Essays in Industrial Organization and Labor Economics

Rodrigo González Valdenegro

A dissertation

submitted to the Faculty of

the Department of Economics

in partial fulfillment

of the requirements for the degree of

Doctor of Philosophy

Boston College

Morrissey College of Arts and Sciences

Graduate School

March 2025

Essays in Industrial Organization and Labor Economics

Rodrigo González Valdenegro

Advisors:

Theodore Papageorgiou (Chair)

Charles Murry (Co-Chair)

Richard Sweeney (Member)

Abstract

"When Competition and Labor Market Policy Collide: The Case of the Minimum Wage"
with Juan Luis Fuentes

Policymakers and competition authorities are concerned about the negative effects that labor market power has on workers and the efficient functioning of the labor market. The minimum wage has the potential to curb firms' power, raising employment and wages (Robinson 1933). However, how does using the minimum wage to mitigate firms' labor market power impact social welfare, when these firms differ in labor productivity and degree of product market power? Do the interests of consumers and workers align or conflict? We answer these questions in the context of sector-specific minimum wages. In our framework, raising the minimum wage affects labor and product market equilibrium. First, it may either increase, reduce or have no effect at all on the marginal cost of a given firm, depending on its level relative to the marginal revenue product of labor. Second, the minimum wage influences the strategic pricing decisions of competing firms in the product market, even when the own-cost effect is null. We extend a model of supply and demand in an oligopolistic industry with rich consumer preferences and endogenous marginal costs, to incorporate firms' production, employment and wages. We estimate the model using a unique dataset from the beer industry in Uruguay, where firms of different sizes producing imperfect substitutes compete. We find that the minimum wage that minimizes employers' power and enhances labor market efficiency also reduces consumer surplus and competition in the product market. The presence of firms with labor and product market power poses a challenge for policymakers aiming to curb employers' power without causing unintended harm.

"Regulation by Public Options: Evidence from Pension Funds" with Pablo Blanchard and Sebastian Fleitas

We study the equilibrium welfare effects of using state-owned enterprises (SOEs) to regulate market power. We estimate a dynamic equilibrium model of Uruguay's individual capitalization pension system, where a high-quality perceived SOE competes with private firms. We find that the presence of an SOE reduces equilibrium fees and increases consumer returns. Eliminating the SOE and replacing it with a private firm would more than double its fee and increase the fees of private firms by 8%. Reducing consumer inertia cannot fully offset privatization. Finally, direct price regulation outperforms the SOE as a competitive force in the market.

Dedication

To the person who loved me, encouraged me, and supported me through this long, challenging, and rewarding journey. To Laura.

Contents

List of Tables

List of Figures

Chapter 1

"When Competition and Labor Market Policy Collide: The Case of the Minimum Wage" with Juan Luis Fuentes

1.1 Introduction

Reducing firms' labor market power is a priority for policymakers and competition authorities. The minimum wage is a tool they can use to achieve this goal, capable of reaching a large number of workers while being straightforward to implement. It is well known that when firms have labor market power, the minimum wage has the potential to raise both wages and employment, thus increasing efficiency and welfare (Robinson 1933, Manning 2021, Luduvice et al. 2024). In practice, firms produce imperfect substitutes, exhibit different levels of product market power, and often display differences in labor productivity. How does employing the minimum wage to mitigate labor market power affect firms' profits and the welfare of consumers and workers in this case? Do the interests of consumers and workers align or conflict?

In this paper, we evaluate how the minimum wage affects social welfare in a specific industry, where firms with varying degrees of product and labor market power compete. We focus on two dimensions that shape the ultimate impact of the policy. First, a higher minimum wage can raise marginal costs in some firms (Friedman 1962) and lower them in others (Robinson 1933), depending on its level

relative to the marginal revenue product of labor. Second, the strategic responses of firms to these cost changes alter equilibrium prices in the product market (Weyl et al. 2013), even when the own-cost effect is null. Combined, these factors make the overall effect of a higher minimum wage on social welfare uncertain.

We find that the minimum wage that increases labor market efficiency and maximizes social welfare also reduces consumer surplus in the industry we study. In the presence of firms with different labor productivity, the minimum wage that improves labor market efficiency and lowers costs in one firm may be too high for another. For consumers, the net effect on welfare then depends on how these firms adjust product prices in equilibrium. When the pass-through of cost savings to prices is limited due to product market power, price increases in fringe firms due to higher labor costs may dominate, reducing consumer welfare. In cases of imperfect product substitution, as in this study, welfare losses can be amplified. Addressing labor market power with the minimum wage can result in lower consumer welfare and higher product market concentration, even when enhancing labor market efficiency and social welfare. The existence of firms with power in labor and product markets poses a significant challenge for policymakers aiming to raise wages without causing negative unintended consequences for consumers and competition in the product market.

In order to illustrate how the minimum wage can impact market structure and affect workers, as well as consumer welfare, this paper considers the Uruguayan beer industry. Our setting is ideal, since it is characterized by three distinctive features that highlight the main trade-offs we study. First, there exists an industry-specific minimum wage, and it is high.[1] As of 2015, the minimum wage was

[1] As in approximately 20% of countries (International Labour Organization, 2021), in Uruguay firms categorized within the same sector of economic activity have to comply with a uniform industry-specific minimum wage. In the case of Uruguay, this applies regardless of firm size or location.

approximately 3.7 times the National minimum wage and 2.5 times the average private sector wage for workers without a college degree. Second, the industry is highly concentrated. The largest firm has an aggregate market share of 90% and offers both domestic and imported products.[2] Third, domestic craft brewers have entered the market and expanded the set of available varieties in recent years. Importantly, craft brewers are more labor-intensive than non-craft ones. Despite having a low market share (3%), craft firms account for more than 15% of employment in the industry.

We use a static, partial equilibrium framework, to examine the role that consumer preferences and productivity differentials between firms play in mediating minimum wage effects. We extend a demand and supply model of oligopolistic competition in the product market, to relate firms' endogenous marginal costs with production, employment and wages. Domestic production requires hiring workers, and labor demand depends on the level of production (which we assume equal to sales), unobserved productivity and relative input prices. Multi-product firms set prices to maximize profits, and post wages consistent with those prices, in order to hire workers in the labor market. Then, consumers and workers make discrete choices over products and jobs, respectively.

The minimum wage affects firms' hiring decisions and marginal costs. When the minimum wage does not bind, the marginal cost of hiring an additional worker is increasing in the stock of workers, consistent with an upward-sloping labor supply curve and a non-discriminatory wage policy (Manning 2003) that gives infra-marginal workers positive rents. In this case, output is inefficiently low and equilibrium wages are below the marginal revenue product of labor due to labor market power. Alternatively, when the minimum wage binds and operates mit-

[2]Market share based on overall beer consumption in the country for 2022. It reached 98% in 2003, after the horizontal merger of the two largest firms of the time was not challenged by the Government. See Section (1.2) for more details.

igating labor market power, the marginal cost of hiring decreases and firms are able to add additional workers without having to raise wages. This constitutes the pro-efficiency effect of the minimum wage (Robinson 1933), and it is exhausted when labor supply, labor demand and the minimum wage intersect. Finally, when the minimum wage is high enough, it leads to standard rationing (Friedman 1962), and firms hire where the minimum wage intersects labor demand.

Simultaneously, the minimum wage affects equilibrium prices in the product market. When the minimum wage reduces marginal costs in a firm, the firm may respond by reducing prices, leading competitors that offer close substitutes to reduce prices as well, even when their own costs remain unaffected. Competition in the product market can intensify as a result of a more efficient functioning of the labor market. In contrast to the case of perfect competition, under imperfect competition consumers also benefit from the pro-efficiency effects of the minimum wage. Crucially, how strong this mechanism is relies on whether firms pass cost efficiencies through to prices, in the availability (or absence) of close substitutes, and on differences in labor productivity between competing firms.

We estimate consumer demand and products' marginal costs following Berry et al. 1995. Demand-side moments are built using a random coefficients nested logit model (Brenkers et al. 2006), and supply-side moments assume that firms set products' prices separately in each market to maximize profits. We augment the estimation by adding "second" choice aggregate diversion moments following Conlon et al. 2024, exploiting a natural experiment that removed certain products from some consumers' choice set. This allows us to estimate rich consumer preference heterogeneity, despite using only aggregate shares, working with a single market per period, and without data on the choices of individual consumers.[3]

In our framework, both prices and the labor component of marginal costs

[3]The benefit is that we avoid multiplicity of equilibrium (Bulow et al. 1985).

are endogenous. We instrument prices using cost data at the product level and products' characteristics. In the case of labor costs, we follow the literature on production function estimation (Ackerberg et al. 2015) and instrument with data on fixed and flexible inputs in previous periods. Finally, we calibrate the firm-level labor supply using previous research on set-ups similar to ours (Card et al. 2018, Kroft et al. 2021, Amodio et al. 2021, Casacuberta et al. 2023) and the National Minimum Wage. In Appendix (3.1), we provide robustness checks that show how our main results hold under a broad range of plausible parameterizations.

We perform two sets of counterfactual to study the welfare consequences of alternative labor market policies for consumers, workers and firms. In the first counterfactual, we search for the industry minimum wage that maximizes social welfare. In the second, motivated by the low degree of substitution between craft and non-craft products and the associated differences in firms' size and productivity, we simulate alternative product and labor market equilibrium under distinct minimum wages for craft and non-craft firms. In both cases, we use as baseline the counterfactual equilibrium without minimum wage. To compare counterfactual scenarios, we calculate social welfare as the sum of consumer surplus, workers' rents and firms profits, and assume that it coincides with private welfare (Conlon et al. 2020).

In the first counterfactual, we find that the industry minimum wage that maximizes social welfare is US$ 1,000 per month.[4] This minimum wage level is compatible with an average wage 18% higher than the average under the no minimum wage scenario, and 94% higher than the National minimum wage. The minimum wage that maximizes social welfare creates winners and losers in the industry we study. Because the efficiency gains of mitigating labor market power in the largest

[4]As a reference, a US$ 50 change in the gross monthly wage is equivalent to US$ 26 cents per hour for a full-time worker.

firm are only partially passed to prices due to product market power, consumers would be better off with a lower minimum wage (US$ 550).

The main driver of the negative welfare effect for consumers is a reduction of 14% in the production and sales of craft products in the high minimum wage equilibrium. Craft firms are more labor-intensive, less productive and face a more elastic residual demand curve, so higher labor costs are passed through one-to-one to prices, and result in more than proportional reductions in sales. The difference in labor productivity between craft firms and the largest (non-craft) firm is such that the minimum wage that delivers the highest efficiency gains in the latter is too high for the former. Lower sales in craft firms also imply reduced employment within these firms. However, at the industry level, total wage increases more than compensate for employment losses and result in higher rents for workers. Because craft and non-craft products are poor substitutes, sales and jobs lost in craft firms are not compensated by meaningful changes of opposite sign in non-craft firms. The best uniform minimum wage in terms of social welfare balances between consumers and small craft firms on one side, and workers and the largest firm on the other. Addressing labor market power with the minimum wage raises product market concentration and equilibrium markups in this industry. This result is robust to alternative beer demand specifications and degrees of firms' labor market power.

Given that a uniform minimum wage tends to work well for the non-craft segment but leads to both lower employment in craft firms and higher prices for their products, we next consider the impact of distinct minimum wage categories for craft and non-craft firms. In this case, the minimum wage combination that maximizes social welfare is US$ 550 for craft firms and US$ 1,000 for non-craft firms. Notably, this policy increases consumer surplus and employment relative to the

best uniform minimum wage, though with a smaller increase in effective wages for incumbent workers in craft firms. The low non-craft to craft diversion ratio implies that the lower equilibrium prices of craft products do not impose a significant competitive pressure on other firms in the product market. The availability of more affordable craft products also does not prompt the largest firm to pass on a higher share of cost savings to consumers. Overall, social welfare is higher with separate minimum wage categories primarily because it boosts consumer surplus, employment and profits at craft firms, with negligible negative effects on profits and employment effects at non-craft firms.

We contribute to two strands of the literature. First, we contribute to the literature that studies the interactions between labor and product markets in specific industries (Kroft et al. 2021, Montag et al. 2022), considering the possibility that certain firms may have price-setting power in both markets (Kroft et al. 2021, Avignon et al. 2022, Rubens 2023). We do not consider entry and exit (Montag et al. 2022), but we incorporate a model of the labor market as in Kroft et al. 2021, with the novelty of studying how minimum wages affect both labor and product market equilibrium. Additionally, we use a richer specification of consumer preferences than Kroft et al. 2021. This allows us to analyze the trade-offs that the minimum wage entails for consumers when firms with different labor productivity produce imperfect substitutes. To our knowledge, we are the first to incorporate a measure of consumer welfare based on observed choices in the study of minimum wage consequences. The idea that minimum wages can affect competition in the product market has been studied before (Williamson 1968, Salop et al. 1983). In our model, this policy is not merely a distortion that negatively affects consumers through higher costs and prices. Instead, the existence of labor market power gives the minimum wage the ability to reduce prices through cost efficiency and

enhanced competition. In practice, this channel is not sufficiently strong in the industry we study, consumer welfare would be higher if the policymaker accepted a certain degree of labor market power.

Second, our work is related to the literature that investigates the consequences of general (Card and Krueger 1993, Aaronson et al. 2018, MaCurdy 2015, Dube 2019, Harasztosi et al. 2019, Berger et al. 2022b, Luduvice et al. 2024)[5] and industry-specific (Haucap et al. 2001, Braun 2011, Martins 2021, Hijzen et al. 2020, Card and Cardoso 2022, Bello et al. 2022, Hermo 2023, Casacuberta et al. 2023) minimum wages. In particular, our work contributes to a recent strand that studies the effects of minimum wage changes beyond employment and prices (Harasztosi et al. 2019, Berger et al. 2022b, Dustmann et al. 2022, Luduvice et al. 2024). Relative to these papers, our framework enables us to focus on the role that product market power plays in mediating minimum wage effects on total welfare. In this case, we contribute by including into the analysis consumer welfare, and not just workers' rents and firms' profits. Additionally, our modeling strategy retains the possibility of studying multiple outcomes simultaneously as in general equilibrium models (Braun 2011, Berger et al. 2022b, Bello et al. 2022), but using rich consumer preferences. In this context, we show how the degree of product differentiation can be a relevant channel that affects the reallocation of workers from smaller/lower productivity to larger/higher productivity firms documented in previous studies (Harasztosi et al. 2019, Dustmann et al. 2022). Finally, the modeling approach we use can rationalize previous findings about the effects of industry-specific labor cost changes. There is evidence that following a change in labor costs, some firms might not only face a higher production cost, but also produce more (Dodini et al. 2023). This can occur naturally in our model as the equilibrium outcome of simultaneous cost increases and positive residual demand

[5]See Neumark et al. 2022 for a recent survey.

shocks.

The paper proceeds as follows. First, we provide institutional details of the industry and of the labor market institutions in place. Second, we describe the data and display the most relevant descriptive information that rationalizes our modeling decision. Third, we present the model we employ in detail, describe the estimation procedure and show the main results. Finally, we study counterfactual scenarios, discuss our results and then conclude.

1.2 Institutional Environment

In Uruguay, minimum wages apply by sector of economic activity, similar to over 20% of countries.[6] Sectoral minimum wages for multiple occupations are established in Collective Bargaining Agreements (CBA) between representative unions and employers associations every 2-3 years.[7]. A National Minimum Wage is in place as well, and serves as a significant reference point in the labor market. Approximately 18% of the workers earned less than 1.5 times the National Minimum Wage in 2018.[8]

We study the sector that covers firms producing and selling beer and non-alcoholic beverages (water, juice and soft-drinks).[9] In particular, we define the beer industry as a subgroup of this sector that contains firms whose primary activity is to produce or to import beer. Craft and non-craft domestic producers are included in this sector and in the beer industry.

For the purpose of this paper, the minimum wage of the sector (and the indus-

[6] International Labor Organization, 2021.

[7] An occupation broadly refers to a set of tasks and responsibilities within a job. Other dimensions such as region, firm size, or trade exposure are not taken into account.

[8] PIT-CNT, 2018

[9] Group 1, Sub-Group 9, Table 1, Tray 1 according to the classification of the Ministry of Labor and Social Security.

try) is the lowest among the four occupation-specific minimum wages that govern the wages of manufacturing workers.[10] As of 2019, the minimum wage we consider was approximately 3.7 times higher than the National Minimum Wage and 2.5 times the average wage in the private sector for workers without a college degree.[11] According to Ramirez et al. 2022, this is among the highest sector-specific minimum wages in the country.

The beer industry is highly concentrated, with a single firm accounting for 90% of sales.[12] The largest firm in the industry is the consequence of two unchallenged horizontal mergers that occurred between 1998 and 2003 between the three national firms of the time.[13] The industry went from an oligopoly with three firms of similar scale to a monopoly over a 5-year period. Along this paper, we will refer to this firm as the largest firm, the largest domestic producer or the largest (and only) non-craft domestic producer.

The largest firm participates in the bargaining of minimum wages for the sector, but does not control the employer's side; other firms participate as well. In contrast, craft brewers are part of the sector but do not participate in negotiations. However, they must comply with minimum wage regulations. The Government extends agreements to non-participating firms within sectors once they are signed. The extension does not have formal requirements on firms' representation or an analysis of the consequences that the extension may have on non-represented firms.

[10]This sector also defines minimum wages for two auxiliary processes (sales and administrative), but we do not consider them as part of the production process.
[11]This will be proven useful in the estimation process.
[12]See summary statistics in Table (1.1)
[13]Melgar et al. 2004

1.3 Data

We leverage multiple data sources to connect labor and product markets. We use these sources to build a dataset that contains prices and quantities of inputs and final products of firms participating in the Uruguayan beer industry.

1.3.1 Data Sources

We use sales data provided by a Uruguayan firm that tracks multiple retail stores to estimate consumer preferences.[14] The data contains monthly observations at the store-brand-size-flavor level between April 2015 and October 2022. The name of the store is anonymous. We complement sales data with production statistics. For craft brewers, we use publicly available data from the brewers' association and proprietary information provided by some of these firms.[15] For non-craft domestic brewers (the largest firm in the industry), we use publicly available data from the Tax Authority about beer volumes associated with excise tax payments by origin (domestic vs imported) and subtract from the total domestic, the output series calculated for craft firms.[16] There are no other non-craft domestic producers in the country. Unlike sales, our production data aggregates by groups of firms (craft and non-craft). However, considering that craft brewers entered the market during the same period and that their scale of operations is similar, this aggregation preserves the heterogeneity between craft and industrial production that is relevant for this study.

We exploit three data sources with information on wages, employment and firms' capital stocks. First, we use anonymized matched employer-employee data

[14] IdRetail
[15] AMAU for Asociacion de Microcervecerias Artesanales del Uruguay.
[16] We check implied figures with other statistics for consistency.

provided by the Ministry of Labor and Social Security for the period 2015 to 2022, covering formal workers in the private sector. The data contains monthly records on workers' basic demographics, employment status, earnings, and sector-job classification for minimum wage purposes. For the case of manufacturing firms, using the number of employees and the ISIC code, we classify them into two groups based on their size.[17] According to our discussions with craft brewers, the group of small firms contains the relevant employment information we need to associate sales in the product market with employment and wages in the labor market. The second source we use is the Annual Survey of Economic Activity (ASEA) conducted by the INE.[18] This survey contains balance sheet data -revenues, expenditures, assets, debt, capital and employment information -number of employees by job category and sector of activity, labor costs- for large firms. Finally, to build a series of the capital stock, we use data from the ASEA, balance sheets and financial statements published by the National Internal Audit Office, as well as customs records of imported capital goods.[19]

We complement previous data sources with information on quantities and prices of materials and other relevant prices (exchange rates, gas). For materials (hops, bottles, and cans), we use customs records at the firm level. We leverage the fact that except for malt and certain bottles, all other inputs are imported. In the case of malt, we use the export price as a measure of the opportunity cost of the input.

[17]International Standard Industrial Classification
[18]Uruguayan National Institute of Statistics
[19]See Appendix (3.1) for details.

1.3.2 Descriptive Information

Beer Sales

We define a product j as a serving size q_j of a combination of brand and flavor (Lager, Ale), similar to Fan et al. 2020. For example, "Heineken Lager" is one product in our data. We define a market m as the consolidated group of large supermarkets in the Capital Montevideo and its metropolitan area in a given period t. The advantage of using supermarkets is that they provide consumers with access to a wide variety of beer products. We assume that market size $\mathcal{M}_m$ is proportional to the highest sales volume observed in our data during the period we study (4.0x), in line with Miller et al. 2017. A product must be available in at least 35 markets to be included in our sample. Sales data contains information on units and revenues, so we calculate monthly prices as the ratio of the two. We standardize product size and express units in servings of 12oz/355mL. Unlike in the US, more than 90% of sales in Uruguay are single units rather than packs, so we focus exclusively on those. We exclude non-alcoholic or low alcohol products.[20] Finally, except for certain imported brands during short periods of time, in our sample products match one to one with firms.

We show sales summary statistics in Table (1.1), which highlights the level of market concentration, the entry of domestic craft brewers, and the increase in the consumption of craft and imported brands over time. The largest firm in the industry is responsible for roughly 9 out of 10 servings sold in the beer market, and offers both domestic and imported products[21] The majority of sales consist of four non-craft brands that this firm produces domestically. To manufacture

[20] Alcohol by Volume $\leq 1\%$.
[21] See the Institutional Section (1.2) for details.

these brands, the company needs to hire workers in the labor market.[22] Imported products are offered by the largest firm and by six non-craft importers that jointly account for 8% of sales.

Table 1.1: Beer Market Summary Statistics

		Quantity			Price			Revenues		
		(# Servings, M)			(US$/Serving)			(US$ M)		
Characteristic		2016	2019	2021	2016	2019	2021	2016	2019	2021
Craft	Non-Craft	14.1	16.2	15.3	1.17	1.27	1.30	20.8	22.0	17.3
	Craft	0.1	0.3	0.4	2.53	2.60	2.53	0.4	0.8	0.5
Style	Lager	13.4	14.3	13.1	1.17	1.24	1.28	19.6	19.0	14.4
	Ale	0.9	2.2	2.6	1.47	1.62	1.55	1.6	3.7	3.3
Origin	Domestic	9.9	10.6	9.0	1.050	1.14	1.18	13.0	12.9	9.1
	Imported	4.4	5.9	6.7	1.49	1.56	1.50	8.2	9.9	8.7
Firm	Largest	13.0	15.0	13.9	1.15	1.25	1.28	18.9	20.2	15.5
	Other	1.2	1.5	1.7	1.51	1.65	1.64	2.4	2.6	2.3
Total		14.3	16.5	15.7	1.19	1.29	1.32	21.2	22.8	17.8
HHI		8,357	8,324	8,076				7,924	7,877	7,624

Note: Supermarkets' sales in the Capital and its Metropolitan Area. Real Uruguayan Pesos converted to US$.

The craft sales in the data consider the sum of the five largest craft breweries in the country. These firms represent more than 80% of craft production and it is on them that we focus in this study. A brewer is considered craft when it receives a certification issued by the craft brewers association, and therefore, a product is characterized as craft when it is produced by a craft brewer.[23] Craft brewers usually advertise their craft condition on the front label of their products to make this characteristic salient to consumers.

[22] Three of the four brands have a longstanding tradition in the Uruguayan beer market, they were introduced in 1890, 1947 and 1956.

[23] The label certifies that the brewer "is an independent Uruguayan craft brewery, which is neither part of nor affiliated in any way with a multinational corporation" Different from the US, the craft certificate does not include a cap on market share. However, it includes a revision of production standards and is more restrictive on the participation of multinationals in ownership (0% in Uruguay versus up to 25% in the US).

