive definite; retaining them would give 930 of 1,296 and would not change the metric-level range.