Although increasing over time, craft beer represents a minor share in this market (1.9%) and is more expensive than non-craft products (94.5% on average in 2021).[24] The increase in the consumption of craft products has been part of a broader trend, as the increase in the consumption of new flavors (Ales) and imported brands shows. Seasonality plays an important role in demand, though seasonal patterns differ between beer types. Craft consumption increases in the colder months, while the consumption of non-craft products increases during the summer.

Productivity, Employment and Wages

Industrial (non-craft) producers offer products at lower prices than craft brewers. Among the reasons that can explain this is the gap that exists in labor productivity. In 2021, the average labor productivity of manufacturing workers in the largest firm (measured as production per worker during a year) was roughly 5 times higher than in craft breweries. The capital-labor ratio is higher in the largest producer too, and this may contribute to explain (partially) the difference in labor productivity. However, craft brewers have invested in physical capital to automate processes that workers used to do, allowing them to reduce the gap of capital per worker during the period we study.[25]

[24] According to (Fan et al. 2020) and (Hidalgo 2023), the price gap was ≈55% for California and 69% for the US, respectively. Similarly, craft products represented ≈ 10% of beer consumption in California (Fan et al. 2020) and 13% in the US (Hidalgo 2023).

[25] Considers only the capital stock used for manufacturing beer.

Table 1.2: Production Summary Statistics

	Productivity		Capital		Wages		Labor Cost	
	(Q/L)		(K/L)		$(\frac{\bar{w}}{\underline{w}})$		$(\frac{wL}{Q})$	
	2016	2021	2016	2021	2016	2021	2016	2021
Non-Craft	276	220	247.1	156.7	2.8	2.6	7.4	8.7
Craft	23	45	52.4	107.2	0.5	0.6	21.7	14.1

Note: Manufacturing statistics for domestic producers. Craft firms consider the five firms contained in sales summary statistics. Labor productivity in liters per worker-year. Capital per worker in constant US$ 000's of 2019, considers manufacturing capital only. The ratio of average to minimum wages considers total labor cost of workers involved in the manufacturing process. Ratio calculated using raw data, does not adjust for reported hours. Average labor cost expressed in constant US$ cents per serving.

The largest firm is not only more productive, but pays higher wages. Although the minimum wage in the industry is high compared to other industries in the country (Ramirez et al. 2022), the largest firm pays average wages above that threshold. This may be related to the fact that in this industry there are multiple manufacturing occupations and to other benefits like tenure within the firm, for example. Still, the minimum wage binds for new workers. According to our data, new workers earn gross entry wages that are 1.0-1.3 times the minimum wage. Craft brewers report gross wages below the minimum, but in most cases associated with part-time workers. These firms seem to comply with the industry minimum wage as well, once we adjust for the number of hours of work they report.

Labor costs are quantitatively more important for craft brewers. The average manufacturing labor cost per serving is 1.5-3.0 times higher in craft than in non-craft domestic firms during the period we study. Despite paying lower average wages, the fact that craft firms are more labor-intensive dominates. This explains how, even when they represent approximately 2% of national beer consumption

(3% of production), they account for 15% of manufacturing jobs in the industry. For the same reason, a uniform minimum wage can increase the costs of craft firms relatively more than those of non-craft firms.

Figure 1.1: Average Labor Cost - US$ per serving

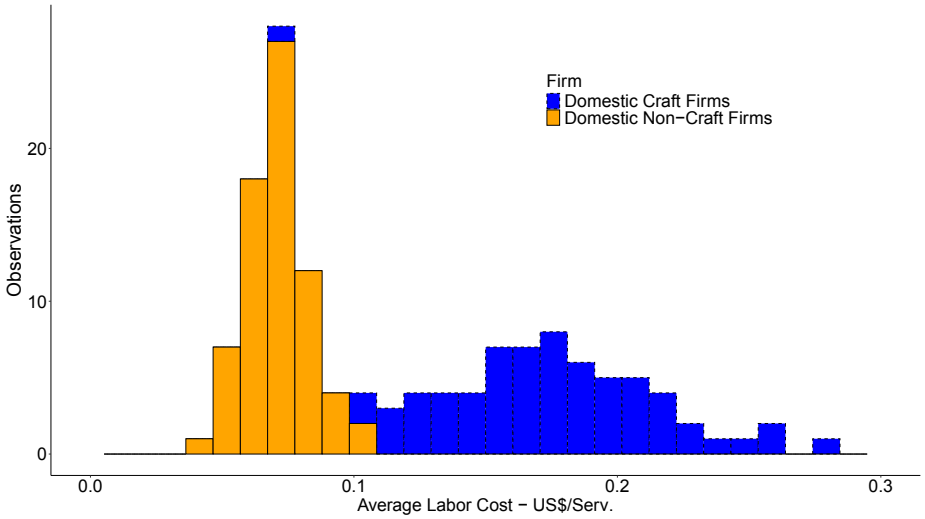

Imported Products and Materials

We use customs data to calculate import prices per serving of imported products. In most cases, we are able to match custom records to products in the output market using the firm name, country of origin, and product descriptions. For those that we cannot, we assign the average import price from that country and firm to all products within the same group. Figure (1.2) shows the relationship between the import price and the retail price in our sample. In addition, we use input import prices at the firm level to estimate cost functions. In Figure (3.1a) we display hop prices, one of the main materials that firms use. The cost differences between domestic producers are also the result of paying different prices for the materials they use.

Figure 1.2: Input prices

(a) Import and Retail Price

(b) Hops Import Price

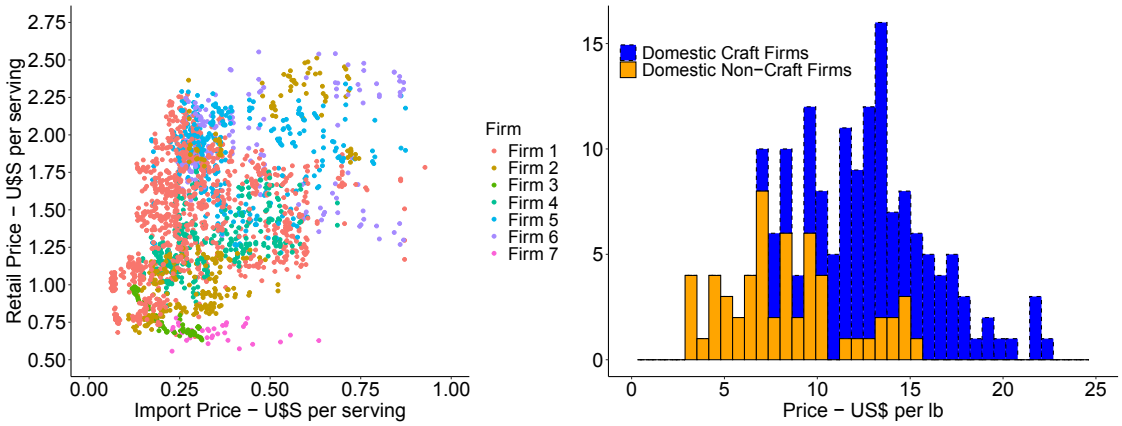

In what follows, we outline the model that we use to represent the most salient characteristics of the industry. Firms, differing in labor productivity and scale, produce imperfect substitutes, charge different prices, and incur varying labor costs, but adhere to the same industry minimum wage.

1.4 Model

We use a partial equilibrium, static framework, where firms in an oligopolistic industry sell differentiated products and compete by setting prices (Berry et al. 1995). Firms that produce locally hire workers in the labor market to manufacture the products they sell to consumers in the product market. Alternatively, firms that sell imported products do not require hiring manufacturing workers. Manufacturing is the only process that affects variable costs and needs labor in our framework. A policymaker sets the minimum wage at the beginning of each period, then firms take it as given when making pricing and hiring decisions. For simplicity, we assume that consumers and workers belong to separate households.

1.4.1 Consumers

Consumers -indexed by i- make discrete choices among available products $\mathcal{J}_{mt}$ and a no-purchase outside option -indexed by 0- in retail market m at time t[26]. Consumers choose the product j that maximizes indirect utility, which is represented by a linear index that considers prices p_{jmt}, exogenous attributes x^d_{jmt}, unobserved quality ξ_{jmt} and idiosyncratic preferences ε_{ijmt}.

$$u_{ijmt} = \gamma_{imt}\mathrm{x}^d_{jmt} + \alpha_{imt}p_{jmt} + \xi_{jmt} + \varepsilon_{ijmt} \tag{1.1}$$

Preferences for products in certain categories h might be correlated (Cardell 1997), so we include a one-level nest specification with $\varepsilon_{ijmt} = \bar{\varepsilon}_{ih(j)mt} + (1-\rho)\bar{\varepsilon}_{ijmt}$ distributed Generalized Extreme Value. We consider parametric distributions of taste parameters (γ_i, α_i) that account for consumers' income y_{imt} and unobserved preference heterogeneity. We interact both types of heterogeneity with product characteristics, indexed by k. We specify the model to estimate flexible substitution patterns and inform the extent of competition between firms. In our setting, this is a relevant channel for understanding the welfare effects of the minimum wage. The full specification for $(\gamma_{imt}, \alpha_{imt})$ is:

$$\alpha_{imt} = -\exp(\alpha_0 + \sigma_\alpha \nu_{i\alpha} + \kappa_\alpha y_{imt}) \tag{1.2}$$

$$\gamma_{imt} = \sum_k (\gamma_{0,k} + \gamma_{y,k} y_{imt} + \sigma_k \nu_{ik}) \tag{1.3}$$

To answer our question we need plausible estimates of how marginal cost changes are passed-through to prices. In principle, the pass-through rate of a particular product depends on the own-price elasticity and on the curvature of residual

[26] A retail market m identifies a geographical location.

demand. However, in multinomial or mixed logit models with constant price sensitivity, the curvature of demand is imposed rather than estimated and the pass-through rate is incomplete. The specification of α_i overcomes this limitation (Miravete et al. 2023) by allowing ex-ante for a broad range of estimable demand curvatures. We normalize the utility of the outside option $u_{i0mt} = 0$ and draw unobserved taste for product characteristics ν_{ik} from the standard normal distribution. For a given month, market shares are given by integrating over the distribution of consumers who vary in income, unobserved tastes for characteristics, and idiosyncratic error terms,

$$s_{jmt} = \int \frac{\exp\left(\alpha_{imt} p_{jmt} + \gamma_{imt} \mathrm{x}_{jmt} + \xi_{jmt}\right)}{1 + \sum_{k \in \mathcal{J}_{mt}} \exp\left(\alpha_{imt} p_{kmt} + \gamma_{imt} \mathrm{x}_{kmt} + \xi_{kmt}\right)} dF(i) \tag{1.4}$$

When estimating the model, we match "second" choice aggregate diversion moments that exploit a natural experiment that removed a group of products from the choice set of certain consumers. We follow the procedure developed by (Conlon et al. 2024). In this case, we calculate "second" choice shares conditional on a given first choice, by integrating consumers' choice probabilities when the group of products is removed, weighted by consumers' probability of making that first choice (conlonmortimer2021). When the group of products $\mathcal{J}_{rem}$ is removed in mt, the diversion ratio to product k is:

$$D_{\mathcal{J}_{rem} \to k,mt} = \frac{s_{kmt \setminus \mathcal{J}_{rem}} - s_{kmt}}{\sum_{j \in \mathcal{J}_{rem}} s_{jmt}} = \frac{1}{\sum_{j \in \mathcal{J}_{rem}} s_{jmt}} \int \frac{s_{ikmt}\left(\sum_{j \in \mathcal{J}_{rem}} s_{ijmt}\right)}{1 - \sum_{j \in \mathcal{J}_{rem}} s_{ijmt}} dF(i) \tag{1.5}$$

The choice probability of an individual consumer (s_{ijmt}) is calculated according to the term inside the integral in (1.4), and the term $s_{kmt \setminus \mathcal{J}_{rem}}$ represents the market share of product k after the removal.

1.4.2 Workers

Workers -indexed by n- make discrete choices over firms in the labor market at time t, considering posted wages w_{gt}, amenities z_{gt} and idiosyncratic unobserved preferences ψ_{ngt} (Kroft et al. 2021, Berger et al. 2022a). We assume that there is a unified labor market, so only the time dimension matters.[27] Firms do not observe workers' preferences and can not wage-discriminate, so every worker in g earns the same wage. However, workers' next best alternatives differ, so infra-marginal workers receive positive rents in equilibrium.[28] Workers' choices maximize an indirect utility function:

$$u_{ngt} = \ln w_{gt} + \ln z_{gt} + \theta\psi_{ngt} \tag{1.6}$$

In this framework, workers view firms (jobs) as imperfect substitutes and firms have to compete to attract them. The degree of imperfect substitution is measured by the variance parameter θ that controls the relationship between jobs' vertical and horizontal differentiation. When there is pure vertical differentiation ($\theta = 0$), the best job is the one that offers the highest $u_{.gt}(w_{gt}, z_{gt})$. In this case, all workers want to work for the firm that offers that job only. Alternatively, when there is pure horizontal differentiation ($\theta \to \infty$) each worker wants to work at the firm where ψ_{ngt} is the highest, irrespective of wages or amenities. In our model, an imperfect labor market implies that both types of differentiation coexist ($0 < \theta << \infty$). To calculate the probability of accepting a job in a particular firm, we assume that ψ_{ngt} follows a Gumbel distribution and that all workers are employed. We use the probability and the labor market size ($\bar{L}_t$) to calculate the inverse labor supply

[27] This assumption is based on the fact that we do not observe workers' locations in our payroll data. However, apart from two cases, all factories are situated in the region that includes the Capital (Montevideo) and its Metropolitan Area

[28] They depend on individual unobserved $\psi_{ng't}$

curve that firm g faces. The function relates the wage it offers (w_{gt}) with the number of workers it attracts (L_{gt}) and information about amenities and general labor market conditions (U_{gt}).

$$w_{gt} = U_{gt} L_{gt}^{\theta} \tag{1.7}$$

In the model, the labor supply elasticity is constant and common to all firms ($\eta_w^L = 1/\theta$), and so is the markdown in the no-minimum wage scenario ($1 + \eta^{-1} = 1 + \theta$).[29] Notice how pure vertical differentiation implies an infinitely elastic labor supply, a competitive labor market and no markdown; while pure horizontal implies an inelastic labor supply, a captive labor market and an infinite markdown.

Alternatively, the term U_{gt} is firm-specific and captures the effect of amenities (z_{gt}) and labor market conditions (Ξ_t).

$$U_{gt} = z_{gt}^{-1} \Xi_t \quad \text{with} \quad \Xi_t = \left(\frac{\sum_{k \in G} (w_{kt} z_{kt})^{1/\theta}}{\bar{L}_t} \right)^{\theta}$$

We assume that when a firm posts a wage, it does not internalize its influence on posted wages in other firms: firms are strategically small in the labor market.[30] This is consistent with our empirical application, where the minimum wage is industry-specific, but the relevant labor market might be broader once we consider workers' mobility. Workers are not constrained to work in the beer industry only.[31]

[29] A constant markdown might still imply different employment and output distortions with respect to the no-markdown equilibrium. The slope of labor demand may -as it is the case in our empirical application- differ between firms.

[30] Formally, this implies that $\frac{\partial \Xi_t}{\partial w_{kt}} = 0$.

[31] Firms in the beer industry are a subset of the total number of firms in the labor market G

1.4.3 Firms

In our model, firms may participate in multiple retail markets -indexed by m- at time t. Nonetheless, production is centralized and firms produce at a single factory: g identifies firms, production facilities and jobs. Their portfolio may include imported and domestic products $\mathcal{J}_g = \mathcal{J}_g^I \cup \mathcal{J}_g^D$. We assume imported products have constant marginal costs, but for domestic products are endogenous and vary based on how much the firm produces at time t. In this case, costs result from a cost function $\mathcal{C}_{gt}^D(\cdot)$ that depends on the level of production Y_{gt}^D required to serve all markets, which we assume equal to sales $Y_{gt}^D = Q_{gt}^D$. To characterize products' marginal costs and solve the pricing problem of the firm, we describe first how firms produce and the cost function of domestic products. The marginal cost of these products is what connects product and labor markets in our model.[32]

Production

The scale of a firm's operation depends on the total production of domestic products at t. The firm uses workers L_{gt}, capital K_{gt} and materials M_{gt}, according to a production function:

$$Y_{gt}^D = H_g(A_{gt}, L_{gt}, K_{gt}, M_{gt}; \beta_g) \tag{1.8}$$

Technology and productivity are summarized by the vector of parameters β_g and an unobserved component A_{gt}. We further assume that workers and capital are homogeneous and interchangeable in the production of different products within the firm. These assumptions represent the beer industry appropriately according

[32]For the same reason, imported products do not require any labor in our framework. Consistently, in our empirical application we separate production workers from the rest, which we consider part of the fixed cost of the firm.

to market participants, and are in line with previous studies for the US and Canada (Grieco et al. 2018, De Loecker et al. 2022). The aggregate production function is consistent with fixed labor and capital proportions that add up to one, once all products in all markets served at t are considered: $\mathrm{a}^D_{jmt} = \frac{L_{jmt}}{L_{gt}} = \frac{K_{jmt}}{K_{gt}}$ (DeLoecker et al. 2021). This aggregate specification allows us to study hiring decisions at the firm level.

Following De Loecker et al. 2022, we parameterize the production function using Leontief technology with fixed proportions between materials and value added. This captures the fact that, for example, hops and malt cannot be replaced for capital or workers. Materials $M_{k,gt}$ may include hops, bottles and malt, and vary according to a composite technology term $\beta^k_{M,g}$. Value added is Cobb-Douglas with Hicks-neutral productivity A_{gt} and constant capital and labor elasticities with respect to output $\{\beta_K, \beta_L\}$. The aggregate production function is:[33]

$$\sum_{m\in\mathcal{M}}\sum_{j\in\mathcal{J}^D_{gmt}} y^D_{jmt} = Y^D_{gt} = A_{gt}L^{\beta_L}_{gt}K^{\beta_K}_{gt}\sum_{m\in\mathcal{M}}\sum_{j\in\mathcal{J}^D_{gmt}}\left(a^D_{jmt}\right)^{\beta_L+\beta_K} + \sum_k \beta^k_{M,g}M_{k,gt} \tag{1.9}$$

Marginal Cost

We use the production function to characterize the cost function behind the marginal costs of domestic products.[34] In this case, marginal costs result from the cost function that solves the short-run cost-minimization problem of the firm. For this problem, we treat labor and materials as variable inputs, and capital as fixed (Ackerberg et al. 2015). We assume that marginal costs are endogenous

[33] We define $\mathcal{M}$ as the set of retail markets (locations) served by the firm and assume it remains constant during the period we study.

[34] For simplicity, we present the model assuming that β_L and β_K are common across firms, that returns to scale are constant ($\beta_L+\beta_K = 1$), and that there are no other product or firm-specific cost components. We relax some of these assumptions in the empirical section.

$c_{jmt}(Q_{gt}^D)$ because labor costs are a function of sales through employment and production, according to the production function (1.9).[35] The variable cost of production of national products in a given month is:

$$\mathcal{C}_{gt}^D(Q_{gt}^D) = w_{gt}\big(L_{gt}\big)L_{gt} + \sum_k \beta_{M,g}^k M_{k,gt} \tag{1.10}$$

Considering that employment is an explicit function of sales $L_{gt}(Q_{gt}^D)$ through equation (1.9), the wage policy w_{gt} is the only element needed to specify how labor and product markets are related through $c_{jmt}(\cdot)$.

We consider three cases. First, when the minimum wage $\underline{w}_t$ binds on the labor demand, the firm faces an excess of labor supply and pays workers the minimum wage applicable in period t. Second, when the minimum wage doesn't bind, the firm faces the labor supply curve of equation (1.7). In this case, the firm hires workers where marginal cost equals marginal revenue, and pays employed workers on the labor supply. Here, firm-level production is inefficiently low and the wage is below the marginal revenue product of labor. Finally, when the minimum wage binds on the labor supply, the firm pays employed workers the minimum wage, but different from the first case, it has an excess of labor demand. However, raising wages is not optimal either, given that marginal cost is higher than marginal revenue due to the no-discrimination wage policy. The policymaker can remove the constraint by increasing $\underline{w}_t$, raising effective wages, employment and firm profit at the same time. The marginal cost of domestic products for cases one and two,

[35]For imported products, marginal costs are constant and do not require special consideration.

respectively, is[36]:

$$\frac{\partial \mathcal{C}^D(Q^D_{gt})}{\partial q_{jmt}} = \frac{1}{\beta_L}\frac{\underline{w}_t L_{gt}}{Q^D_{gt}} + \sum_k \frac{p^M_{k,gt}}{\beta^k_{M,g}} \tag{1.11}$$

$$\frac{\partial \mathcal{C}^D(Q^D_{gt})}{\partial q_{jmt}} = \frac{(1+\theta)}{\beta_L}\underbrace{\frac{w_{gt} L_{gt}}{Q^D_{gt}}}_{lc_{gt}(\cdot)} + \sum_k \frac{p^M_{k,gt}}{\beta^k_{M,g}} \tag{1.12}$$

Notice that the marginal cost of q_{jmt} contains a labor cost element lc_{gt} that is common to all national products in all markets served at t.[37] The wage policy and the minimum wage affect the cost of domestic products and equilibrium in the product market through this term. When the minimum wage does not bind (1.12), the term $(1+\theta)$ represents the equilibrium markdown $(1+1/\eta^L_w)$, which is decreasing in the elasticity of labor supply.

Moreover, in the model higher capital K_{gt} and higher productivity A_{gt} firms have lower marginal costs (higher Q_{gt}/L_{gt}). In our empirical application, the labor productivity of non-craft firms is approximately ten times larger than in craft ones, so a uniform minimum wage may have different implications in marginal cost levels. Finally, marginal costs might differ also because firms pay -as we observe in the data- different prices for materials $p^M_{k,gt}$ or use different technology $\beta^k_{M,g}$.[38]

Marginal costs might contain other elements unrelated with the production process as well.[39] In order to parameterize marginal costs, we isolate the en-

[36]The difference with case three is related to the shadow cost it introduces to reflect the hiring constraint and prevent firms from hiring additional workers. In this case, the marginal cost has a discontinuity and the equilibrium markdown is in the range $(0,\ 1+1/\eta^L_w)$. We handle this case in the empirical section.

[37]The result uses the fact that Q_{gt} is the sum of q_{jmt} across products j in $\mathcal{J}^D_{gmt}$ and across markets m in $\mathcal{M}$, so $\frac{\partial Q_{gt}}{\partial q_{jmt}} = 1$

[38]In the case of prices differences, they may originate in quality -craft brewers may use better quality hops- or bulk buying. For this reason, we do not use the price of hops when instrumenting for endogenous prices in the demand model.

[39]Different from other countries like the US, firms in the beer industry can be vertically integrated and distribute and sale their own products. As a result, marginal cost may contain distribution and retail related costs.

dogenous labor cost lc_{gt}, arrange the full set of exogenous attributes in $\mathrm{x}_{jmt}^{s,ex}$, and include an unobserved cost shock ω_{jmt}:

$$mc_{jmt}(Q_{gt}) = \lambda_j^{lc} lc_{gt}(Q_{gt}) + \lambda_j^{ex} \mathrm{x}_{jmt}^{s,ex} + \omega_{jmt} \tag{1.13}$$

The vector of coefficients $\{\lambda_j^{lc}, \lambda_j^{ex}\}$ is allowed to vary by product, to acknowledge the difference between domestic and imported products in our framework.[40] According to our specification of the production process, the coefficient λ_j^{lc} has a structural interpretation: it depends on the elasticity of output with respect to labor (β_L), and potentially, on the elasticity of labor supply ($\eta = \theta^{-1}$), depending on whether the minimum wage has an active or passive role in the equilibrium under consideration.[41] Additionally, the vector $\mathrm{x}_{jmt}^{s,ex}$ may contain both product specific (the price of materials for domestic products or the import price for imported products) and common attributes (gas prices or taxes).

Firm problem

We use previous results on products' marginal cost to solve the pricing problem of the firm. We assume that firms play a static, full information, simultaneous move pricing game each period t. Firms know the distribution of shocks for consumers (ε) and workers (ψ) but not their realizations.

The timing of events is as follows. First, firms choose prices p_{jmt} to maximize profits in each product market, and make simultaneous wage offers w_{gt} in the labor market, consistent with those prices.[42] Consistency implies that, for any candidate vector of equilibrium prices $\mathbf{p}_t$, firms need to hire workers to produce enough

[40] In fact, $\lambda_j^{lc} = 0$ for imported products.

[41] This notation can accommodate the case of firm-specific production technology $\beta_{L,g}$.

[42] For firms offering imported products only, this is the traditional pricing game except for the fact that the minimum wage affects them through competition in the product market.

output in order to serve the associated demand for domestic products $Q_{gt}^D(p_t)$ at those prices. Afterwards, consumers and workers observe their preference shocks and decide. Payoffs—utilities, wages, and profits—are realized.

We work under the assumption that observed prices and wages constitute a Nash equilibrium in the pricing and wage posting game, and that marginal costs in (1.13), sales, production and employment are consistent with it. The first order condition of a firm's problem in market m at time t is:

$$q_{jmt}(p_{mt}) + \sum_{k \in \mathcal{J}_{gmt}} \Big(p_{jmt} - mc_{jmt}(\cdot)\Big) \frac{\partial q_{jmt}(p_{mt})}{\partial p_{kmt}} = 0 \tag{1.14}$$

The assumption regarding the price setting behavior of firms implies that the first order conditions are standard, and that we can estimate consumer demand using only a subset of the markets firms serve each period.[43] We use these first order conditions and previous modeling assumptions to estimate jointly consumer demand and marginal costs in the next section.

1.5 Estimation and Main Results

1.5.1 Specification Details and Instruments

We estimate consumer demand in a single retail market m per period t. The retail market is the group of large supermarkets in the Capital and its Metropolitan area, which accounts for 79% of national sales in this type of store.[44] For the purpose of this paper, defining more granular markets and including smaller regions would introduce the problem of multiple equilibrium in the counterfactual

[43] Formally, $\frac{\partial q_{km't}(p_{m't})}{\partial p_{jmt}} = 0 \;\; \forall m \in \{\mathcal{M} \setminus m\} \;\; \& \;\; \forall k \in \mathcal{J}_g$

[44] The market definition (region × time) is in line with the one used by (Miller et al. 2017), for example.

analysis (Bulow et al. 1985), with little additional gain in terms of identifying consumer heterogeneity.[45] As a result, we omit the subscript m in what follows, except where necessary.

We include unobserved preference heterogeneity for the inside good (σ_0), the price sensitivity (σ_α), and craft characteristic (σ_C). We further consider a nesting parameter (ρ) to account for correlated preferences for lager products, in line with observed consumer behavior in the removal experiment. In the baseline specification, we don not consider income heterogeneity due to lack of meaningful time variation.

To account for price endogeneity, we use a vector of demand side instruments (z^d_{jt}) that include product attributes (x^d_{jt}) and a set of excluded instruments. We decompose unobserved quality (ξ_{jt}) into product and time fixed effects, and an unobserved deviation term ($\xi_{jt} = \xi_j + \xi_t + \Delta\xi_{jt}$), in line with Hausman et al. 1994. Product fixed effects (ξ_j) absorb consumer mean valuation ($\gamma_{0,k}$) for attributes that are constant over time. Nevertheless, we use product characteristics (brand, beer style, and dummies for craft, domestic, alcohol content and bitterness) to create price instruments.[46]

The set of excluded instruments contain prices in previous periods ($p_{j,t-1}$) and exogenous input prices interacted with product characteristics. For imported products, we use their import prices, and for domestic products we use firm-level prices of inputs uncorrelated with unobserved quality (bottles and malt).[47] Excluded instruments also contain interactions between beer style and time, to account for the fact that Lager (Ale) products are consumed more intensively in

[45] In Appendix (3.1) we display the results of using a similar model using a more granular market definition and show that the main summary statistics do not differ greatly.

[46] For Alcohol by Volume (ABV) and International Bitterness Units (IBU), we classify products based on whether they are below or above the median in our sample

[47] We use CIF prices. CIF (Cost, Insurance, and Freight) is the import price reported to customs, including the cost of goods, insurance, and transport cost to the destination country.

the warm (cold) season. Joint with the traditional within-group share instrument, this instrument also contributes to identify ρ in our model.

In our supply specification, the labor cost of national products (1.13) is an explicit function of quantity and prices. In this case, we address endogeneity concerns following the literature on production function estimation (Olley et al. 1996, Ackerberg et al. 2015), and use as excluded instruments fixed inputs in current period (K_{gt}), and flexible and fixed inputs in previous periods ($K_{g,t-1}, lc_{g,t-1}$).[48] We estimate the first-stage by least-squares using these instruments and work with predicted values $\hat{lc}_{gt}$ during the estimation routine.[49] Additionally, we consider exogenous product attributes ($\mathrm{x}_{jt}^{s,ex}$) that vary according to the origin of the product. For domestic products, we include the import price of hops and bottles and the export price of malt, while for imported products we use the import price of the final product.[50] We decompose cost shocks into product, month and year fixed effects, and an unobserved deviation term ($\omega_{jt} = \omega_j + \omega_{mo} + \omega_{yr} + \Delta\omega_{jt}$).[51]

Finally, we include "second" choice aggregate diversion moments, that consider the removal of a subset of products from the choice set of certain consumers between June 2020 and November 2021.[52] The removal resembles a real experiment, where certain products are excluded from a selling point for a period of time. In this case, the experimenter is unable to observe the decision of individual consumers, but only shares before and after the removal takes place. We match the diversion ratios calculated with the model for markets in the first quarter of

[48] We build a series of observed labor cost using data on gross wages, employment, workers categories and production. Due to our inability to identify craft firms by name in our employment data, our measure of labor cost aggregates across them. See (ref Data)

[49] See first-stage.

[50] In months without transactions, we use data on previous months. The export price is FOB (Free On Board), in the custom of origin.

[51] We don't consider market fixed effects (ω_t) because, jointly with ω_j, they would absorb most of the variation contained in $\hat{lc}_{gt}$

[52] See details in Appendix (3.1).

2020, with the ones observed after June 2020 for the outside good (no-purchase option) and relevant inside products, following Conlon et al. 2024. This approach allows us to uncover rich consumer preference heterogeneity, despite estimating demand with aggregate market shares, a single market per period, and no data on individual consumers.

1.5.2 Estimation

We estimate the model with GMM, following the procedure outlined by Berry et al. 1995, and incorporate the "second" choice aggregate diversion moments according to the procedure developed by Conlon et al. 2024.

Moments are built using a two step procedure that incorporates demand (z^d_{jt}) and supply (z^s_{jt}) instruments. The goal is to search for the combination of parameters of the model that match predicted and observed market shares. First, we use a candidate solution for the parameters that drive consumer heterogeneity and compute mean utility (δ_{jt}), in order to match shares and infer the vector of unobserved quality (ξ_{jt}). Then, we compute marginal cost (mc_{jt}) and cost shocks (ω_{jt}) using the model and the first order conditions of the firm. Second, we estimate the fixed effects and build demand and supply moment conditions, by interacting structural error terms with their corresponding instruments:

$$g^d(\Theta) \equiv \mathbb{E}[\Delta\xi_{jt}(\Theta)z^d_{jt}] \tag{1.15}$$

$$g^s(\Theta) \equiv \mathbb{E}[\Delta\omega_{jt}(\Theta)z^s_{jt}] \tag{1.16}$$

The vector Θ contains the full list of parameters included in the model. "Second" choice aggregate diversion moments $g^{sc}(\Theta)$ are included at this stage and match observed "second" choices with the ones calculated using the model (1.5). When

calculating $g^{sc}(\Theta)$, we weight consumers according to the ratio between sales in treated selling points and total sales, in the period before the removal took place. If the model is correctly specified, moment conditions $\{g^d(\Theta), g^s(\Theta), g^{sc}(\Theta)\}$ are equal to zero when evaluated at the true value of the parameters Θ_0. We estimate the model by minimizing the following objective function, where W is a weighting matrix and $g(\Theta)$ is a vector that stacks the moment conditions described before:

$$\min_{\Theta} \; q(\Theta) = g(\Theta)' W g(\Theta) \tag{1.17}$$

We estimate demand and supply using PyBLP (Conlon et al. 2020). We employ the two step GMM procedure with a derivative based optimization algorithm, approximate the market share using individuals' nodes and weights according to the Gauss-Hermite product rule, and use an error tolerance of 10^{-10} for the contraction mapping.

1.5.3 Estimates

Table (1.3) reports the estimates of the parameters associated with consumer preference heterogeneity, mean price sensitivity[53] and marginal costs, and main summary statistics.

Demand

We estimate relevant unobserved heterogeneity in the price sensitivity (σ_α) and a significant distaste for prices (α_0). Notice that, while the negative sign is imposed, the magnitude of the coefficient is not. Furthermore, unobserved heterogeneity in the preference for the inside good is sizeable. This is the result of observed consumer behavior in the removal experiment. Despite excluding products that

[53]To be precise, the price sensitivity of the consumer with $\nu_{i,\alpha} = 0$

Table 1.3: Demand and Supply Estimates

Variable	Parameter	RCNL-1	RCNL-2	Logit
Demand				
Price	$\exp(\alpha_0)$	-2.20 (0.37)	-2.38 (0.18)	-3.31 (0.08)
RC - Price	σ_α	0.95 (0.18)	1.04 (0.16)	
RC - Inside Good	σ_0	2.24 (0.46)	2.71 (0.45)	
RC - Craft	σ_C	0.94 (0.16)	0.87 (0.15)	
Nest - Lager	ρ	0.66 (0.03)	0.63 (0.03)	
Supply				
Labor Cost - Dom ($\mathcal{J}^D$)	λ^{lc}	3.81 (0.22)	3.83 (0.23)	4.7 (0.42)
Hops - Dom ($\mathcal{J}^D$)	λ_H	0.009 (0.001)	0.01 (0.001)	0.015 (0.002)
Imp. Price - Imp ($\mathcal{J}^I$)	λ_I	1.36 (0.04)	1.47 (0.12)	3.0 (0.092)
Summary Statistics				
Median Own-Price Elasticity		-3.49	-3.55	-2.34
Median Market Price Elasticity		-0.11	-0.08	-0.12
Median Outside Good Diversion		12.4%	16.5%	89.8%
Fixed Effects				
Product - Demand	ξ_j	Y	Y	Y
Market - Demand	ξ_t	Y	Y	Y
Product - Supply	ω_j	Y	N	N
Month, Year - Supply	$\omega_{m(t)}$	Y	Y	Y

Note: Standard errors are displayed in parentheses. We use monthly data between April 2015 and June 2022, and only consider products that were available in at least 50 markets. N: 3,722 observations, 87 markets and 54 products. The price sensitivity of the mean consumer ($\exp(\alpha_0)$) is evaluated at $\nu_{i,\alpha} = 0$.

accounted for 92% of sales, most consumers remained in the market: observed diversion to the outside good was only 38%.[54] Consumers exhibit unobserved preference heterogeneity for craft products and correlated taste for Lager products, as indicated by the estimates of the associated random coefficient and nest parameters. In contrast, they do not show significant heterogeneity for Ale or Imported attributes. Overall, we are able to estimate rich consumer preferences for the product's characteristics that are important for this study.

The estimates of price elasticity are similar to the ones found in previous studies of the beer market, though in our framework consumers seem to be less

[54]The plain-vanilla logit model predicts a 95% diversion rate in this case.

price sensitive. The median estimate of the own-price elasticity is similar to romeo2016incorporating, Hidalgo 2023 and slightly below Miller et al. 2017, Fan et al. 2020; while the median diversion to the outside good is in line with Miller et al. 2017.[55] We estimate a market-level price elasticity of -0.11, lower in absolute terms than previous studies, but below one as in Miller et al. 2017. Regarding product categories, we find that craft products have a higher median own-price elasticity than non-craft (-4.2 vs -3.1), in line with the findings of Hidalgo 2023, Fan et al. 2020 for the US.[56]

Supply

We use our supply side estimates to recover the elasticity of output with respect to labor ($\hat{\beta}_L$). According to our specification of the production process and the assumption on firms' short-run cost-minimizing behavior, the coefficient on labor cost $\hat{\lambda}^{en}$ has a structural interpretation.[57] However, without additional information regarding whether the equilibrium we observe is consistent with the marginal costs of Equation (1.11) or (1.12), in principle, we are not able to distinguish between the two possible cases $\{\frac{1+\theta}{\beta_L};\ \frac{1}{\beta_L}\}$. To overcome this limitation, we assume that, for the number of workers we observe in the data, the current minimum wage is higher than the marginal cost of hiring and binds on firms' labor demand (1.11).[58] This assumption is also supported by the high level of the industry's minimum wage compared to those in other industries (Ramirez et al. 2022).

[55] Median own-price elasticity of -4.74 in Miller et al. 2017 preferred specification, and between -5.8 and -10.1 in Fan et al. 2020

[56] (Hidalgo 2023) estimate a median own-price elasticity of -3.4 for flagship products and -8.6 for craft ones. Alternatively, (Fan et al. 2020) estimate a median of -6.3 for non-craft and -9.7 for craft.

[57] See Equation (1.13).

[58] Observed wages would need to compensate for very low levels of amenities, and this is not a plausible explanation according to market participants. In fact, according to the owner of one of these firms, they "receive hundreds of application each time they post a wage offer at the current minimum wage level $\underline{w}$".

In this context, we interpret $\hat{\lambda}^{en}$ as containing information only about the output-labor elasticity ($\hat{\beta}_L = 0.33$). The implied coefficient, once the value added tax has been accounted for, is similar to De Loecker et al. 2022, who estimated the same production function for the US using cost data only. The coefficients on materials are significant and have the expected sign as well. The magnitude of the coefficient on hops imply a unitary cost for craft products of 25 cents per serving, which aligns with the average cost of 20 cents per serving as reported by craft brewers. The coefficient on imported products is consistent too, as more expensive products are sold to consumers at higher prices.

Marginal Costs, Markups and Diversion Ratios

The marginal cost of national products is higher for craft products compared to non-craft (industrial) products. In our preferred specification (RCNL-1), the marginal cost ratio of craft relative to non-craft products was 5.6 during 2020/21, on average[59]. From our model, we can explain part of the difference through heterogeneous labor costs and input prices. First, craft products are more expensive to produce because the labor cost element of marginal costs is 55% higher. Second, craft firms pay for hops twice what non-craft firms pay. These production components explain approximately 25% of the total gap in marginal cost between these two types of products (US$1.53 per serving). Arguably, there might be other factors associated with distribution or retail activities that explain the remaining part, and are contained in supply structural cost shocks in our model. Market prices do not reflect fully these cost differences and equilibrium markups ($\mu_u = \frac{p}{mc}$) in craft products are lower than in non-craft. During 2020/21, the average markup in craft products was ($\mu_u^C = 1.3$) and in non-craft ($\mu_u^{NC} = 3.1$),

[59]This is above the 4.8 estimated by (Hidalgo 2023) for the US in 2016. However, different from their context, in Uruguay the price differences are larger (≈95% versus 55-70% in US).

in line with the estimates Hidalgo 2023 for craft (1.2) and flagship (3.5) products.

The median non-craft to craft diversion ratio is 5%, lower than the 25-35% range than Hidalgo 2023 finds for the US. Diversion ratios inform the extent of direct competition in the product market (2023 U.S. Merger Guidelines). In our study, this is relevant because it reveals how much price changes in the products of one firm, affect residual demand and production in competing firms. Estimated diversion ratios points towards a low level of competition between craft and non-craft products. In the Uruguayan market, craft products do not seem to be an important competitive threat to traditional industrial brewers. Our findings are consistent with those of Fan et al. 2020, who estimated that consumers of craft products tend to exhibit strong loyalty to this category.

Summarizing, results show sizable marginal cost differences between craft and non-craft products, weak direct competition as implied by low non-craft to craft diversion ratios, and higher elasticity on the segment of residual demand where craft firms operate. These facts are central to understand the results of the counterfactual analysis we conduct next.

1.6 Alternative Labor Market Policies

We perform two sets of counterfactual analyses to study the welfare consequences of alternative labor market policies for consumers, workers and firms. In the first counterfactual, we search for the industry minimum wage that maximizes social welfare, and discuss the role that the positive correlation between labor and product market power plays in our results. In the second, motivated by the low degree of substitution between craft and non-craft products that we estimate, and differences in firm size and productivity, we simulate alternative product and labor market equilibrium under different industry minimum wages for craft and

non-craft firms. In what follows, we first describe briefly the basic implementation details and characterize the measure of welfare we use, and then present the results of our counterfactual analysis.[60]

1.6.1 Implementation details

Labor supply and wages

We calibrate the inverse labor supply curve that firms would face in the absence of the industry-specific minimum wage.[61] In the case of the labor supply elasticity, we follow recent evidence for Uruguay (Casacuberta et al. 2023) and assume that small craft (large non-craft) firms face a more elastic (more inelastic) labor supply curve and have less (more) labor market power.[62] For amenities ($U_{gt'}$), the calibration allows small craft firms to attract a certain number of workers when offering the National minimum wage.[63] [64] Notably, the core results of our study do not depend on the particular details of this calibration.[65]

We make three assumptions about how the level of the minimum wage translates into effective wages. First, we relate changes in the minimum wage to changes in effective wages using an elasticity of one, in line with the findings of Casacu-

[60] See Appendix (3.1) for implementation details.

[61] We describe the inverse labor supply curve in Equation (1.7).

[62] Formally, the inverse of the elasticity in craft firms is lower than in non-craft ($\theta_C \leq \theta_{NC} \iff \eta_C \geq \eta_{NC}$). Studies for other developing (Amodio et al. 2021) and developed (Yeh et al. 2022) countries find a similar relationship between firm size, markdowns and labor supply elasticity.

[63] This term contains information about firm level amenities and the general situation of the labor market. Because we assume firms to be strategically small relative to the labor market, we describe them only by referring to their amenities component.

[64] Approximately 18% of employed workers earn a wage in the range 1.0-1.5x the National minimum wage. For workers with less than 12 years of formal education, this is a relevant benchmark.

[65] In the baseline specification, we use ($\theta_C = 0.1$) for craft and ($\theta_{NC} = 0.25$) for non-craft firms, consistent with wages that would be marked down 9% and 20%, respectively, absent the industry minimum wage. In the case of amenities, we use ($U_C = 440$) and ($U_{NC} = 276$), consistent with the largest non-craft firm offering 20% better amenities than craft. We show results using alternative parameterizations in Appendix (3.1).

berta et al. 2023 and Hermo 2023 for Uruguay and Argentina, respectively.[66] Second, based on the descriptive evidence, for craft firms we assume that when the industry minimum wage binds, it coincides with effective wages. Third, for non-craft firms we assume that the ratio between the minimum and the average wage that we observe in the data remains constant.[67] Despite the existence of four occupation-specific minimum wages related to manufacturing activities, the ratio between categories remained constant in the period we study. The minimum indexes the entire wage structure in these firms.

Social Welfare

We calculate social welfare as the unweighted sum of consumer surplus, workers' rents and producer surplus (firms' profits).[68] In our framework, all measures are a function of the minimum wage ($\underline{w}$) through equilibrium prices ($\mathbf{p}_t$) at time t. We refer to the minimum wage that maximizes this measure as the optimal minimum wage $\underline{w}^*$.

$$\mathrm{SW}_t(\underline{w}) \quad = \quad \mathrm{CS}_t(\mathrm{p}_t(\underline{w})) + \mathrm{WR}_t(\mathrm{p}_t(\underline{w})) + \mathrm{PS}_t(\mathrm{p}_t(\underline{w})) \tag{1.18}$$

We calculate consumer surplus using the log-sum formula (McFadden 1981), integrating over the estimated distribution of consumers' marginal utility of income (α_i). We assume that consumers in the beer market do not experience any income effect from price variations, so consumer surplus and the compensating variation

[66] In Uruguay, internal studies conducted in the Ministry of Labor and Social Security find the same elasticity.

[67] Even though the minimum wage binds for new (entry) workers, in practice the average wage is above this level. This is the consequence of the existence of more than one production occupation, and of other benefits or payments related, for example, to tenure within firms.

[68] We abstract from fixed cost considerations.

coincide. The measure we use is:[69]

$$CS_t = \int \frac{1}{\alpha_{it}} \log \left[1 + \sum_{j \in \mathcal{J}_t} \exp \left(\alpha_{it} p_{jt} + \gamma_{it} \mathrm{x}_{jt} + \xi_{jt} \right) \right] dF(i) \tag{1.19}$$

We define workers' rents in line with Kroft et al. 2021, who follow Robinson 1933 and Rosen 1986: the rent of a given worker is the excess payment that she receives over what would be required to change her job choice. In practice, all workers have the same weight and total rents in a given firm are the difference between the equilibrium wage bill of the firm and the area below the inverse labor supply curve (1.7). In each period t, we measure rents in every firm g and then aggregate at the industry level according to:

$$WR_t = \sum_g \left(w_{gt} L^*_{gt} - \int_0^{L^*_{gt}} w^s_{gt}(L) dL \right) = \sum_g \left(w_{gt} L^*_{gt} - \frac{U_{gt} (L^*_{gt})^{1+\theta}}{1+\theta} \right) \tag{1.20}$$

Finally, we assume that social and private welfare coincide (Conlon et al. 2020). As a result, social welfare can be expressed equivalently as the total area below the demand curve, net of total variable (non-labor) costs, and the value of workers' next best alternatives.[70]

Simulation procedure

We simulate labor and product market equilibrium for every minimum wage ($\underline{w}$) set by the policymaker. From our description of the problem of the firm in Section (1.4.3), in equilibrium the prices that firms set in the product market and the wages they post in the labor market are consistent with each other. Consistency

[69] We normalize the constant of integration to 0. Its value is not relevant when calculating welfare changes.

[70] We assume that firms' profits include Government revenues from value added and excise taxes. Notice that while we can infer revenues associated with the value-added tax, we cannot do the same for the excise tax. We assume that the excise tax is part of the marginal cost of a product through $\hat{\omega}_{jt}$, and that it remains constant in our counterfactual analysis.

implies that there is no profitable deviation in any market: the vector of market prices is an equilibrium in the product market given a vector of marginal costs, and the vector of wages associated with those marginal costs is an equilibrium in the labor market. At the vector of posted wages, firms can hire the workers they need to produce and satisfy the demand they receive from consumers at the vector of equilibrium prices.

We use the estimated model and the calibrated labor supply curve to search for the equilibrium vector of prices (p_t) that would arise in each case, and then recover wages, employment, production and sales.[71] We operationalize the search using the concept of consistency defined above. We employ a fixed-point iteration over product and labor markets' equilibrium with a tolerance of 10^{-9} for the maximum difference between shares in successive rounds. We simulate equilibrium for all markets in the period 2019-2021, using a grid of gross minimum wages in the range [525; 1,500] in intervals of US$ 25.[72]

Finally, to isolate inefficiencies related to labor and product market power from input misallocation, in our counterfactual, we allow firms to adjust their capital-labor ratio before engaging in short-run price competition. We do this because changes in the minimum wage affect relative input prices and firms' optimal capital-labor ratios.[73]

1.6.2 Optimal Uniform Minimum Wage

The first labor market policy we study is allowing the policymaker to establish the minimum wage that maximizes social welfare. We find that the gross minimum

[71] Conservatively, we assume that remaining retail markets m' are not affected, so we focus on the market where we estimated the demand model only.

[72] Equivalent to 13 cents in the gross hourly wage of a full time worker.

[73] We explain the procedure in detail in Appendix (3.1) and discuss the consequences of this assumption in the Counterfactual Section (1.6.2).

wage that maximizes social welfare is US$ 1,000 per month, compatible with an industry average wage 18.0% higher than the average under the no minimum wage scenario, and 93.7% higher than the National minimum wage.[74] We display the change in social welfare relative to baseline at each minimum wage level in Figure (1.3) and summarize the welfare gains that the optimal level ($\underline{w}^*$) achieves for consumers, workers and firms in Table (1.4). We use as a baseline an equilibrium where the industry minimum wage coincides with the National minimum wage.[75]

The optimal minimum wage results in both winners and losers. Compared to baseline, workers and the largest firm are better off under $\underline{w}^*$, while consumers and craft brewers are worse off. On one hand, the minimum wage is high for craft brewers, increasing marginal costs and prices, and reducing employment, production and sales. On the other hand, the optimal minimum wage mitigates labor market power in the largest firm and improves efficiency. This boosts effective wages, employment and profits at this firm, as well as in the industry due to its dominant size. For consumers, the welfare loss associated with the first effect outweighs the gains from the second, leading to a reduction in consumer surplus (US$ 0.16 Million).

Consumers receive a modest share of the efficiency gains that mitigating labor market power in the dominant firm delivers. Although the optimal minimum wage reduces marginal costs 5.5%, equilibrium prices of domestic products fall by a smaller amount (0.6%). Figure (3.4) shows changes in the equilibrium marginal costs and prices of domestic products by firm relative to the baseline. Furthermore, once the equilibrium response in the prices of imported products is accounted for, the aggregate price reduction is even lower (0.3%). The degree of product market power of this firm is large, so the savings in marginal costs of domestic

[74] Relative to the current equilibrium, the observed minimum wage is 72.2% above the optimal.
[75] Since the National minimum wage is not binding in any equilibrium, the baseline scenario is equivalent to a situation without an effective minimum wage.

products are passed through to prices at a rate of 12%, on average. Operating over the inelastic part of the demand curve weakens the incentives of this firm to reduce prices and increase sales after a cost's reduction. The minimum wage that maximizes consumer surplus (US$ 550) is lower than the optimal minimum wage.[76] The main driver of the difference in consumer surplus between the two minimum wage levels is the reduction in the consumption of craft products. For minimum wages above US$ 550, and except in a neighborhood of $\underline{w}^*$, consumer surplus decreases monotonically. Consumers would be better off if the policymaker where to accept certain degree of labor market power in this industry.[77]

Table 1.4: Simulation Results - Uniform Minimum Wage

Optimal Minimum Wage $\underline{w}^*$ by Objective - (Variation vs Baseline No-MW Equilibrium)				
Welfare (US$ Million)	Δ Cons. Surp.	Δ W. Rents	Δ Profits	Δ Soc. Welfare
Social Welfare - MW: US$ 1,000	-0.16	1.52	0.49	1.85
Consumer Surplus - MW: US$ 550	0.00	0.03	-0.00	0.03
Equilibrium (Δ %)	$\Delta\bar{p}$	ΔQ_t	$\Delta\bar{w}$	ΔL_t
Social Welfare - MW: US$ 1,000	-0.3%	0.1%	18.0%	4.0%
Consumer Surplus - MW: US$ 550	-0.0%	0.0%	-0.0%	0.1%

Note: Welfare effects for markets in 2019/2021.

In craft firms, the optimal minimum wage binds on the labor demand and reduces sales (-14.2%), due to higher marginal costs (5.6%) and associated equilibrium prices (5.1%). Craft firms are more labor-intensive, less productive and face a more elastic residual demand curve, so higher labor costs are passed through one-to-one to prices, and result in more than proportional reductions in sales. According to our demand estimates[78] consumers have a taste for craft products, so the reduction in consumption entails a significant welfare loss. Therefore, when

[76] This level is 7% above the National MW.

[77] Importantly, our assumption on the specific degree of labor market power (θ_{NC}) changes the size of the gap between the two minimum wage levels, but not the underlying low pass-through rate that drives it.

[78] See Table (1.3)

the price reduction in non-craft firms and the reduction of craft consumption are combined, consumers are worse off under the optimal minimum wage than under the no minimum wage equilibrium. The minimum wage that mitigates labor market power and improves efficiency, reduces the welfare of consumers and craft firms' profits.

In contrast, the optimal minimum wage leads to improved outcomes for both workers and the largest firm. In the case of workers, despite lower employment in craft firms, average wages (18.0%), aggregate employment (4.0%) and total workers' rents in the industry are higher. The optimal minimum wage increases efficiency in the sense of Robinson 1933 and raises employment and effective wages simultaneously. Industry profits raise because the increase in the profits of the largest firm dominates the reduction in craft firms (-10.9%). For the former, the low equilibrium pass-through rate implies that the price-cost margin increases. Still, the increase in profits at the largest firm is modest (2.6%). The reason is that labor market power represents a minor distortion relative to product market power in this industry, so addressing this inefficiency does not entail large aggregate welfare gains, except for workers. Then, even under the optimal minimum wage, efficiency gains do not translate into sizeable reductions in products' prices. In fact, equilibrium markups and product market concentration raise (marginally) under the optimal minimum wage.

Jobs lost at craft firms under the optimal minimum wage are compensated by an increase in equilibrium employment at the largest firm. The reason is that mitigating monopsony reduces the marginal cost of labor, and incentivizes firms to hire additional workers. Lower hiring costs affect equilibrium employment through two channels in our framework. First, lower labor costs reduce marginal costs and prices, expanding sales, production and employment. This effect, which

materializes in an increase in the manufacturing of domestic non-craft products, explains approximately 18% of the increase. The second channel is a consequence of the assumption that there is no input misallocation. We assumed that firms adjust their capital-labor ratio when input prices change before engaging in price competition. This implies that a lower (higher) equilibrium wage makes labor relatively less (more) expensive and decreases (increases) the equilibrium capital-labor ratio. The remaining share of the employment gain in the largest firm is due to the effect of the minimum wage on input combination, an not to lower products' prices and higher sales.

Figure 1.3: Welfare Change - vs Baseline at each Min. Wage - US$ Million

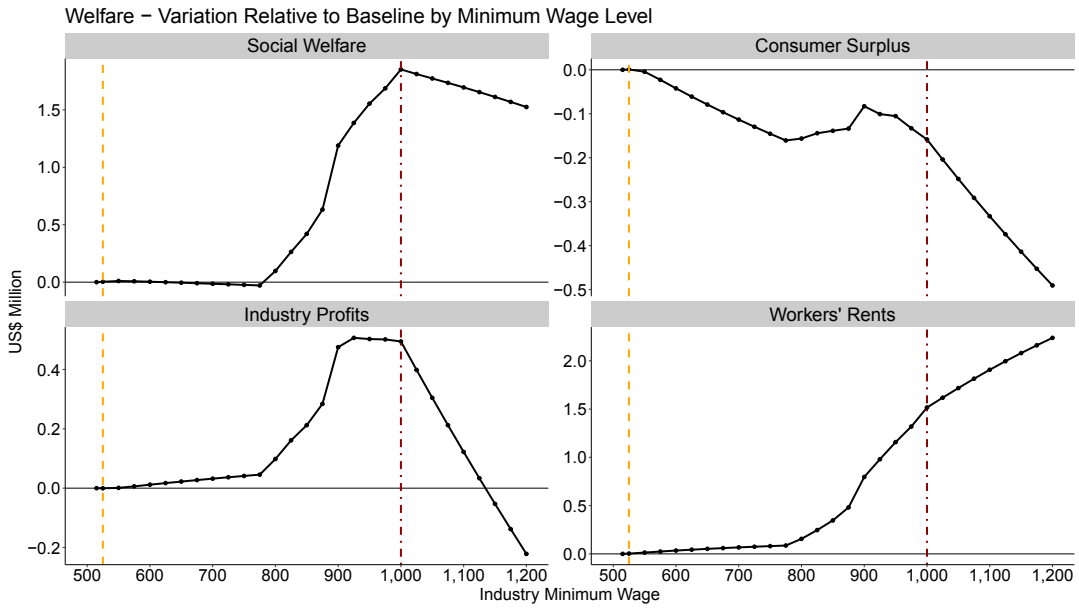

The reallocation of workers from craft (small/less productive) to non-craft (large/more productive) firms driven by consumer substitution plays a minor role in this industry. Recent work shows that following a minimum wage increase, workers reallocate from small/less productive firms to larger/higher productivity firms (Harasztosi et al. 2019, Dustmann et al. 2022)[79]. While possible in our

[79]Dodini et al. 2023 find similar results for plausibly exogenous labor cost increases.

model, sales lost in craft firms are not compensated by meaningful changes of opposite sign in non-craft firms, due to the low estimated craft to non-craft diversion ratio. Approximately 9 out 10 consumers that stop purchasing craft beer following a price increase choose to leave the market, instead of purchasing non-craft options.[80] The shift in the residual demand of non-craft products when the prices of craft ones increase is not large. This is consistent with consumers' taste for craft and rationalizes the reduction in consumer surplus we described before under the optimal minimum wage.

Our results imply that technology and consumer preferences are both potentially relevant in understanding the reallocation process that has been documented after a minimum wage increase. Crucially, a lower elasticity of substitution between capital and labor that the one implicit in Cobb-Douglas technology may make jobs lost at craft firms irrecoverable. This complements the work of Berger et al. 2022b for the US. In a model that assumes perfect competition in the product market, the authors find that the largest efficiency gain of using the minimum wage to curb labor market power would come from workers reallocating to more productive firms. Our work shows how product market power and consumer preferences over imperfect substitutes may ameliorate the pace of the reallocation process, and ultimately, reduce the efficiency gain of such policy.

Alternative Correlation between Product and Labor Market Power

The positive correlation between labor and product market power is the main driver behind the conflict between the interests of consumers and workers in our framework. In this industry, the minimum wage that addresses labor market power leads to limited price savings for consumers, while also raising the price of

[80] For minimum wages in the range US$ 550 - US$775, the market shrinks at the pace of the sale reduction in craft brewers.

craft products. To study how our results depend on this correlation, we simulate a situation where product market power diminishes while labor market power remains constant. This would resemble firms producing in separate plants with local labor market power, while competing in the product market at a higher (national) level.

The simulation consists in changing the ownership of certain products currently in the portfolio of the largest firm, and establishing an oligopoly with three large national producers instead, but retaining the degree of labor market power and productivity that we estimate for the largest firm.[81] A key aspect that explains the degree of product market power of the largest firm is that it offers products with high diversion ratios with each other. Therefore, by changing the ownership of some of these products, we make competition in the product market more intense.

Figure 1.4: Divestitures - Welfare Change vs No-MW Equilibrium

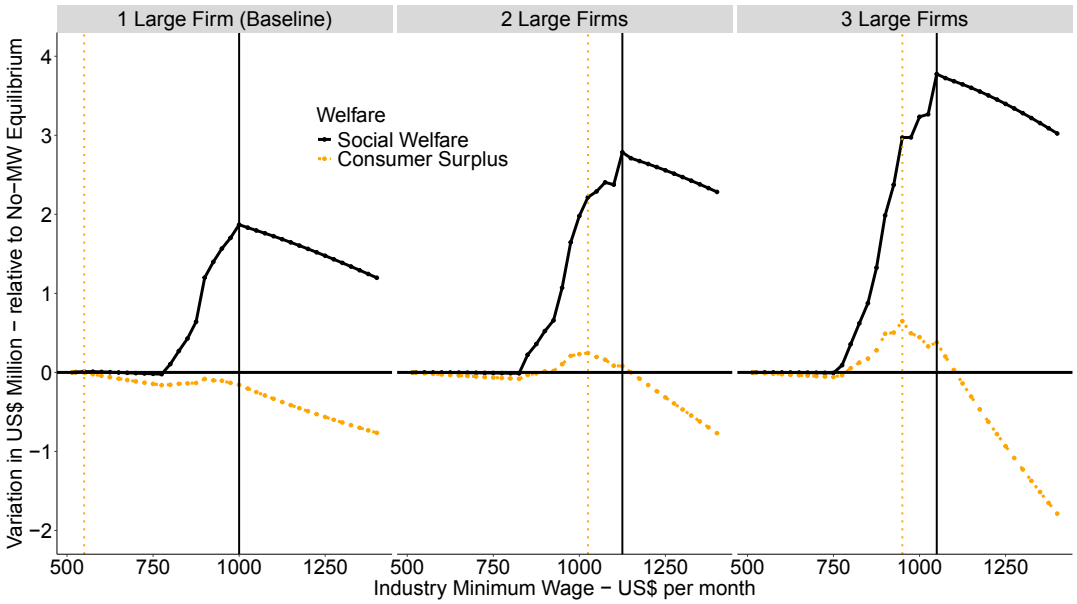

[81] In practice, we simulate a product market with 2, 3, 4, and 5 large firms. Part of the exercise implies restoring in steps the brand ownership observed before 1998 when the first horizontal merger took place and the market went from an oligopoly with 3 firms to a duopoly with just 2. The second merger occurred in 2003 and established a monopoly in production and sales. We describe here the oligopoly with 3 national producers and one large importer. The complete exercise is described in Appendix (3.1).

We find that the optimal minimum wage ($\underline{w}^*_{sim}$) is US$ 975 and that now the interests of consumers and workers are better aligned.[82] The minimum wage that maximizes consumer surplus is US$ 850 in this case. The gap between the minimum wages that maximize social welfare and consumer surplus decreased from US$ 450 to US$ 75 relative to the estimated for the observed product market structure. Consumers now benefit from the efficiency gains that mitigating labor market power delivers due to increased product market competition. When firms experience the cost efficiency that the optimal minimum wage entails, they have the incentive to pass it through to prices to increase sales and market share. The equilibrium pass-through rate of national non-craft products in this simulated oligopoly is 56%, on average. Fixing labor market inefficiency produces a larger positive welfare effect for consumers because the markups of national non-craft products are lower (1.9-2.0) relative to the estimated ones in observed equilibrium (3.1). Therefore, removing the markdown distortion is relatively more important. Notice that the optimal minimum wage in this simulation is still high for craft brewers, but different from before, the pro-efficiency effects that the minimum wage delivers lower the prices of national non-craft products to a larger extent. In a more competitive product market, the minimum wage that mitigates labor market power in large firms simultaneously raises the welfare of consumers and workers.

[82]The optimal minimum wage is lower for two reasons. First, the labor demand of the new firms shifts inward as a result of the change in product ownership. The beer market expands due to more competition, but firms individually produce less. Second, since divested firms have the labor productivity and face the labor supply curve of the largest firm, the marginal costs' curves coincide. Combined, firms produce on a smaller scale and the pro-efficiency effects of the minimum wage materialize at lower wage levels.

1.6.3 Minimum Wages for Craft and Non-Craft Firms

Considering that the optimal uniform minimum wage leads to lower consumer surplus and to both lower employment of craft firms, as well as higher prices for their products, we next consider the impact of distinct minimum wages for craft and non-craft firms.[83] Notably, this labor market policy was approved by the Ministry of Labor and Social Security during the course of this project and will take effect in 2025. In this context, our counterfactual provides a first measure of the potential welfare gains that would be achieved under an optimal implementation for the group of firms currently in the industry.

From our previous results, we know that craft and non-craft products are considered poor substitutes from consumers' point of view, and that the labor productivity in craft firms is lower than the productivity in non-craft firms. This results in lower consumer surplus when the minimum wage is established to mitigate labor market power. In this context, our solution consists in allowing the policymaker to establish the optimal combination of minimum wages for craft and non-craft firms that maximizes social welfare. In Table (1.5) we display the welfare variation relative to baseline and compare it with the two optimal levels we discussed previously.

[83] Furthermore, introducing a structural remedy in the product market to deal with a merger that occurred more than twenty years ago seems unlikely.

Table 1.5: Simulation Results - Optimal Minimum Wage

Optimal Minimum Wage $\underline{w}^*$ by Objective - (Variation vs Baseline No-MW Equilibrium)				
Uniform Minimum Wage				
Welfare (US$ Million)	Δ Cons. Surp.	Δ W. Rents	Δ Profits	Δ Soc. Welfare
Social Welfare - MW: 1,000	-0.16	1.53	0.49	1.87
Consumer Surplus - MW: 550	0.00	0.03	-0.00	0.03
Equilibrium (Δ %)	$\Delta\bar{p}$	ΔQ_t	$\Delta\bar{w}$	ΔL_t
Social Welfare - MW: US$ 1,000	-0.3%	0.1%	18.0%	4.0%
Consumer Surplus - MW: US$ 550	-0.0%	0.0%	-0.0%	0.1%
Minimum Wages for Craft and Non-Craft Firms				
Welfare (US$ Million)	Δ Cons. Surp.	Δ W. Rents	Δ Profits	Δ Soc. Welfare
Social Welfare - MWs: {550, 1,000}	0.12	1.42	0.5	2.04
Equilibrium (Δ %)	$\Delta\bar{p}$	ΔQ_t	$\Delta\bar{w}$	ΔL_t
Social Welfare - MWs: {550, 1,000}	-0.5%	0.3%	12.3%	9.6%

Note: Welfare effects for markets 2019–2021.

We find that the combination of minimum wages that maximizes social welfare is US$ 550 and US$ 1,000, for craft and non-craft firms, respectively. This policy increases consumer surplus by US$ 0.12 M and employment by 9.6% relative to baseline. Compared to the optimal uniform minimum wage equilibrium, consumer surplus improves even further (+US$ 0.28 M). Consumer benefit from lower prices in craft products (-4.8%) and marginal additional price reductions in domestic non-craft (-0.2%) and imported products (-0.1%) relative to the uniform optimal case. The low non-craft to craft diversion ratio prevents lower craft prices from imposing significant competitive pressure on other firms in the industry.

Worker welfare increases less than under the uniform minimum wage. The employment gain in craft firms is insufficient to compensate for the more moder-

ate wage increase. Overall, social welfare is higher with distinct minimum wage categories primarily because it raises consumer surplus, employment and profits at craft firms, with negligible negative impacts on profits and employment at non-craft firms. This policy framework provides policymakers with an additional degree of freedom in setting minimum wages. Therefore, when used to account for structural differences in labor productivity and consumer preferences between craft and non-craft products, it can enhance social welfare.

1.7 Conclusion

Policymakers and competition authorities are concerned about the negative effects that labor market power has on workers and the efficient functioning of the labor market. In this paper, we investigate the trade-offs they face when using the minimum wage as a tool to curb firms' power. Using a model of simultaneous product and labor markets' equilibrium, we find that using the minimum wage to mitigate labor market power reduces consumer welfare. This occurs due to a positive correlation between labor and product market power, and the existence of firms with varying levels of labor productivity. We highlight the role that product market power plays in the distribution of the efficiency gains that the minimum wage generates in the labor market. Finally, in contexts where sector-specific minimum wages apply, grouping firms with different levels of labor productivity that produce imperfect substitutes in a common sector may lead to unnecessary welfare losses.

We make a series of relevant assumptions that influence our results. First, by studying only the intensive margin (no entry-exit) without involuntary unemployment, we may underestimate the negative welfare effects of the optimal minimum wage on consumers and on workers due to possible job losses. Incorporating these

dimensions could be a direction for future research. Second, the conduct assumptions in the product market might underestimate the share that consumers receive from mitigating labor market power in the largest firm, or overestimate the extent to which the prices of craft products vary when sold through retailers. Third, the utilization of a production technology with constant elasticity of substitution between capital and labor may affect the employment variation we estimate and ultimately workers' rents. Considering a more flexible technology could improve the ability of this study to evaluate similar situations in other contexts.

Chapter 2

"Regulation by Public Options: Evidence from Pension Funds" with Pablo Blanchard and Sebastian Fleitas

2.1 Introduction

Governments often rely on State-Owned Enterprises (SOEs) to enhance competition and regulate markets. In the market for Pension Fund Administrators (PFAs) where firms charge workers a management fee in exchange for investing their contributions, weak competition has led to high fees and reduced workers' net savings (OECD, 2018). As several countries use this institutional design[1], concerns over low retirement benefits have sparked policy debates on whether alternative policies can mitigate market power and improve workers' outcomes. This paper examines the welfare effects of using a public option —a SOE that competes with private firms in a market— as a regulatory or competition-enhancing instrument in the market for PFAs.

The argument in favor of introducing public options[2] is that, because they

[1] In Latin America, these systems were introduced decades ago, and most of them are still in place: Chile (1980), Peru (1993), Colombia (1994), Uruguay (1996), Bolivia (1997), Mexico (1997), El Salvador (1998), Costa Rica (2001), Dominican Republic (2003), Nicaragua (2004), Ecuador (2004).

[2] Except where specifically noted, we will use the terms state-owned enterprises and public options interchangeably throughout the paper.

are not pure profit-maximizers, they can compete more aggressively with private firms by charging lower fees, thereby increasing workers' savings. However, while the participation of public options can contribute to solving market failures, the existent literature shows how the equilibrium welfare effects of this policy are a-priori uncertain Kang 2022. Market segmentation, price-increasing competition, and inefficient provision (Hastings, Hortaçsu, et al. 2017, Chen et al. 2008, Duggan et al. 2006) can negatively affect the welfare of market participants instead of increasing it (Jiménez-Hernández et al. 2021, Atal et al. 2021, Fonseca et al. 2022). Therefore, whether the public option outperforms alternative policies in raising workers' savings or not is an empirical question.

In this paper, we study the welfare consequences of the participation of a state-owned enterprise (SOE) in the Uruguayan market for Pension Fund Administrators (PFAs), where three private firms and a public option compete. Using rich administrative data, we estimate a structural model of demand and supply in which forward-looking PFAs compete to enroll workers and manage their savings. In Uruguay, PFAs charge a single management fee to administer and invest workers' social security contributions until retirement, as part of the defined-contribution retirement subsystem[3]. We analyze three distinct equilibria, driven by two policy shifts: a change in the SOE's shareholder preferences in 2005 and the introduction of a cap on management fees in 2018. We use the estimated model to conduct counterfactual policy evaluations.

We develop and estimate a dynamic model of demand and supply in the market for pension fund administrators. On the demand side, we specify a two-stage decision model for inattentive individuals, consistent with the observed inertia in

[3]PFAs participate in the defined-contribution subsystem while workers are active in the labor market. Once they retire, insurance companies pay retirement benefits in the form of annuities, based on the retiree's age and accumulated savings. The defined-contribution subsystem is part of the broader retirement system, which also includes a defined-benefit (pay-as-you-go) subsystem.

workers' initial choices. In the first stage, workers receive an awareness shock (Ho et al. 2017) that, if large enough, leads them to re-optimize and reconsider their enrollment decision. In the second stage, they choose where to enroll using a discrete choice model in the tradition of (Berry 1994). We also assume that workers are myopic (Luco 2019), expecting observed fees and average returns to remain constant until retirement.

On the supply side, we use a dynamic model of forward-looking firms that compete for enrollees by setting fees and mean portfolio returns, with no possibility of fee or return discrimination between new and existing cohorts. Private firms maximize the present discounted value of economic profits, while the public option considers both profits and workers' savings in its objective function. Since workers already enrolled exhibit inertia and in each period new cohorts are small relative to the existing stock of enrollees, firms face an investing-harvesting trade-off (Beggs et al. 1992). Furthermore, given the observed stability in management fees and return differentials in the periods we study, we assume that equilibrium outcomes are consistent with a stationary, "no-sales" Nash-Bertrand equilibrium (Farrell et al. 2007).

We use detailed administrative data provided by the Uruguayan social security administration (Banco de Previsión Social). The data consists of a panel of a representative sample of workers who made enrollment decisions between 1996 and 2020. We observe monthly data on gross wages, basic employer characteristics, spells in and out of the formal labor market, PFA enrollment decisions, and the enrollment mechanism, along with basic demographics (sex and date of birth). We complement these data with publicly available market-level information on market shares, management fees, investment returns, switchers, sales force agents, and PFAs' financial statements.

We use market-level data to estimate a probit model for the first stage of the decision model, where the probability of being aware depends on the gap between the fees and returns of private firms and those of the SOE[4]. For the second stage, we estimate a conditional logit model that allows workers' heterogeneity in fee and return sensitivities based on wage quartiles and whether they have access to an outside option[5]. To estimate preferences, we exploit the fact that fees and returns are common for all enrollees within a firm, while management costs and the potential accumulated savings from an enrollment decision vary across individuals, due to differences in gross wages (Hastings, Hortaçsu, et al. 2017). Our estimates show that higher-wage workers are more sensitive to management fees than to returns, and that on average make enrollment decisions that deliver higher savings relative to those of low-wage workers. These findings are consistent with evidence that financial literacy is positively correlated with income (Lusardi 2008).

We use estimated preferences and the assumption that firms set fees and mean returns to maximize the present discounted value of their objective function to recover marginal enrollment costs and investment costs. In contrast to the previous literature that studied this market assuming zero marginal costs Hastings, Hortaçsu, et al. 2017; Luco 2019; Illanes 2016, we estimate positive enrollment marginal costs, aligned with the variable payment per enrollee that sales force agents receive from PFAs in practice. Additionally, we leverage the existence of a minimum rate of return regulation and use estimated preferences and the risk of capitalization that firms face to estimate the variable cost of obtaining investment returns for their enrollees. Finally, to recover the non-profit motives of the public option, we assume its enrollment marginal cost is similar to that of private firms, since marginal cost and the conduct parameter cannot be separately identified in

[4]Consistent with what we observe in the data. We modify this assumption in our counterfactual simulations.

[5]This depends on their wage level.

this case.

We conduct counterfactual policy evaluations to quantify the equilibrium welfare effects of the public option and assess the role of alternative regulatory instruments. We consider two main counterfactuals. First, we simulate the privatization of the SOE by converting it into a private firm with standard profit motives and similar cost structures. We also explore how the effects of privatization change when it is complemented with a demand-side policy that reduces workers' inertia. Second, we examine the effects of strengthening the SOE's non-profit motives and compare the results with the effects of introducing a cap on management fees—a regulatory tool implemented in Uruguay in 2020 and commonly used in other pension systems.

The privatization counterfactual shows that the SOE contributes to reducing fees in the market and increasing workers' savings. Eliminating it would result in a 44% increase in average fees and a 3% decrease in mean returns, leading to a 6.7% decline in expected retirement savings for a representative worker. These effects are more pronounced for former SOE enrollees, who experience a 150% increase in fees and a drop of over 10% in savings. We also simulate a demand-side policy that complements privatization, in which workers become more responsive to fee differences. Although increased sensitivity improves overall savings and reduces fees, the degree of responsiveness required to fully offset the effects of privatization is unrealistically high. Moreover, this combination still creates winners and losers, with former SOE enrollees being worse off.

In the second counterfactual, we explore the consequences of increasing the non-profit motives of the SOE and study the distributive profile of the policy relative to a cap on management fees. We find that raising non-profit motives increases market segmentation, leaving low-wage workers enrolled in private firms

worse off due to higher equilibrium fees cause by the incentive to harvest on lower elasticity stayers. When compared with the cap on management fees introduced in 2020, the cap outperforms the non-profit motives in terms of aggregate savings and equity: it increases savings for all workers, especially the lowest earners, and narrows the savings gap between SOE and private firm enrollees from 91.3% to 98%.

In summary, an oligopoly with a public option competing alongside three private firms delivers welfare gains by lowering fees and raising returns relative to an oligopoly with four private firms, particularly for the SOE's own enrollees. Demand-side policies can enhance competition but are unlikely to fully compensate for the privatization of the SOE. Furthermore, strengthening the SOE's non-profit motives can raise savings for its own enrollees, but also exacerbate market segmentation and inequality. By contrast, a cap on management fees improves savings across the board—especially for low-income workers—and promotes convergence in outcomes between private and public enrollees.

In this paper, we contribute to two strands of the literature. First, we contribute to the empirical literature that analyzes the welfare effects of state-owned enterprises Fonseca et al. 2022; Jiménez-Hernández et al. 2021; Atal et al. 2021; Handbury et al. 2021; Curto et al. 2019; Busso et al. 2019; Cunha et al. 2019. In this context, we contribute by evaluating the equilibrium effects of a public option on welfare in a market with forward-looking, single-product firms that possess market power, in an environment where workers place a high valuation on the SOE. Our results show that the presence of the SOE benefits not only workers enrolled in it but also those enrolled in private PFAs, due to increased competition, lower fees, and higher savings, relative to a privatization scenario in which four private firms compete. However, using its non-profit motives more aggres-

sively also raises market segmentation and can leave low-wage workers enrolled in private firms worse off.

Second, we contribute to the literature that studies the regulation of the market of pension fund administrators in individual capitalization retirement systems Hastings and Tejeda-Ashton 2008, Hastings, Hortaçsu, et al. 2017, Illanes 2016, Luco 2019. In this case, different from previous papers, we observe a public option in action, instead of estimating its effects as a counterfactual (Hastings, Hortaçsu, et al. 2017). Furthermore, the changes in the regulatory environment and in shareholders' preferences allow us to understand how the welfare effects of the public option change when the institutional configuration also changes. Furthermore, we leverage the financial regulation to endogeneize the returns that firms obtain when investing workers' savings. This makes our analysis more comprehensive than previous studies on the topic by dealing with the two main outcomes that matter for workers' retirement benefits: fees and returns.

2.2 Institutional Environment

2.2.1 Retirement System Overview

In 1996, Uruguay reformed its retirement system by introducing a defined-contribution subsystem based on individual capitalization to complement the existing defined-benefit (pay-as-you-go) system. While the latter is administered by the social security administration, Banco de Previsión Social (BPS), and is mandatory for all workers[6], the former is mandatory only for high-wage workers. However, in practice most active workers participate in the defined-contribution subsystem and have an individual account in a Pension Fund Administrator (PFA). This subsys-

[6]With the exception of small groups covered by special retirement subsystems.

tem is partitioned into two. Before retirement, workers' contributions to individual savings accounts are managed and invested by PFAs, financial institutions regulated by the Central Bank (BCU). After retirement, an insurance company pays an annuity based on workers age at retirement and accumulated savings[7].

2.2.2 Enrollment and Contributions

Enrollment in a pension fund administrator is optional for workers earning less than USD 1,535[8] and mandatory for those earning more. If they are not already enrolled when they first cross this threshold, they have two months to choose a PFA. If they do not, the social security administration assigns one by default[9]. Enrollment outside this mechanism requires contact with sales force agents employed by PFAs

Workers contribute a fixed 15% of their gross wages to social security (with no contributions required on earnings above USD 4,605)[10]. The distribution of these contributions between the defined-benefit subsystem administered by the Social Security Administration (SSA-DB) and the defined-contribution subsystem administered by pension fund administrators (PFA-DC) depends on a worker's gross monthly wage and whether they choose to apply Article 8 of Law 16.713. For enrollees earning less than USD 1,535, if they do not opt for Article 8, 100% of their contribution goes to the defined-benefit subsystem[11]. If they do opt for

[7]Pension fund administrators do not participate in this segment.

[8]Expressed in real US Dollars of April 2021. Thresholds are adjusted yearly according to a Nominal Average Wage Index.

[9]Until 2014, the assignment was made by lottery in proportion to firms' market share. Since then, accounts are assigned to the two firms with the lowest management fees, unless the gap between them exceeds 20%. In that case, default affiliations are made only to the firm with the lowest fee. In practice, this enrollment mechanism accounts for 10–12% of new enrollments each year, and since 2014 the SOE has been the sole beneficiary.

[10]Employers' contributions (also a share of workers' gross wages) go exclusively to the defined-benefit subsystem.

[11]See the right bar of Figure 3.9.

Article 8, the contribution is split equally between SSA-DB and PFA-DC for income up to this threshold.

If a worker earns more than USD 1,535, their contribution is automatically split between SSA-DB and PFA-DC, regardless of whether they choose to apply Article 8. However, the specific distribution across subsystems depends on this choice. With Article 8, contributions from USD 1,535 to USD 2,303 go to SSA-DB, and those from USD 2,303 to USD 4,605 go to PFA-DC. Without it, 100% of contributions above USD 1,535 are allocated to the defined-contribution subsystem. In practice, more than 90% of eligible workers have opted for Article 8.

PFAs receive workers' monthly gross contributions, from which they pay a disability and survivor insurance premium on behalf of the worker[12] and deduct a management fee. The remaining amount is credited to the worker's individual account in the PFA and invested by this institution in legally authorized assets.

2.2.3 PFAs' Market Structure and Regulation

There are four active firms in the market for pension fund administrators: three private firms and one public option. The SOE operates under the same rules as private firms, but its shareholders are other state-owned public institutions. It has been active since the inception of the system in 1996, with ownership divided among BROU (a state-owned commercial bank, 51%), BPS (the Social Security Administration, 37%), and BSE (a state-owned insurance company, 12%)[13]. The private segment of the market originally included five firms, but reached its current

[12] Also as a share of the gross wage. The premium is set by a monopolist state-owned insurance company as a share of workers' gross contributions. The differences between premiums across PFAs are negligible.

[13] BROU: Banco de la República Oriental del Uruguay, BSE: Banco de Seguros del Estado and BPS: Banco de Previsión Social.

structure after four firms merged into two in 2001.

Firms charge management fees for their services. Until 2008, they were allowed to charge both a fixed fee (a flat monetary amount) and a variable fee (a percentage of the gross contribution). Since then, only a single variable fee is permitted. PFAs are not allowed to price discriminate, so all enrollees are charged the same fee. In 2018, Parliament introduced a cap on fees, setting it at 1.5 times the lowest fee in the market[14]. This cap was fully implemented in 2020, following a two-year transition period.

Firms invest workers' savings in legally authorized assets and manage investment portfolios based on the age of each enrollee. Unlike in other individual capitalization systems, workers in Uruguay cannot choose how their savings are allocated across funds. Until 2014, each PFA operated a single investment fund. A second fund was introduced that year, and since then, the allocation between the two has been determined entirely by the enrollee's age.

Firms are required to provide a minimum level of investment return to their enrollees. By law, PFAs must cover any shortfall using their own capital when their rate of return falls below a specified threshold. This threshold, which varies with the overall level of returns, is defined as the lower of: 1) 2%, or 2) the system's average return minus 200 basis points. The average is calculated by weighting each firm's return by the value of the assets under management.

Finally, workers may switch between PFAs, provided they have made at least six months of contributions to the one they wish to leave. Unlike initial enrollment, switching cannot be done through a sales force agent; instead, it requires an in-person procedure at the PFA's office. Although this requirement has been simplified in recent years, switching rates remain relatively low compared to other Latin American countries with capitalization systems.

[14]The cap is set at 1.5 times the lowest fee in the market.

Having described the institutional setting and regulatory framework, we now turn to the data sources and key patterns that motivate our empirical approach.

2.3 Data and Descriptive Statistics

2.3.1 Sources

We combine several data sources on workers' characteristics and choices in the pension system, along with PFAs' financial statements and data on fees, sales force agents, and portfolio rates of return. First, we use a novel database of administrative records collected by BPS for a representative random sample of workers, covering the period from the market's inception in 1996 through 2020. The dataset is a monthly panel of worker records containing information on wages, employer characteristics, demographics (date of birth, sex, area of residence), and PFA enrollment decisions (enrollment mechanism and date, selected PFA, and an indicator for whether the worker ever switched). We complement this data with publicly available market-level data published by the Central Bank on market shares, fees, investment returns, contributions, switchers, and sales force agents for the 1996–2022 period. Finally, we use firms' financial statements for 2001–2020, which are also publicly available on the regulator's website.

2.3.2 Summary Statistics and Enrollment Patterns

We begin by summarizing the characteristics of the worker sample and enrollment behavior over time. In Table 2.1 we show the descriptive statistics of the random sample of workers. The median entry age to the formal labor market is 23.5 years and the median age at enrollment is 24.9 years. Also, conditional on enrollment and excluding mandatory enrollees, 75% of workers enroll within the first two

years of entering the labor market[15].

Table 2.1: Workers' Sample Summary Statistics

Individuals	125,453
Gender (female)	0.48
Age when entering the market (median)	23.2
Age when enrolling (median)	24.9
Gross wage (median, US$)	834
Share with enrollment gross wage above threshold (US$ 1,535)	0.15
Outside option (conditional on gross wage below US$ 1,535)	0.26

Notes. The Table reports descriptive statistics for selected demographics for the available sample. Average 1996-2020. UYU expressed in US$ 2017.

The median gross wage was USD 834, and 15% of workers earned gross wages above the first contribution threshold when entering the labor market. Among individuals whose wage eventually exceeded the mandatory enrollment threshold, 24% enrolled by default. Finally, among those who were not required to enroll (i.e., with gross wages below the first threshold), 74% voluntarily chose to save in an individual account within the capitalization sub-system. Beyond individual-level characteristics, several market-level patterns provide insight into the structure and dynamics of competition among PFAs.

2.3.3 Market Dynamics and Stylized Facts

We now turn to market-level dynamics, highlighting key patterns that inform the structure of our empirical model. Individual capitalization retirement systems evolve in two phases: an initial stage with mostly unenrolled workers and intense competition for new enrollees[16], and a later stage with a large stock of existing

[15]See Figure 3.11 in the Appendix

[16]A model for firms' competition in this period is proposed in Hastings, Hortaçsu, et al. 2017.

enrollees and a smaller inflow. This paper focuses on the latter phase, beginning in 2002 when new enrollees dropped below 10% of the total.

In Table 2.2, we present descriptive statistics at market level for the average year in the period we study. The average number of enrolled workers is 1,146,540, while the average cohort of new enrollees entering the formal labor market is 63,317, representing 5.5% of the stock. Additionally, the average number of switchers per year is low. Notably, the switching rate (0.31%) is even lower than those observed in similar regimes in other Latin American countries[17]. Low switching rates in this market have been documented and analyzed in previous studies (Luco 2019, Illanes 2016). Furthermore, conditional on switching, most enrollees leave private firms and switch (95%) to the public option. Workers leaving the public option or switching from one private firm to another are almost nonexistent.

Table 2.2: Market Summary Statistics

	Average Year	% Enrollees
Enrollees	1,146,540	
New enrollees	63,317	5.52%
Switchers	3,497	0.31%
Switchers towards SOE (95%)	3,322	
Switchers towards PF (5%)	175	

Notes. Average during period of analysis. Net switchers on average per year.

During the period analyzed in the paper, we identify three distinct market equilibria. These arose from structural changes, including a shift in the preferences

[17]See https://www.aiosfp.org/ for detailed information on the percentage of switchers over affiliates by country.

of the public option's shareholders following the 2005 change in government, and the introduction of a cap on market fees[18]. The shift in shareholder preferences led to lower fees, with clear effects on accounting profits. While we acknowledge that economic profits are the relevant concept for firm behavior and guide the analysis in the following sections, we highlight accounting profits due to their marked variation over time, which helps distinguish the three stationary equilibrium we characterize in our supply and demand model.

- Equilibrium 2002-2005: Relatively high SOE fee, no cap on management fees. During this period, the public option charged slightly lower fees than private firms but achieved similar accounting profits, as measured by Return on Equity (ROE).

- Transition 2006-2013: SOE reduces its fee, no cap on management fees. In 2006, the SOE began reducing its fee —a shift in behavior that we argue reflects a change in the preferences of its shareholders, as documented in minutes from public shareholder meetings. The policy aimed to benefit workers by allowing fee reductions as long as the ROE of the SOE remained above a minimum threshold set by the majority shareholder.[19].

- Equilibrium 2014-2017: Low SOE fee, no cap on management fees. The SOE continued to operate under the same shareholder mandate, but ROE became a binding constraint on further fee reductions. Compared to the first equilibrium, the fee of the SOE was cut by half, while private PFAs reduced theirs only slightly. As a result, during this period, the ROE of private PFAs increased, while that of the public option declined from 40%

[18]These equilibria are depicted in gray in Figure 2.3.

[19]Shareholders' meeting for fiscal year 2017: "(...) distributable profits stood at a ROE of 14.4%, exceeding the minimum requirement of 12% established by the majority shareholder." (...) "it is requested to continue with the fee reduction policy (...)"

to 12%. This provides preliminary descriptive evidence of a shift in the SOE behavior. However, in our model, we do not impose this behavioral change ex ante; rather, we recover it from the data.

- Transition 2018-2019: Low SOE fee, progressive implementation of a cap on management fees. Between 2018 and 2019, the regulator implemented a transition phase that allowed firms to adjust from pre-regulation fee levels to the new legal maximum. Once fully implemented, the cap limited fees to no more than 50% above the lowest fee in the market. The discussion surrounding the potential introduction of this cap began in December 2017, so it does not affect our analysis of the earlier period. The reduction in PFA fees observed during this phase is attributed to the new regulation. Since 2019, private PFAs have charged fees equal to the legal maximum.

- Equilibrium 2020: Low SOE fee, fully implemented cap on management fees. The cap on management fees led to a new equilibrium in which private PFAs' fees fell by half compared to the 2014–2017 average, and their ROE declined from approximately 60% to 20%.

Figure 2.1: Management Fee and Firms' Accounting Profits

(a) Management Fee (% Gross Wage) (b) Accounting Profits (ROE, %)

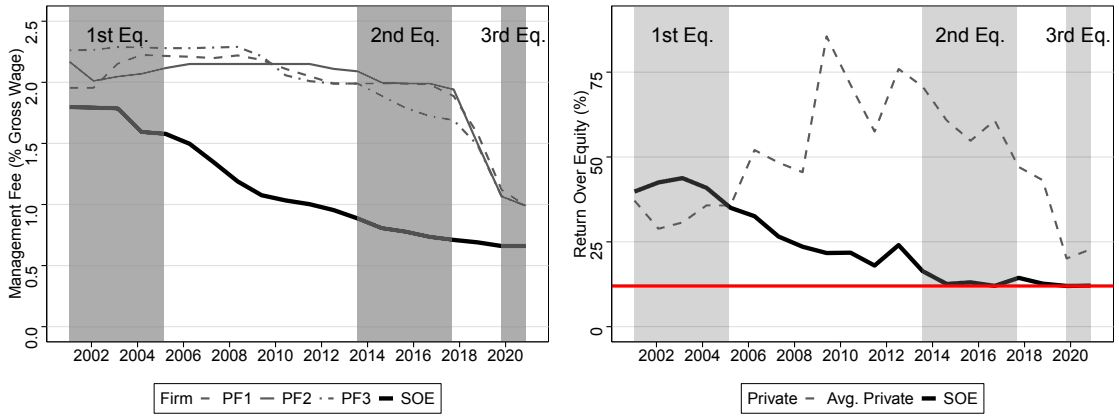

Note. Management fee as a share of the gross wage component relevant for social security contributions to the defined-contribution subsystem. The Return-Over-Equity is the ratio between distributed profits in t and equity in $t-1$. The red line represents the minimum acceptable ROE imposed by the main shareholder on the SOE. Shaded areas indicate the 3 equilibrium periods: 2002-2004, 2014-2017, and 2020.

Within each equilibrium period, the fees of private PFAs remain relatively stable and show little difference between them. Additionally, Figure 3.10 shows that the share of old workers evolves without significant variation over time. The public option is the market leader, with nearly 40% of total enrollees. There is also heterogeneity in market shares across the wage distribution, with the SOE capturing a larger share among higher-wage workers[20]. In fact, the public option is not the leading firm among low-wage workers.

To complete the description of firms' characteristics, Figure 2.3 shows the evolution and distribution of the real annual rate of return that firms generate for enrollees by investing their savings. Two main patterns emerge. First, returns are

[20]See shares by income bracket in Table 3.5 in the Appendix.

highly correlated over time[21]. Based on informal discussions with PFA managers and regulators, this correlation reflects both regulatory constraints —such as limits on the types and proportions of financial assets PFAs may hold— and behavioral responses. On the behavioral side, firms face a minimum rate of return regulation intended to protect workers from under-performance. In practice, firms respond by closely replicating each other's portfolios, resulting in similar returns.

Figure 2.3: Investment Real Returns (1 Yr)

(a) Investment Returns (Monthly) (b) Investment Returns (Fitted Density)

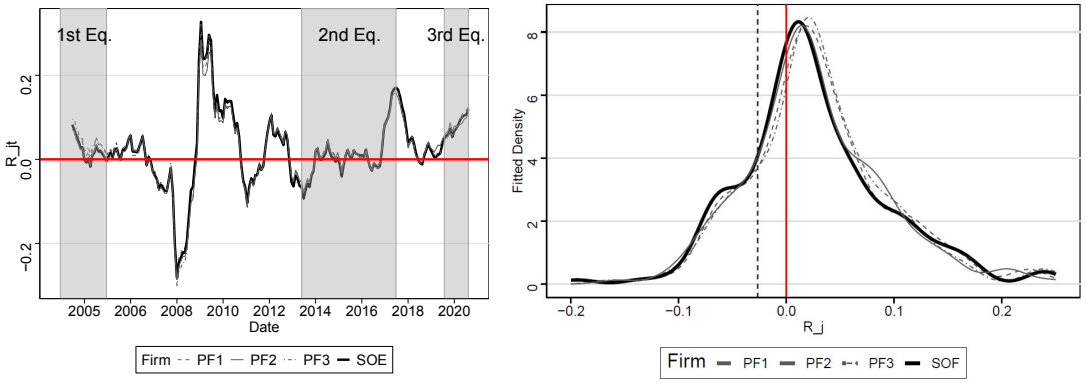

Second, some private firms are slightly more successful at generating higher returns. While short-term rankings fluctuate, certain private PFAs have consistently delivered superior performance[22]. Leveraging this fact, we use the minimum rate of return regulation to estimate firms' variable costs of "producing" those returns. Finally, regarding the sale force, the public option has the larger number of agents, with approximately 35%, a share share that is stable between the three periods[23].

[21]The lowest correlation coefficient in 1 Yr. real returns between any two firms during this period was 0.98.

[22]Importantly, whether these performance differences are sufficient to offset fee disparities is a question we examine in Section 3.2.2.

[23]In Table 3.6 of Appendix 3.2.2 we show the average share by firm in each period

2.4 Model

The model characterizes the competition between forward-looking firms in the oligopolistic market of Pension Fund Administrators to enroll workers and invest their social security contributions. Given the observed inertia in workers' initial enrollment decisions, we develop a model based on the literature on dynamic competition with switching costs (Beggs et al. 1992), where firms face an invest-harvest trade-off due to the presence of both new and existing consumers each period.

2.4.1 Workers

Workers' enrollment decisions d_{ijt} are the result of a two-stage decision process that repeats every period t and characterizes the choices of myopic and inattentive agents that display inertia, in the spirit of Ho et al. 2017. In the first stage, workers may be unaware and ignore the decision problem, or may receive an awareness shock that "wakes them up" and makes them go through the problem of deciding where to enroll. In the second stage, conditional on being aware, workers choose one of the available options.

Enrollment Decision - 2^{nd} Stage

For enrollment decisions we use a demand system in the tradition of Berry 1994, Berry et al. 1995. Workers make discrete choices between the $\mathcal{J}$ firms available at time t, and when available[24], a non-enrollment outside option -indexed by 0-. Worker i chooses the firm j that maximizes her indirect utility u_{ijt}, represented

[24]For workers with a gross monthly wage above USD 1,535 enrollment in a PFA is mandatory. For workers below that threshold, it is not, but more than 75% of workers enroll in one during the first two years of entering the labor market. Furthermore, enrollment is a terminal state. Once enrolled, a worker can switch between PFAs but cannot leave the market.

by an index function that depends on the management cost C_{ijt}, the potential stock of savings S_{ijt}, the number of sales force agents sf_{jt}, firm η_j and market ζ_t characteristics, and unobserved idiosyncratic preferences ϵ_{ijt}.

$$u_{ijt} = \underbrace{\theta \times C_{ijt}(\{w_{ik}\}_t^{t^r}, f_{jt})}_{\text{NPV}_i \text{ of Management Cost}} + \underbrace{\gamma_S \times S_{ijt}(\{w_{ik}\}_t^{t^r}, \mu_{jt})}_{\text{NPV}_i \text{ of Savings}} + \underbrace{\gamma_{SF} \times sf_{jt} + \eta_j + \zeta_t}_{\text{Common Components}} + \epsilon_{ijt} \tag{2.1}$$

Since we assume that workers initially believe they are making a permanent decision, the management cost and the potential savings stock are monetary net present value terms that depend on the flow of workers' gross wages $\{w_{ik}\}_t^{t^r}$, fees $\{f_{jk}\}_t^{t^r}$, and returns $\{\mu_{jk}\}_t^{t^r}$ between enrollment at t and retirement at t^r. We assume that workers are myopic in the sense that they believe that the fees and returns of period t will remain constant until retirement. However, we assume that they have better information about the path that their own wages will follow in the future and that they discount the future according to a common factor δ. Notice that even though fees and returns are common for all enrollees, the management cost and the potential savings stock are worker-specific. Cost and savings terms are calculated as follows:

$$C_{ijt} = \sum_{k=t}^{k=t^r} \delta^{k-t} f_{jt} w_{ik}, \tag{2.2}$$

$$S_{ijt} = \tau^{ssc} \sum_{k=t}^{k=t^r} \delta^{k-t} w_{ik} \times (1 + (\mu_{jt} - \bar{\mu}_t))^{t^r - k} \tag{2.3}$$

Furthermore, we allow workers to have different sensitivities for each term (θ, γ_S), as one can be more salient than the other at the decision stage. In this sense, the savings' stock measures the potential monetary gains (or losses) from selecting a firm with returns μ_{jt} above (or below) the market average $\bar{\mu}_t$, scaled by the social security contribution rate τ^{ssc}, with $f_{jt} \in [0, \tau^{ssc}]$. We also assume that

the sales force deployed in the market sf_{jt} can influence workers' decisions with sensitivity γ_{SF} and that unobserved idiosyncratic preferences are distributed Type One Extreme Value. The individual choice probabilities pr_{ijt} are

$$pr_{ijt}\big(\{w_{ik}\}_t^{t^r}, \mathrm{f}_t, \boldsymbol{\mu}_t\big) = \frac{\exp[\theta C_{ijt} + \gamma_S S_{ijt} + \gamma_{SF} sf_{jt} + \eta_j + \zeta_t]}{1 + \sum_k \exp[\theta C_{ikt} + \gamma_S S_{ikt} + \gamma_{SF} sf_{kt} + \eta_k + \zeta_t]} \tag{2.4}$$

Awareness - 1^{st} Stage

To reach the second stage, workers must be aware of the decision problem. A worker's awareness is represented by a binary random variable A_{ijt}, which takes the value 1 if the worker is aware and 0 otherwise. For workers entering the labor market for the first time (i.e., potential new enrollees), we assume full awareness; thus, they make a choice consistent with the model described in the previous section. In contrast, workers who enrolled before period t (old enrollees) receive an awareness shock ε_{ijt} each period. When this shock exceeds a threshold ($\varepsilon_{ijt} \geq \bar{\varepsilon}_{ijt}$), the worker re-optimizes. We assume the shock follows a standard normal distribution $\varepsilon_{ijt} \sim \mathrm{N}(0,1)$. The awareness probability aw_{ijt} depends on a firm fixed-effect υ_j and a dummy variable indicating whether the simplified switching procedure introduced in 2013 is in effect (Post_{2013}):

$$aw_{ijt} = \Phi(\beta_0 + \beta_1 \times (f_{jt} - f_{SOE,t}) + \beta_3 \times (\mu_{jt} - \mu_{SOE,t}) + \beta_4 \times \mathrm{Post2013} + \upsilon j) \tag{2.5}$$

Consistent with what we observe in the data[25], in the baseline model we work under that assumption that conditional on receiving a shock that makes them re-optimize, workers decide to move to the state-owned enterprise. Therefore, only old enrollees in a private firm may re-optimize in the baseline model. Furthermore, this is the reason why the cumulative distribution function that calculates the

[25]See Figure 3.12

probability of being aware aw_{ijt} is a function of the difference between the fees and returns of private firms relative to those of SOE. Notice that while being aware and switching from a private PFA to the SOE are equivalent in this model, we revisit this assumption in our counterfactual scenarios[26].

2.4.2 Firms

In the model, oligopolistic, forward-looking firms compete to enroll workers and manage their social security contributions. The strategic behavior of firms manifests in two strategic variables: management fees f_j and mean portfolio returns μ_j[27].

Since firms can charge only a single management fee and enrolled workers exhibit inertia, firms face an invest-harvest trade-off (Beggs et al. 1992). On the one hand, firms have incentives to "harvest" by charging a high fee to profit from their existing base of enrollees. On the other hand, by charging a lower fee, they have incentives to "invest" in obtaining a higher market share of new workers that will remain enrolled between t and t^r -with high probability - and therefore to make more profits from them in the future.

The model captures a second trade-off related to the portfolio returns that firms obtain from investing workers' savings. On the one hand, obtaining higher returns is good for firms for two reasons. First, because returns attract workers and increase revenues through enrollment decisions[28]. Second, because higher returns reduce the probability of capitalization, and therefore reduce firms' expected costs associated with the minimum rate of return regulation[29]. However, achieving

[26] In those cases, workers that re-optimize do choose according to the second stage model we stated in the previous section.
[27] While we allow the sales force to influence the enrollment probability, in the model we treat it as non-strategic.
[28] See Equation 2.1.
[29] See Institutional Section 2.2.3

higher returns is costly. We assume that firms can control the mean return of their portfolio through their asset allocations but that a higher return exposes them to higher risk. We take advantage of the institutional environment to estimate the implicit cost of generating portfolio returns.

Revenues

Firms obtain revenues each period t by charging a management fee f_{jt} to their enrollees, expressed as a share of their gross wage[30]. The fee affects the choices of new and old enrollees according to the decision model of Section 2.4.1. We work with a stationary model of the labor market and re-express choice probabilities as monetary shares that aggregate across workers[31]. These shares depend on choice probabilities pr_{ijt}, but re-weight workers according to their wages w_{it}, consistent with the fact that firms obtain higher revenues by enrolling higher wage workers.

In the stationary model of the labor market, the mass of wages of new workers is a fixed fraction α of the aggregate wage mass M_t. Under these assumptions, we express expected revenues as the management fee multiplied by the sum of two components: 1) the share of the wage mass from new enrollees, $s^n_{jt}(\mathrm{f}_t, \boldsymbol{\mu}_t)$, applied to αM_t, representing workers entering the labor market; and 2) the share of the wage mass from existing enrollees, $s^o_{jt}(\mathrm{f}_t, \boldsymbol{\mu}_t)$, applied to $(1-\alpha)M_t$, accounting for continuing workers net of those who retire. Furthermore, for the case of private firms, revenues also consider the probability that old enrollees in firm j may become aware and leave to enroll in the SOE[32]. The following equation describes

[30]Because the social security contribution rate is fixed (15%) and the management fee is deducted from these contributions, the fee can be also expressed as a share of the social security contributions ($f_{jt}/0.15$).

[31]See Appendix 3.2.3 for additional details.

[32]Consistently, these switchers represent additional revenues for the SOE. See Equation 3.5 in the appendix for a description of this case.

firms' expected revenues each period.

$$\mathbb{E}\big[Y_{jt}\big] = f_{jt} \times \Big(\underbrace{\alpha\, M_t\, s^n_{jt}(\mathrm{f}_t, \boldsymbol{\mu}_t)}_{\text{New Workers' Wages}} + \underbrace{(1-\alpha)\, M_t\, s^o_{jt}(\mathrm{f}_t, \boldsymbol{\mu}_t)\,(1 - aw_{jt}(\mathrm{f}_t, \boldsymbol{\mu}_t))}_{\text{Old Workers' Wages}} \Big) \quad (2.6)$$

Costs

Firms incur economic costs from their regular business activities and from complying with the regulation in place. The model considers the variable costs of 1) enrolling new workers, 2) obtaining portfolio returns, and 3) complying with the regulations regarding 3.1) the minimum rate of return that firms must achieve for enrollees and, 3.2) the equity they must put aside and invest mirroring the portfolio that enrollees have.

The cost of enrolling new workers is a variable cost related to the payments that firms make to their sales force agents as a reward when they bring in new enrollees. In practice, this payment is tied to the "quality" of the new enrollee in terms of the wage, the expected density of future contributions, etc. To better reflect this feature, in the model the marginal cost of enrolling an additional worker depends on the wage of the new enrollee: we express the cost as the product of the enrollment probability $pr_{ijt}(f_t, \mu_t)$, the wage of the new enrollee w^n_{it}, and a firm-time specific cost per dollar CPD_{jt}.

Firms also incur a variable cost $f(\mu_{jt}|\kappa_j)$ to achieve the mean return μ_{jt} for their enrollees. This cost is related to optimal portfolio theory and reflects the increasing volatility risk firms must manage when seeking higher mean returns. This is a reduced-form mechanism that summarizes the economics of an efficient frontier without risk-free assets (Markowitz 1952). Furthermore, based on what we observe in the data, we allow firms to have different investment abilities, as summarized by the firm-specific parameter κ_j.

Firms benefit from higher returns not only because they increase the enrollment probability, but also because they reduce -ceteris paribus- the expected cost of capitalization. The regulation requires firms to compensate workers with their own equity when realized returns R_{jt} fall below a threshold $r_{min,t}$[33]. Since both returns are random variables, the economic cost associated with this risk is determined by their joint distribution[34]. When this event occurs ($R_{jt} \leq r_{min,t}$), the capitalization cost depends on the gap between the realized returns and the stock of workers' savings managed by firm j (PSF_{jt}, Pension Savings Fund).

Finally, to further incentivize firms to achieve good returns for enrollees while simultaneously mitigating the moral hazard of investing others' money, the regulation requires firms to invest their own equity, mirroring the asset portfolio they select for enrollees. The required equity investment is a fixed fraction (0.5%) of the total stock of savings they manage (PSF_{jt}). This regulation imposes an opportunity cost on firms by constraining their portfolio allocations along multiple dimensions, including asset type, origin, and credit rating, among others, thereby limiting their ability to allocate their equity freely. We calculate this cost using the benchmark rate of return r_t^* that would be obtained from an unconstrained allocation.

Taken together, the elements of the expected cost of firm j in period t described

[33]The threshold $r_{min,t}$ is the minimum between 2% and the average return of the market minus 2%: $r_{\min,t} = \min\{\bar{R}_t - 2\%; 2\%\}$. See Appendix 3.2.3 for details.
[34]See Appendix 3.2.3 for details of this probability distribution.

previously are summarized in the following equation:

$$\begin{aligned}\mathbb{E}\big[C_{jt}\big] &= \underbrace{\sum_{i_{c=t}} pr_{ijt}(\mathrm{f}_t, \boldsymbol{\mu}_t) \cdot w_{it} \cdot CPD_{jt}}_{\text{Enrollment Cost of New Workers}} + \underbrace{f(\mu_{jt}|\kappa_j)}_{\text{Investment cost}} \\ &+ \underbrace{\mathbb{E}\Big[\big|(r_{\min,\mathrm{t}} - R_{jt})\big| \cdot \mathrm{PSF}_{jt}\Big]}_{\text{Expected Capitalization Cost}} + \underbrace{\mathbb{E}\Big[0.5\% \cdot (r^*_t - R_{jt}) \cdot \mathrm{PSF}_{jt}\Big]}_{\text{Expected Opp. Cost of Firm } j\text{'s Equity}}\end{aligned} \tag{2.7}$$

Firm Problem and Equilibrium

We solve the problem of the firm using the elements related to workers' preferences and firms' revenues and costs described above. We assume that firms play a dynamic, full information, simultaneous move game each period. In the game, firms set the sequences of management fees $\{\mathrm{f}_{jk}\}_t^T$ and mean returns $\{\mu_{jk}\}_t^T$ between t and a terminal period T, that maximize the present discounted value of their objective function $\mathcal{W}(\{\mathrm{f}_{jk}, \mu_{jk}\}_t^T)_{jt}$.

For private firms, the objective function considers only economic profits, whereas for the state-owned enterprise, it incorporates a weighted sum of profits and the stock of its enrollees' savings. We denote the weight on workers' savings as Non-Profit Motives (NPM), represented by the parameter λ, with the remaining weight on profits given by $(1 - \lambda)$.

$$\mathcal{W}(\{\mathrm{f}_k, \mu_k\}_t^T)_{jt} = (1-\lambda) \underbrace{\left(\mathbb{E}\Big[V(\{\mathrm{f}_k, \mu_k\}_t^T)_{jt}\Big]\right)}_{\text{NPV Profits}} + \lambda \underbrace{\left(\mathbb{E}\Big[\text{Savings}(\{\mathrm{f}_k, \mu_k\}_t^T)_{jt}\Big]\right)}_{\text{NPV Workers' Savings}} \tag{2.8}$$

$$1)\ \ \mathbb{E}\big[V(\mathrm{f},\mu)_{jt}\big] = \sum_{k=t}^{k=T} \beta^{k-t}\mathbb{E}\big[\pi_{jk}\big] = \sum_{k=t}^{k=T} \beta^{k-t}\mathbb{E}\big[Y_{jk} - C_{jk}\big]$$

$$2)\ \ \mathbb{E}\Big[\mathrm{Savings}(\mathrm{f},\mu)_{jt}\Big] = \sum_{k=t}^{k=T} \beta^{k-t}\mathbb{E}\big[\sum_{i} \sum_{l=k}^{l=k+(t_i^r-t)} w_{il}(\tau^{ssc} - f_{jl}) \prod_{l}^{l=k+(t_i^r-t)} (1+R_{jl})\big]$$

We work under the assumption that observed fees and mean returns constitute a pure-strategy, stationary "no-sales" Nash equilibrium of the game, in the spirit of Farrell et al. 2007. This implies that equilibrium fees and mean returns are constant in every period k between t and T: $\{\mathrm{f}_k^*, \mu_k^*\}_t^T = (\mathrm{f}_t^*, \boldsymbol{\mu}_t^*)$. Firms know the distribution of the preference shocks (ϵ) and awareness (ε) shocks that workers receive but not their realizations. The timing of events is as follows. First, firms choose simultaneously fees and mean returns to maximize their objective function. Afterwards, workers observe their preference and awareness shocks and decide.

2.5 Estimation and Results

2.5.1 Demand

Specification Details

We estimate workers' demand in the second stage following Hastings, Hortaçsu, et al. 2017. In order to allow for flexible preference heterogeneity, we estimate conditional logit models separately for five demographic brackets (b) using workers' micro data. First, we divide the population into those with and without an outside option. For those with an outside option, we further classify individuals into four quartiles based on the distribution of gross wages. For individuals without one, we construct a single group. Therefore, the vector of sensitivity parameters $\{\theta^b, \gamma_S^b, \gamma_{SF}^b\}$ and the firm η_j^b and time ζ_t^b fixed effects are bracket-specific.

To compute the net present values of the management cost $C_{ijt(i)}$ and of the

potential savings stock S_{ijt}, we assume that workers are active in the labor market for 40 years before retirement. Additionally, for the wage w_{ik} we use the average wage over workers' life cycle $w_i = \frac{1}{40}\sum_l^{l+40} w_{il}$. To estimate this wage, we first regress observed wages against age and age-squared, and then compute the average wage consistent with the estimated wage curve[35]. Finally, to calculate the stock of potential savings S_{ijt}, we assume that the mean return rate μ_{jt} is equal to the average 12 months' return during the year before the enrollment decision.

We use a probit model to estimate the awareness probability in the first stage. Due to the fact that we only have aggregate switchers' data at the firm level, we use data from the entire period for which we have information (2002-2020). Furthermore, we assume that the shock ε_{ijt} is common to all enrollees within a firm ($\varepsilon_{ijt} = \varepsilon_{j(i)t}$).

Identification

The identification of the cost and savings stock sensitivities $\{\theta^b, \gamma_S^b\}$ is based on the fact that fees and returns are plausibly exogenous[36]. In this setting, firms set national fees and have single portfolios for their enrollees, but costs and potential savings are worker-specific and vary with gross wages and spells in the formal labor market. This individual-level variation of the costs and benefits that every PFA delivers, even among individuals with similar demographic characteristics, gives us arguably exogenous variation to estimate the target parameters.

2.5.2 Results

We estimate the conditional logit model by maximum likelihood separately for workers in each bracket b. In Figure 2.4 we summarize the results by displaying

[35]See Appendix 3.2.4 for details.
[36]Hastings, Hortaçsu, et al. 2017 use a similar argument for the identification of the fee sensitivity.

the median elasticity to the management cost C_{ijt} and the potential savings stock S_{ijt} for individuals in each estimation bracket. We calculate these using estimated parameters, observed fees and returns, and workers' characteristics.

Figure 2.4: Median Elasticity by bracket

(a) Management Cost (b) Potential Savings

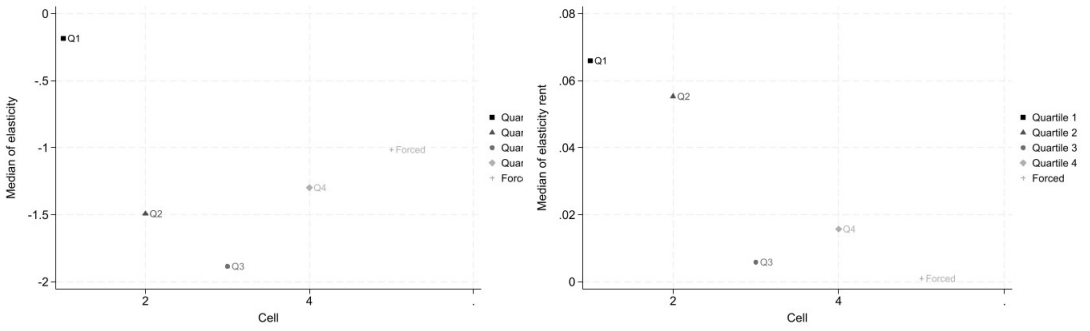

Note. Within-bracket median PFA elasticity, all periods. Elasticities are calculated using observed fees, mean returns, and gross wages.

The estimates show that higher-wage individuals are more elastic than lower-wage ones with respect to changes in the management cost[37], but this pattern reverses for the savings component. In particular, the correlation between the wage level and the sensitivity to the management cost is in line with the previous literature on individual capitalization pension systems in other Latin American countries Hastings and Ashton 2008, Hastings, Hortaçsu, et al. 2017, and on financial literacy Lusardi 2008.

Furthermore, while both terms are monetary metrics, management costs appear to be more salient than potential savings to workers when making enrollment decisions. This is reflected in the magnitude difference between the estimated elasticities of both terms. Considering the differences in management fees that exist

[37]In Figure 3.15 we show the management fee elasticities for each period.

between private firms and the SOE, the estimates are consistent with the fact that higher wage workers enroll disproportionately more in the SOE than low wage workers[38]. This pattern is relevant for understanding the welfare implications of alternative regulations and market configurations.

We estimate the first-stage awareness probit with maximum likelihood[39]. In the baseline model, the results imply that an increase of one standard deviation in the fee differential ($f_{jt} - f_{SOE,t}$) raises the probability of leaving firm j to the SOE by 9 p.p. Similarly, an increase of one standard deviation in the mean return differential ($\mu_{jt} - \mu_{SOE,t}$) reduces the probability of leaving firm j to the SOE by 2 p.p.

We new use the system of first-order conditions with respect to fees and mean returns from firms' problem, the estimated demand, and market level data to recover the vector of enrollment cost per dollar (CPD_{jt}), the parameters of the investment cost function (κ_j) and the weight related to the state-owned enterprise Non-Profit Motives (λ).

2.5.3 Enrollment Marginal Cost and Non-Profit Motives

We recover the cost per dollar of enrolling a worker with average wage w_i in three different periods: 2002–05, 2014–17, and 2020. As described in Section 2.3.3, we interpret these as three stationary equilibrium, reflecting the change in shareholder preferences on the SOE's board in 2005 and the introduction of a cap on management fees in 2018. In Table 2.3 we show the enrollment marginal cost for the first two equilibrium and the three private firms. For private firms, we estimate enrollment marginal costs between USD 10 and USD 30 in the first period, and between USD 40 and USD 67 in the second.

[38]See Figure 3.16 in Appendix 3.2.4
[39]In Table 3.7 we display the estimation results.

Table 2.3: Enrollment Marginal Cost and Non-Profit Motives

	Average CPD (USD)				SOE NPM ($\hat{\lambda}$)
Equilibrium	PF 1	PF 2	PF 3	SOE ($\lambda = 0$)	(Max. MC_{PF})
2002-05	10	29	21	-275	0,63
2014-17	45	67	40	-394	0,86

Note: Cost of enrolling a new worker with the average gross monthly wage, inflation-adjusted USD (2017).

Remarkably, we cannot separately identify the cost per dollar CPD and non-profit motives λ for the state-owned enterprise. However, ignoring non-profit motives would require accepting a negative marginal enrollment cost to rationalize the observed equilibrium variables under the implicit pure profit-maximizing behavior. Given that firms in this market incur variable costs when enrolling additional individuals, this scenario seems unlikely. The marginal cost in this market is mainly associated with paying sales force agents to enroll new workers entering the labor market[40]. Therefore, we proceed by assuming that the marginal cost of the SOE is equal to the marginal cost of the less efficient private firm (PF 2). The direct implication of this assumption is that non-profit motives effectively increased between the first and the second equilibrium.

[40]Although we do not directly observe the variable wage component that sales force workers receive for each affiliation, using data on the sales force average productivity (calculated as the ratio between monthly new enrollees per firm over monthly total sales force agents), and aggregated wages, and on their minimum wage established in Collective Agreements, we can approximate the observed variable wage they earn. For 2017 this calculation implied a mean variable wage component of USD 63 dollars per enrollee, a figure close to our enrollment marginal cost estimates.

2.5.4 Variable Investment Cost

The remaining unobservable primitive we recover is the vector of parameters κ that controls the variable investment cost that firms pay to obtain units of portfolio returns. To do so, we use a parametric cost function of the form $f(\mu_{jt}|\kappa_j) = \exp(\kappa_j \cdot \mu_{jt})$. In this case, different from the enrollment marginal cost, we assume that firms have a time-invariant ability to obtain those returns. Table 2.4 we display the mean return, the expected capitalization probability consistent with observed returns and the minimum rate of return regulation, and the estimated κ for each firm.

Table 2.4: Investment Cost

Variable	PF 1	PF 2	PF 3	SOE
μ_j^*	1.84	1.17	1.71	1.53
$\mathbb{E}[\text{Pr(Cap)}_j]$	7.6%	10.7%	9.8%	4.3%
$\hat{\kappa}$	944	1,518	1,089	1,126

$\omega = [0.18; 0.09; 0.16; 0.57]$, $\sigma = [0.089; 0.088; 0.095; 0.094]$,

$\rho_{jk} = 0.9$

Notice how the equilibrium effect of the regulation results in smaller firms having a higher capitalization probability, and this occurs despite obtaining higher returns on average. Due to its size, the SOE has the lowest probability (4.3%), which is almost half the probability of the private firm that achieves the highest mean return (7.6%). Therefore, the fact that some private firms achieve higher returns for their enrollees while facing higher capitalization probabilities is, in our model, consistent with a lower cost to generate those returns and a lower value of the associated parameter $\hat{\kappa}$.

2.6 Counterfactual Policies

We use our estimates about preferences and marginal costs to understand the value of the public option and the effect of the regulation on fees in the market. In particular, we analyze two counterfactuals. In the first one, we substitute the public option with a private firm that resembles the other private options. Here we try to capture the value of the public option by comparing the observed market equilibrium with an alternative configuration with private PFAs only. In the second, we analyze how far the observed equilibrium was from a benchmark where fees were set such that PFAs have zero economic profits. We separate the analysis between 2014-2017 and 2020 to account for the effects of the introduction of the fee regulation in the later period. We calculate counterfactual equilibrium fees using the demand primitives of the period of interest, but always with the marginal cost parameters of 2014-2017.

2.6.1 Privatization

The privatization counterfactual implies simulating the equilibrium we would observe in an oligopoly with four private firms, instead of the one we observe with three private firms and one public option. The approach we follow implies asking "what if the SOE is sold to a private investor, a new board takes over, and all enrollees were informed about the change". To do so, we must address how a privatization affects certain demand and supply side aspects of our model.

Implementation Details

In the case of supply, we consider changes in non-profit motives and in harvesting incentives due to a large pre-existing base of enrollees. In the first case, in the

counterfactual firms care only about economic profits, so we set the value of the parameter that summarizes non-profit motives in the public option to zero ($\lambda = 0$). In the second, we deal with the sizable harvesting incentives by redistributing the share of old enrollees among all four firms[41].

In the case of demand, we consider how the privatization affects the awareness probability (first stage) and the enrollment decision (second stage). For the probability that worker i enrolled in firm j becomes aware[42], we assume that the chance of receiving a large attention shock and re-optimize is now a function of the differences between firm j's fee and mean return, and their associated market averages[43]. For the enrollment decision, we change the estimated intrinsic mean valuation η^b_{SOE} that workers have for the public option, and replace it with the average valuation they have for private firms $\overline{\eta_j}^b$. These changes in the two-stage decision problem imply that awareness and switching are not equivalent events any more[44]. Furthermore, for enrollees that decide to switch, we assume that the firm pays the enrollment marginal cost CPD.

Privatization Results

The counterfactual indicates that, in equilibrium, the public option reduces fees, increases portfolio returns, and raises the expected savings of all workers. In our preferred specification,[45] the privatization of the SOE would increase average fees by 8% and reduce average returns by 3% (4 bp), relative to the stationary

[41]The public option is the largest firm, so only setting non-profit motives to zero would imply a massive increase in the equilibrium fee of this firm. In the main specification we redistribute the old enrollees of the SOE in equal parts among all four private firms.

[42]We retain the assumption that new workers are aware with certainty.

[43]In the baseline, this probability was a function of the difference between private firms' fees and mean returns relative to those of the SOE. More importantly, awareness and switching coincided but now they do not.

[44]For old enrollees, being aware and switching to from a private firm to the SOE were equivalent events in the baseline model. This was motivated by the switching pattern that we observe in the data (See Figure 3.12).

[45]Consistent with the specification details described in the previous section.

equilibrium in 2014–17. This implies that the expected savings at retirement for a new worker who enrolls in the new market equilibrium and remains active for forty years would decrease by 6.7%, on average. In contrast, the profits of the four firms more than double. We present these results in Table 2.5, where we decompose the effects of the different implementation assumptions we adopt.

Table 2.5: Privatization - Four Private Firms Oligopoly

	Fees f_j^* (% Gross Wage)		Returns μ_j^* 1Yr (%)		$\mathbb{E}[\pi_t]$ (US$ Mill.)		$\mathbb{E}[\text{Savings}]^*$ (US$ '000)	
	PFs	SOE	PFs	SOE	PF (Tot.)	SOE	PFs	SOE
Baseline Equilibrium (Avg. 2014/17)	1.94	0.8	1.36	1.33	40.97	6.87	40.75	44.61
Counterfactual								
1) $\rightarrow \hat{\lambda} = 0$	1.9	4.57	1.35	1.03	37.25	142.03	40.82	28.85
2) $\rightarrow \hat{\lambda} = 0 \ + \ s_{SOE}^{old} = s_j^{old} = 0.25$	2.1	2.51	1.34	1.25	85.10	22.48	40.04	37.65
3) $\rightarrow$ 2) $+ \eta_{SOE} = \bar{\eta}_{pf}$	2.1	2.01	1.34	1.27	84.83	13.73	40.06	39.64

* Mean savings for a worker facing equilibrium f^* and μ^* for 40 years. PF averages are weighted by enrollment shares s_j. Awaking probability is a function of the difference between a firm's fee and the market average. When consumers awake, they meet the sales force (as in the first choice). Marginal costs apply every period.

In terms of heterogeneity, while the direction of the changes is common across firms, the magnitude of the effects is not. For workers enrolled in the SOE, privatization implies a 150% increase in fees and a 4.5% (6 bp) reduction in the return rate, which together reduce the expected savings of the average worker by 11.1%. The magnitudes are more modest in the case of private firms: privatization results in an average increase in fees of 8% and a 1.5% (2 bp) reduction in the mean return rate, with an associated 1.7% decrease in workers' savings. Overall, the results show that competition in the market is modest and that the benefits delivered by the SOE's non-profit motives are largely concentrated among its own enrollees. The combination of a modest demand sensitivity to fee differences —which leads sometimes enrollees to select dominated options— and the absence of strong entry

incentives for new firms[46] imply that the privatization of the SOE would reduce savings across the board.

2.6.2 Inattention

A typical argument against supply-side regulations is that demand-side policies can make the market more competitive through changes in consumer behavior. To examine whether such policies could offset the negative effects of privatization in this market, we simulate the effects of a policy that reduces inertia. To conduct this analysis, we increase the sensitivity of enrollees to the difference between the fee of their current option j and the market average, which, in the context of our model, implies an increase in the parameter β_1.

As displayed in Table 2.6, results show that reducing inattention would increase competition and benefit all enrollees. Because old enrollees are more likely to re-optimize when a firm charges a fee above the market average, harvesting incentives are moderated, and therefore the equilibrium fees of all firms are lower in equilibrium. However, for this to compensate for the negative effects of privatization, the increase in β_1 must be sizable. For the average fee to decrease by 27.5% and the stock of savings to increase by 5.4%, the sensitivity must increase by a factor of four.

[46]In Appendix 3.2.5 we analyze these incentives in more detail.

Table 2.6: Privatization + Demand-Side Policy

	Fees f_j^* (% Gross Wage)		Returns μ_j^* 1Yr (%)		$\mathbb{E}[\pi_t]$ (US$ Mill.)		$\mathbb{E}[\text{Savings}]^*$ (US$ '000)	
	PFs	SOE	PFs	SOE	PF (Tot.)	SOE	PFs	SOE
Baseline Equilibrium (Avg. 2014/17)	1.94	0.80	1.36	1.33	40.97	6.87	40.75	44.61
Privatization	2.1	2.01	1.34	1.27	84.83	13.73	40.06	39.64
Counterfactual (Privatization + ↑ Sensitivity)								
$2 \times \hat{\beta}_1$	1.86	1.81	1.35	1.28	71.85	10.22	40.98	40.39
$3 \times \hat{\beta}_1$	1.66	1.64	1.35	1.28	61.51	7.09	41.72	41.07
$4 \times \hat{\beta}_1$	1.49	1.48	1.35	1.29	52.81	4.34	42.36	41.67

* Mean savings for a worker facing equilibrium f^* and μ^* for 40 years. PF averages are weighted by s_j.

Furthermore, even if a demand-side policy could effectively achieve this, the combination of privatization and inattention reduction policies would still create winners and losers relative to the baseline environment with three private firms and one public option. Note that while the former implies higher savings for the enrollees of private firms, it leaves enrollees in the former SOE worse off.

Due to the fact that demand-side policies have not been able to reduce equilibrium fees substantially in this market before, we consider the required increase to compensate for the privatization negative effects for workers to be unlikely. Looking from a different angle, it is highly likely that the public option increases retirement savings for everyone, at the same time that it creates winners and losers among the pool of workers that participate in the market. For this reason, we explore now how it performs relative to another common policy in this market: a cap on the management fee.

2.6.3 Regulation by public options

Having shown that a public option with non-profit motives increases workers' savings relative to an oligopolistic market with only private firms, we now examine how far such motives can go as an instrument to improve equilibrium outcomes for workers. We also compare these outcomes with those achieved under a cap on management fees. This is motivated by the fact that many countries rely on fee caps as a regulatory instrument in the market of pension fund administrators, and that we observe the consequences of its introduction in the case we study.

In our baseline specification, we found that the reduction of the management fee of the SOE observed between the first equilibrium (2002-05) and the second (2014-17) was consistent with an increase in its non-profit motives (the weight the SOE assigns to workers savings in its objective function) from 0.63 to 0.86. Accordingly, in this counterfactual, we use the second equilibrium as baseline ($\lambda = 0.86$), and increase the SOE's non-profit motives to match two potential targets: (1) an extreme case with a zero management fee (free provision), and (2) a less extreme case with an equilibrium fee that implies zero economic profits. In Table 2.7 we display these results and compare them with the observed equilibrium after the cap on fees was introduced.

Table 2.7: Non-Profit Motives and Cap on Management Fees

	Fees f_j^* (% Gross Wage)		Returns μ_j^* 1Yr (%)		$\mathbb{E}[\pi_t]$ (US$ Mill.)		$\mathbb{E}[\text{Savings}]^*$ (US$ '000)	
	PFs	SOE	PFs	SOE	PF (Tot.)	SOE	PFs	SOE
Baseline Equilibrium (Avg. 2014/17)	1.94	0.8	1.36	1.33	40.97	6.87	40.75	44.61
Counterfactual: Non-Profit Motives λ								
$\lambda: f_{\text{SOE}} = 0$	2.03	0.00	1.34	1.34	43.56	-22.79	40.27	47.58
$\lambda: \mathbb{E}[\pi_{\text{SOE}}] = 0$	1.95	0.61	1.36	1.34	41.27	0.00	40.68	45.31
Cap on Fees								
Observed Equilibrium (2020)	0.99	0.66	1.36	1.34	8.59	2.13	44.27	45.13

* Mean savings for a worker facing f^* and μ^* for 40 years. PF averages are weighted by enrollees s_j.

Though magnitudes differ across targets, the main effect of raising non-profit motives is to increase market segmentation. In both cases, private firms raise their equilibrium fees in response to the reduction in the SOE's fee. High-wage workers have greater elasticity to changes in management costs ($C_{ijt}(f_{jt}, w_{it})$), so even if they are equally likely as lower-wage workers to receive an attention shock, they are still more likely to switch to the SOE when f^*_{SOE} decreases. As a result, the pool of enrollees that remain in private firms becomes more inelastic, and harvesting incentives increase, as reflected in higher management fees relative to the baseline equilibrium.

The previous result suggests that using non-profit motives to enhance competition may have redistributive consequences and exacerbate wage inequality in retirement. Figure 2.5a illustrates the aggregate equilibrium effect on savings when non-profit motives are set to achieve the zero-profits target, relative to the baseline. The fact that the savings of low-wage workers remain almost unchanged reflects a net effect: lower savings for low-wage workers who stay in private firms are fully offset by higher savings for those who enroll in the SOE.

Figure 2.5: Savings at Retirement by Wage Decile - Counterfactual vs Baseline

(a) Non-Profit Motives ($\lambda_{\mathbb{E}[\pi_{SOE}]=0}$) (b) Cap on Management Fees

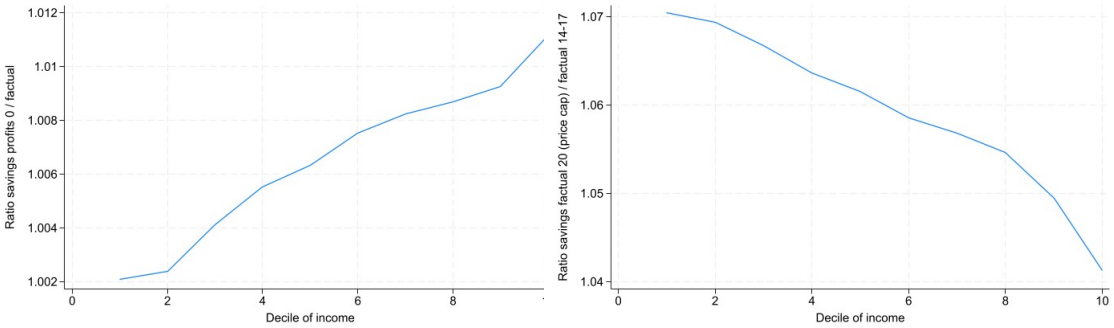

Note: "Baseline" stands for the 2014-17 stationary equilibrium before the introduction of the cap. The figure shows the policy effects at the market level, aggregating enrollees of private firms and the public option.

Cap on Management Fees

Finally, we compare the welfare effects of non-profit motives with those of introducing a cap on management fees. This policy was implemented in 2020 following a two-year transition period and has been adopted by other regulators worldwide to increase workers' net savings. We simulate the policy as designed by the Uruguayan regulator: the maximum management fee cannot exceed 1.5 times the lowest fee in the market,[47] which in our case corresponds to the SOE.

In our model, the equilibrium effects of the cap are ambiguous. On the one hand, the policy limits the harvesting incentives of private firms and may reduce fees across all firms, including the public option. On the other hand, it reduces the fee gap of privates firms with respect to the SOE, and thus reduces the likelihood that workers in private firms receive an attention shock that triggers re-optimization. Additionally, as the cap alters enrollment decisions and market

[47] As the SOE had the lowest management fee, this design may, in theory, give the SOE an incentive to lower its fee to the point where private firms incur losses and exit the market. However, since neither the SOE nor the regulator aims to establish a monopolistic public option, we abstract from this possibility.

shares, the resulting effect on mean returns is also uncertain.

The results show that the cap increases savings for all workers, with relatively larger gains for low-wage workers. More importantly, while private firms now charge the maximum allowed fee, differences in expected savings across firms shrink, primarily because returns are—marginally—higher in private firms. Before the cap, an average worker enrolled in a private firm was expected to accumulate 91.3% of the savings she would have earned in the SOE. After the cap, this figure rises to 98%.

2.7 Conclusion

This paper examines the welfare consequences of introducing a state-owned enterprise in the market of pension fund administrators to compete with private firms. This market is usually characterized by workers' inertia in enrollment decisions and low sensitivity to fees and returns, features that translate into high management fees. Using rich data on workers' enrollment decisions, we estimate a demand and supply model where forward-looking firms compete to enroll workers and the SOE maximizes an objective function that considers profits and non-profits' motives. We then use the model to evaluate counterfactual policies in this market.

First, we simulate a privatization scenario in which the SOE is replaced by a private firm with standard profit motives and cost structures. We also explore whether a demand-side policy that reduces inattention among workers can offset the negative effects of privatization, finding that although such a policy increases competition and savings, it requires unrealistically high workers' responsiveness to fee differentials to fully compensate for the replacement of the public option. Next, we assess how enhancing the SOE's non-profit motives affects outcomes, showing that stronger non-profit behavior reduces fees and raises savings for its

own enrollees but also increases market segmentation and inequality. Finally, we compare this intervention with the introduction of a cap on management fees—a policy implemented in Uruguay in 2020 and used in other pension systems—which proves more effective at increasing savings across the board and narrowing gaps between public and private enrollees, especially benefiting low-wage workers.

Our findings suggest that creating a public option with high not-for-profit motives helps to increase workers' savings through lower fees and higher returns. However, due to moderate demand elasticities to fees and lack of entry of new competitors, the presence of the public option is not enough to meet the increase in savings that the cap in management fees generates for workers. In the environment we study, the combination of a public option with a cap on fees seems to perform better than an alternative regulation aiming at privatize and reduce workers inertia.

Chapter 3

Appendix

3.1 Appendix Chapter 1

3.1.1 Data

Inputs

For imported products, we use customs data to retrieve information on products' import prices and quantities at the product level. We remove imports in bulk and calculate import prices per serving for those products where we can identify the number of units and volume. We assume products take 30 days to arrive from the border to the store and consider the import price as the main product cost for the upstream firm. When a price is missing during a month, we use the last available one. In most cases, we are able to match custom records to products in the output market using firm name, country of origin and products' descriptions. For those that we can't, we assign them the average import price from that country and firm to all products within the same group. Figure (3.1) contains the price of bottles at the firm level and average malt prices during the period we study.

Figure 3.1: Input prices

(a) Bottles Import Price

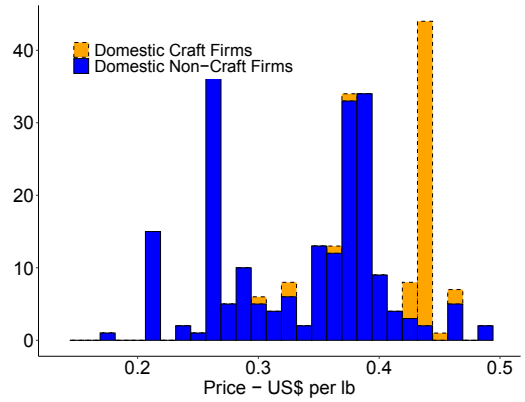

(b) Malt Price

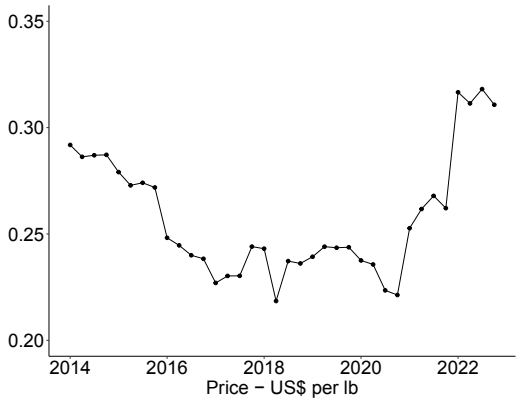

Craft Firms

We do not have balance sheet information for craft firms. Therefore, in order to overcome this limitation and construct a series of the capital stock, we exploit the fact that all capital goods that affect firms' production capacity such as automatic bottling machines, steel tanks and beer barrels, have to be imported from abroad. We build a capital stock index at the firm level that uses import prices and augments them conservatively for a factor of 1.2 to account for transport and installation costs. Then, over time we consider a depreciation rate of 5% per year. In Figure (3.2) we display the evolution of sales, the stock of capital and employed workers for craft firms.

Figure 3.2: Craft brewers

(a) Retail Sales (q) (b) Capital stock (c) Workers

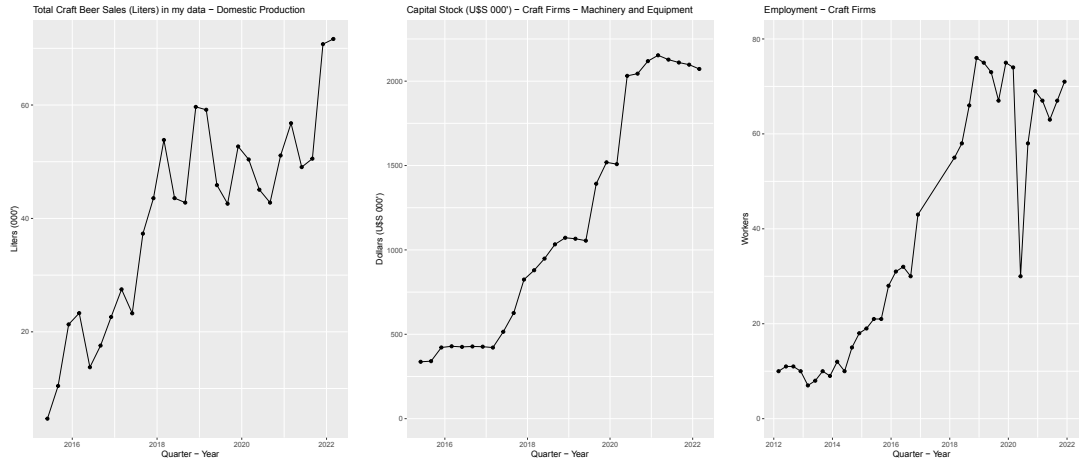

3.1.2 Estimation Results

Labor Costs - First Stage

We instrument labor costs ($lc_{\bar{g},t}$) at time t following the literature on production function estimation (Olley et al. 1996, Ackerberg et al. 2015). We aggregate domestic producers by firm type $\bar{g}$, depending on whether they produce craft or non-craft products. We assume that labor is a flexible input, so firms can adjust its quantity based on current period shocks. We use labor costs in the previous period ($lc_{\bar{g},t-1}$) and capital in the current period ($k_{\bar{g},t}$). We include time and group fixed effects. The model of the first stage is:

$$lc_{\bar{g},t} = \kappa_{lc} lc_{\bar{g},t-1} + \kappa_{\bar{g}} k_{\bar{g},t} + \sigma_t + \tau_{\bar{g}} + \nu_{\bar{g}t} \tag{3.1}$$

Estimation results using ordinary least squares are displayed in Table 3.1: Results have the expected signs, with labor costs (capital) positively (negatively) correlated with labor costs. We use the predicted values of $\hat{lc}_{\bar{g},t}$ from model IV in the estimation of demand and marginal costs.

Table 3.1: First Stage - Dep. Var. $lc_{\bar{g},t}$

	I	II	III	IV
$\hat{\kappa}_{lc}$	0.95***	0.59***	0.65***	0.47***
	(0.02)	(0.064)	(0.06)	(0.07)
$\hat{\kappa}_{\bar{g}}$	0.002	-0.014***	-0.012***	-0.0063***
	(0.0018)	(0.0032)	(0.003)	(0.0012)
FE				
Year	N	Y	Y	Y
Month	N	N	Y	Y
Firm	N	N	N	Y
Num.Obs.	164	164	164	164
R2	0.94	0.94	0.95	0.96

* p < 0.1, ** p < 0.05, *** p < 0.01

Demand Estimates - Alternative Market Definition

We estimate demand using a more granular market definition, exploiting regional and retail chain variations in sales. These extra dimensions of heterogeneity allow us to introduce cross-section income variation using the Household Survey conducted by the National Institute of Statistics. In this case, data is aggregated at the quarter level and a product is defined as a brand only. The estimates of this demand model are displayed in the following table. Importantly, the medians of market and own-product elasticities are similar to the estimates we obtain using a more aggregated model.

Variable	Parameter	Logit	RCL-1	RCL-2	RCNL-1
Price	α	-1.38 (0.09)	-4.75 (1.44)	-5.08 (1.54)	-2.85 (0.75)
RC	σ_0		1.49 (0.16)		
Nest	ρ_0				0.41 (0.09)
Demographic Interactions					
Income × Price	κ_α		1.46 (0.44)	1.51 (0.49)	0.86 (0.23)
Income × Constant	κ_0			1.46 (0.21)	0.54 (0.23)
Fixed Effects		(P,R,T,C)	(P,R,T,C)	(P,R,T,C)	(P,R,T,C)
Median Own Price El.		-2.39	-3.41	-3.38	-3.44
Median Market Price El.		-0.29	-0.44	-0.44	-0.34
Median Outside Div.		74.8%	51.2%	46.1%	38.1%

Models estimated by GMM, standard errors are shown in parentheses. Data was aggregated at the quarter level. Products should be available in at least 100 markets and more than 1 retail chain between 2015Q2 and 2022Q1. Prices per serving (355mL-12oz) are expressed in US$; Income per HH member is expressed in real UYU and converted to US$ 000'. Fixed Effects: (P) Product, (R) Region, (T) time and (C) Chain for P = 34, R = 6, T = 28 and C = 8. 15,930 observations of 721 markets.

3.1.3 Additional Results

Uniform Minimum Wage

The welfare consequences of a uniform minimum wage were described in Section (1.6). In Figure (3.3) we show results for average prices (wages) and total sales (employment) in the product (labor) market.

Figure 3.3: Labor and Product Market Equilibrium relative to Baseline

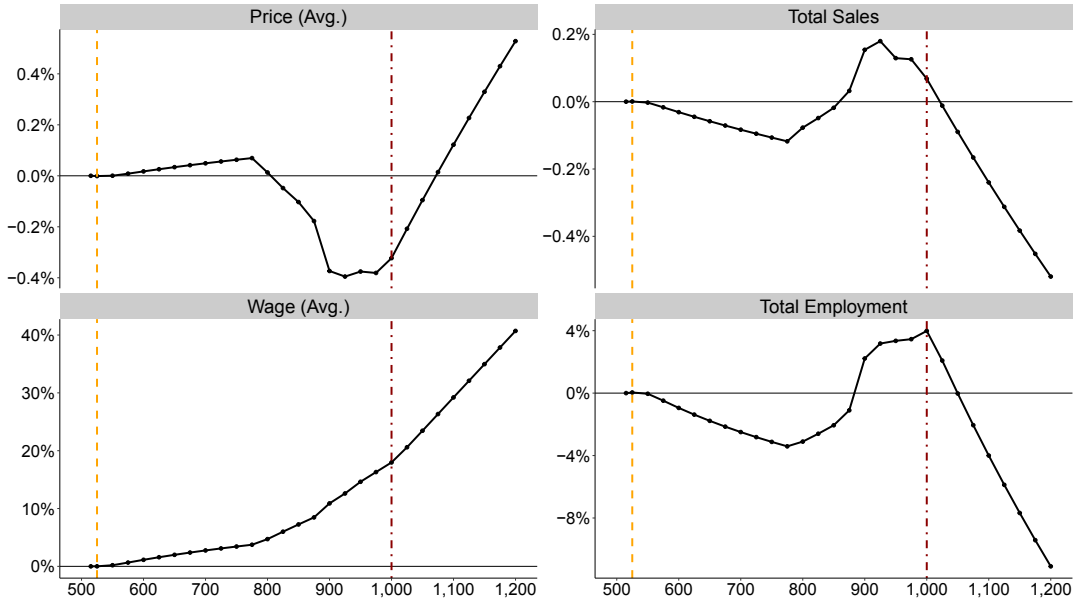

Table (3.2) contains counterfactual results by firm type and origin of the product.

Table 3.2: Simulation Results by Firm - Optimal Uniform Minimum Wage

Optimal Minimum Wage by Objective										
	Prices		Sales		Wages		Employment		WR	π
	$\Delta\%^{\dagger}$	US$ ‡	$\Delta\%^{\dagger}$	$s_{g(j),ins}$	$\Delta\%^{\dagger}$	US$ ‡	$\Delta\%^{\dagger}$	# L	US$ M	US$ M
Social Welfare - MW: US$ 1,000										
Craft Firms	5.1	2.56	-14.3	1.2	94.2	1,000	-43.7	9	0.2	0.46
Largest Firm	-0.6	1.18	0.5	87.8	11.4	2,059	10.8	121	6.1	47.6
Domestic Products	-0.5	1.06	1.8	61.9	-	-	-	-	-	33.1
Imported Products	1.5	1.62	-3.0	21.7	-	-	-	-	-	14.5
Importers	0.3	1.21	-0.1	15.1	-	-	-	-	-	2.9
Consumer Surplus - MW: US$ 550										
Craft Firms	-1.5	2.43	4.6	1.4	10.8	550	6.8	16	0.03	0.51
Largest Firm	-0.0	1.19	0.0	87.5	0.0	1,848	0.00	109	4.7	47.0
Domestic Products	-0.1	1.07	0.03	60.7	-	-	-	-	-	32.3
Imported Products	0.0	1.59	0.0	22.6	-	-	-	-	-	14.8
Importers	0.0	1.20	-0.0	15.3	-	-	-	-	-	2.9

Note: Welfare for markets we study in 2019/2021. † Variations ($\Delta\%$) relative to a baseline scenario with National Minimum Wage. ‡ Average price per serving, gross monthly wage. Inside share $s_{g(j),ins}$. Employment considers only the jobs needed to produce and serve the markets we study.

Figure 3.4: Price and Marginal Cost Variation - Domestic Products, Optimal MW vs No-MW

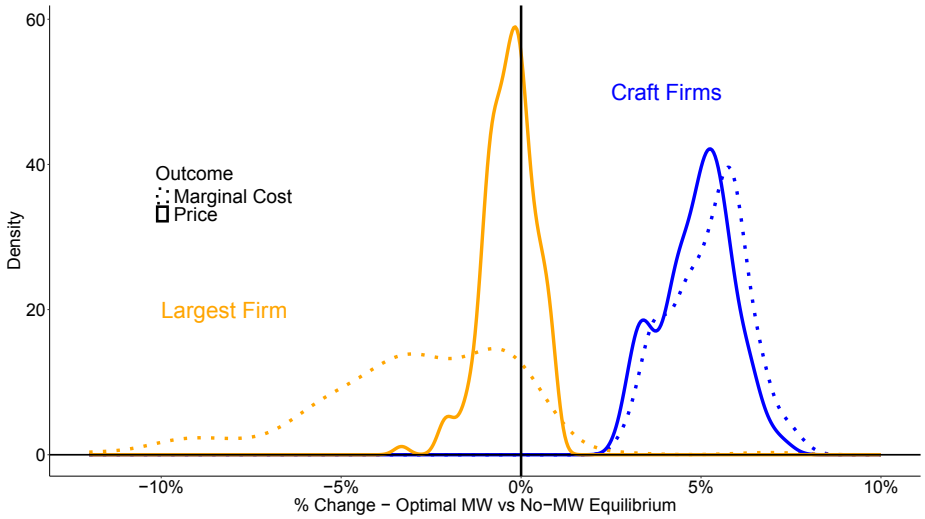

Note: kernel density estimates, marginal cost and price variation of domestic products by product and market, under optimal minimum wage relative to baseline equilibrium.

Uniform MW - Alternative Labor Supply Elasticity

The baseline specification assumes that, absent the minimum wage, the largest firm has more labor market than craft brewers. In this sense, we assumed $\theta_{NC} = 0.25$ and $\theta_C = 0.1$, consistent with labor supply elasticities of $\eta_{NC} = 4$ and $\eta_C = 10$. We tested the sensibility of our results to alternative parameterizations of the labor supply elasticity, in this case using common values for craft and non-craft domestic firms.

In Figure (3.5) we display consumer surplus variations relative to the equilibrium without minimum wage. Vertical lines represent the minimum wages that maximize social welfare in each case. Notice that, while consumer surplus is maximized in a range of US$ 550-600, social welfare is maximized at minimum wages that increase with the degree of labor market power. Again, efficiency gains in the largest firm are realized at higher minimum wages and that improves welfare for

this firm and workers. Consumers do not receive a large share of these gains, so jointly with craft brewers, they would be better off under lower minimum wages.

Figure 3.5: Consumer Surplus - Change relative to No-MW Equilibrium

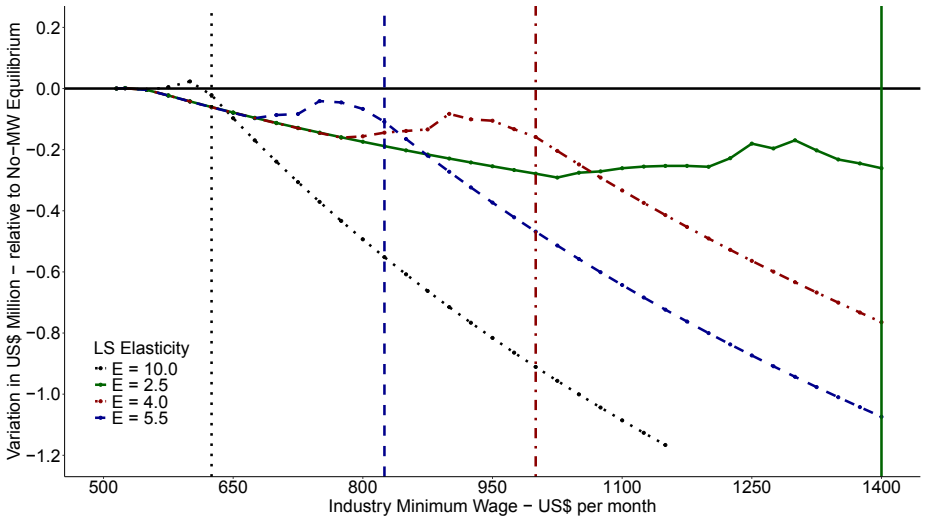

Uniform MW - Alternative Labor and Product Market Correlation

Table 3.3 contains the simulated changes in ownership, where we divest the property of the main products offered by the largest firm. In each case, we assign all imported products to a different firm and keep it constant in alternative simulations. The case we discuss in the counterfactual section is Domestic Oligopoly (A) with 3 large domestic firms. Results for other ownership scenarios are displayed in Figure (3.6), in all cases under the a common labor supply elasticity $\eta = \theta^{-1} = 4$.

Table 3.3: Simulated Ownership of Divested Products

# Domestic Firms →		# 1		# 2	# 3	# 4
Brand	Mkt. Share (Inside)	Baseline	Domestic Monopoly	Domestic Duopoly	Domestic Oligopoly (A)	Domestic Oligopoly (B)
HHI:		7,521	4,316	2,552	2,152	1,761
Brand 1	22.4%	1	1	1	1	1
Brand 2	15.6%	1	1	2	2	2
Brand 3	12.6%	1	1	2	2	3
Brand 4	8.9%	1	1	1	3	4
Domestic	59.5%					
Brand 5	11.6%	1	2	3	4	5
Brand 6	5.2%	1	2	3	4	5
Brand 7	4.8%	1	2	3	4	5
Brand 8	3.3%	1	2	3	4	5
Brand 9	1.1%	1	2	3	4	5
Rest	1.0%	1	2	3	4	5
Imported	27.0%					
Total	86.5%					

Figure 3.6: Welfare (US$ M) - Change relative to No-MW Equilibrium

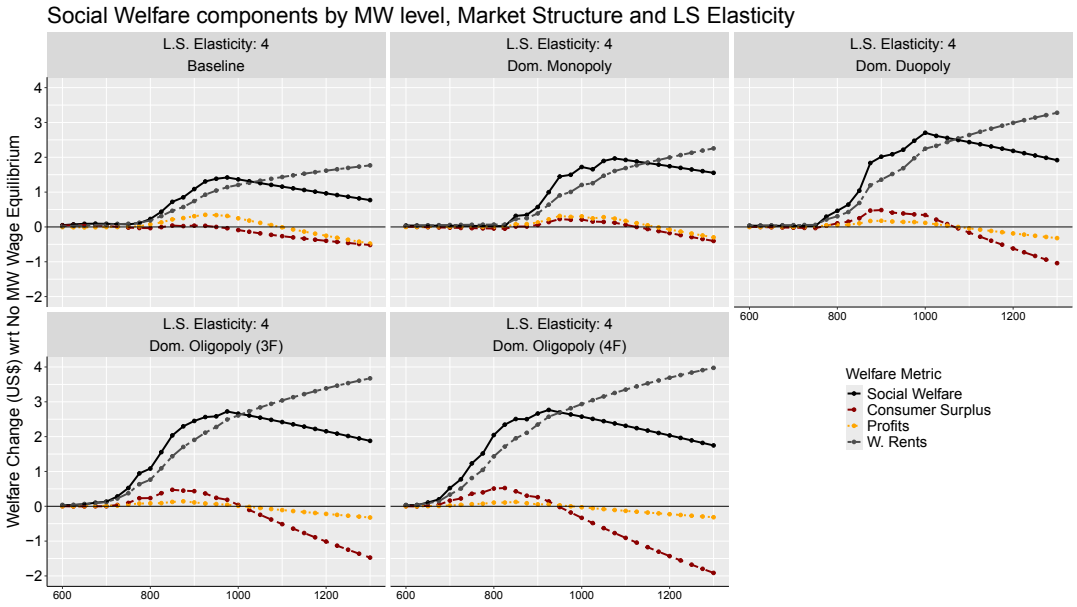

3.1.4 Second Choice Diversion Moments

We match aggregate second choice diversion moments using an unexpected removal of all products offered by the largest firm ($\mathcal{J}_{\text{large}}$) in a particular retailer between June 2020 and November 2021.

In Figure (3.7) we display the evolution of sales over time in retail chains affected and unaffected by the removal. In the case of affected chains (right panel), a large share of consumers previously choosing products of the largest firm, switched to products of competitors that had a minor share before. The diversion to the outside good we estimate is a consequence of the consequence that most consumers remained in the market. Figure (3.8) shows consumer substitution patterns in the affected retail chain. Although only non-craft products were removed, the consumption of craft products was unaffected. In the case of lager, most consumers remained buying lager after the event.

We follow the procedure developed by Conlon et al. 2024, exploiting the

Figure 3.7: Total Sales by Retail Chain

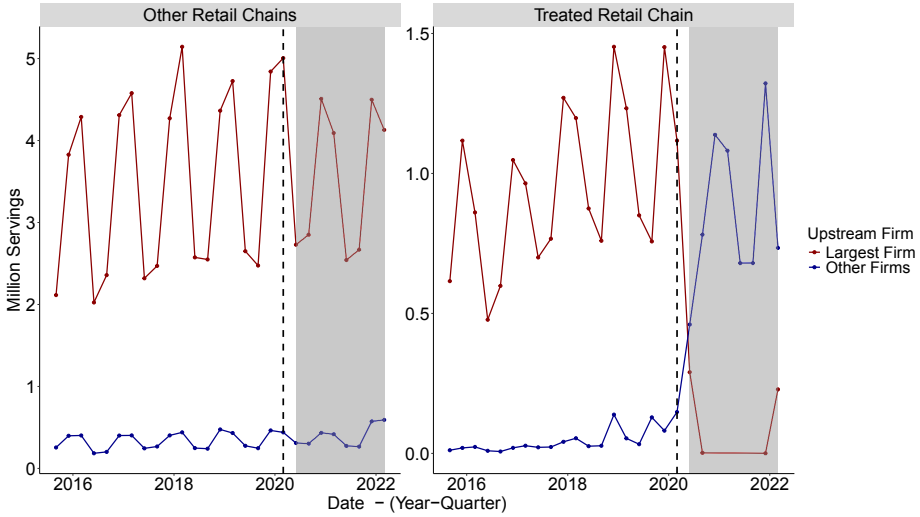

Figure 3.8: Sales by Characteristic, Treated Retail Chain

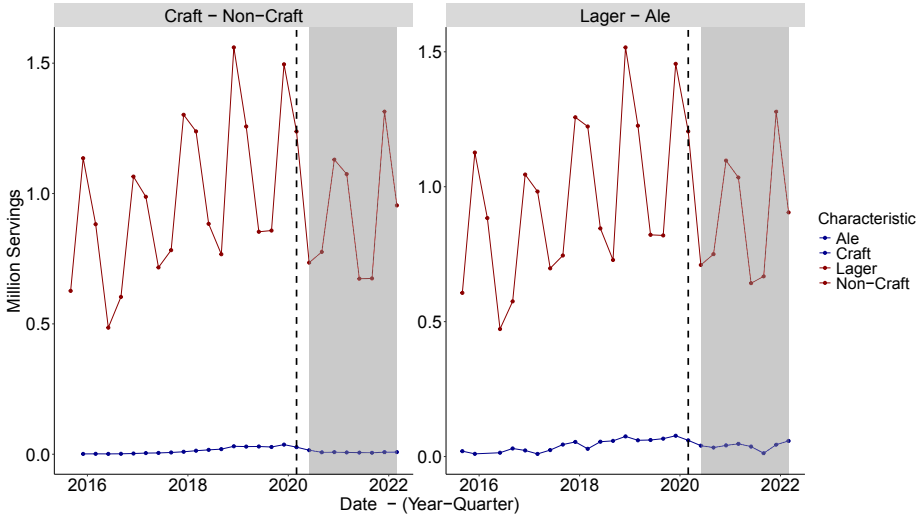

model assumption that consumers make mutually exclusive and exhaustive discrete choices. When product j is removed from the choice set, consumers previously choosing j (treated), can either select one of the remaining products $\mathcal{J} \setminus \{j\}$ or the outside good q_0. The procedure consists in selecting valid controls to compute raw diversion ratios, based on the fact that the level of sales of alternative products can neither decrease nor increase more than the volume of sales of the removed product j. Finally, we use a Bayesian procedure to obtain a posterior

distribution of the diversion ratio to each non-removed product k. The posterior restricts diversion ratios to be in the unit simplex and to sum one. In this case, the removal included more than one product ($\mathcal{J}_{rem} = \mathcal{J}_{\text{large}}$), but the estimation procedure is the same. For estimation purposes, we match the model-based diversion calculated using Equation (1.5) with the mean of the posterior distribution.

In practice, we match diversion ratios for the outside good and for a subset of relevant products spanning multiple characteristics. Table (3.4) contains the mean of the posterior distribution of the non-parametric diversion ratios.

Table 3.4: Diversion Ratio - Posterior Distribution - Mean

Product	Div. Ratio	Lager	Imported	Craft
Outside Good	0.38			
Product 1	0.20	✓	✓	
Product 2	0.18	✓	✓	
Product 3	0.09	✓	✓	
Product 4	0.05	✓	✓	
Product 5	0.006	✓	✓	
Product 6	0.005	✓	✓	
Product 7	0.004	✓	✓	
Product 8	0.002	✓	✓	
Product 9	0.002	✓		✓
Product 10	0.001			✓

3.1.5 Counterfactual - Implementation details

Simulation procedure

We simulate labor and product market equilibrium for every minimum wage ($\underline{w}$) set by the policymaker. From our description of the problem of the firm (1.4.3), in equilibrium the prices that firms set in the product market and the wages they post in the labor market are consistent with each other. Consistency implies that there is no profitable deviation in any market: the vector of market prices is an equilibrium in the product market given a vector of marginal costs, and the vector of wages associated with those marginal costs is an equilibrium in the labor

market. At the vector of posted wages, firms are able to hire the workers they need to produce and satisfy the demand they receive from consumers at the vector of equilibrium prices. We use the estimated model and the calibrated labor supply curve to search for the equilibrium vector of prices (p_t) that would arise in each case, and then recover wages, employment, production, and sales. Conservatively, we assume that remaining retail markets m' are not affected, so we focus only on the market where we estimated demand.

The core of our counterfactual analysis consists in changing minimum wage levels, and this affects (in most cases) relative input prices. Therefore, to isolate the inefficiencies we study (labor and product market power) from input misallocation, in our counterfactual, we allow firms to adjust their capital-labor ratio before engaging in short-run price competition when wages change. We proceed as follows: First, we assume that the capital-labor ratio in the equilibrium we observe is the optimal (short and long run marginal cost coincide), and that the technology firms use has constant returns to scale ($\hat{\beta}_K = 1 - \hat{\beta}_L$). Using these assumptions and information on capital stocks, we recover firm-specific implicit capital prices $r_{gm(t)}$. Then, similar to the case of other inputs (bottles, hops), we assume that firms can buy and sell all the capital they want at that price when the equilibrium wage changes.

We operationalize the simulation with a fixed point iteration over product and labor markets' equilibrium. When the minimum wage binds on firms' labor demand, the marginal cost is constant and the simulation is standard. When the MW does not bind, the marginal cost is increasing in production due to the upward sloping labor supply curve[1], and the fixed point solves for the intersection of marginal revenue and marginal cost. Finally, when the minimum wage binds on the labor supply, we solve for the shadow value of the hiring constraint (ref.

[1]This holds even under the assumption of constant returns to scale

Appendix) that makes the level of employment and production consistent with marginal revenue and sales[2]. To search for an equilibrium, we use a fixed-point iteration with a tolerance of 10^{-9} for the maximum difference between shares in successive iterations, using the concept of consistency defined above. We simulate equilibrium for all markets we study in the period 2019-2021, using a grid of gross minimum wages in the range [525, 1,500] in intervals of US$ 25[3].

[2]See (1.4.3)
[3]Equivalent to 13 cents in the gross hourly wage of a full time worker.

3.2 Appendix Chapter 2

3.2.1 Gross Wage Distribution for Contributions

Figure 3.9: Gross Wage Distribution for Social Security Contributions

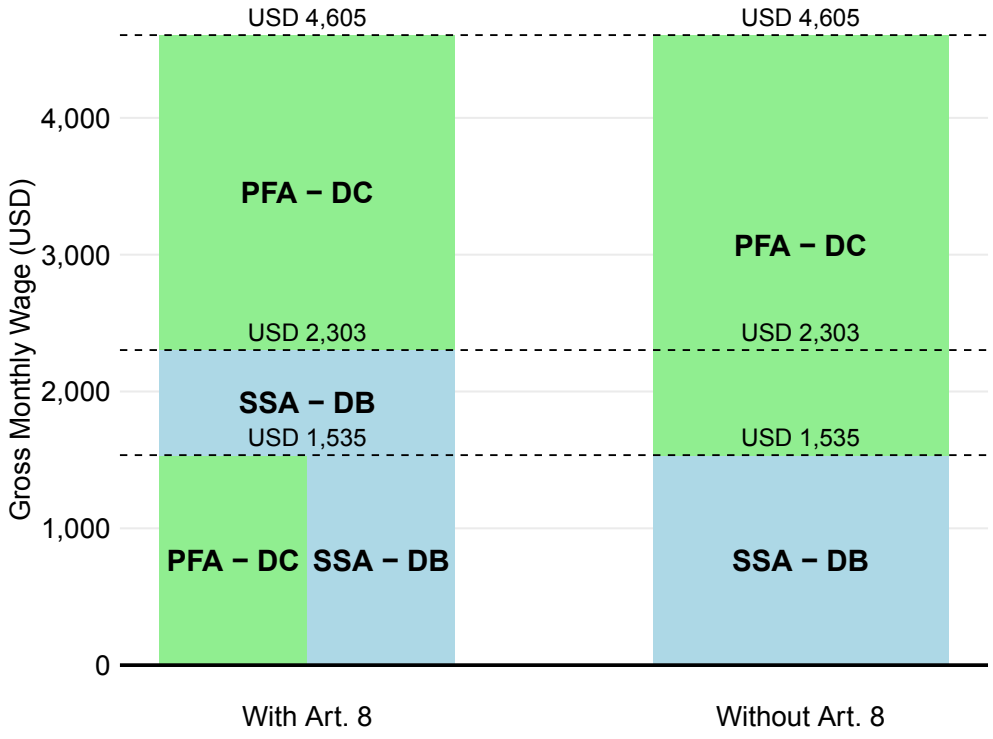

Note. PFA-DC: Pension Fund Administrator - Defined Contribution subsystem, SSA-DB: Social Security Administration - Defined Benefit subsystem. Expressed in real US Dollars of April 2021. Thresholds are adjusted yearly according to a Nominal Average Wage Index. Contributions to each subsystem are calculated as 15% of the gross wage allocated to each subsystem.

3.2.2 Additional Descriptive Information

Market Shares

Table 3.5: Market Share by Firm, Gross Wage Tertile and Period - New Enrollees

	PF 1	PF 2	SOE	PF 3
Period 2002-2005				
Optional T.1	0.21	0.27	0.19	0.32
Optional T.2	0.18	0.25	0.31	0.27
Optional T.3	0.17	0.17	0.48	0.18
Mandatory	0.12	0.10	0.64	0.14
Period 2014-2017				
Optional T.1	0.25	0.28	0.21	0.26
Optional T.2	0.21	0.23	0.26	0.30
Optional T.3	0.16	0.22	0.36	0.27
Mandatory	0.10	0.15	0.57	0.18
Period 18-20				
Optional T.1	0.24	0.24	0.23	0.28
Optional T.2	0.20	0.22	0.24	0.34
Optional T.3	0.21	0.17	0.33	0.29
Mandatory	0.08	0.02	0.85	0.05

Notes. For each level of gross wages we display the average share of new enrollees by PFA. The three tertiles are composed of workers whose wages are below the compulsory enrollment threshold. Optional stands for enrollees with gross wages above the mandatory enrollment threshold.

Figure 3.10: Market Share by Firm - All Enrollees

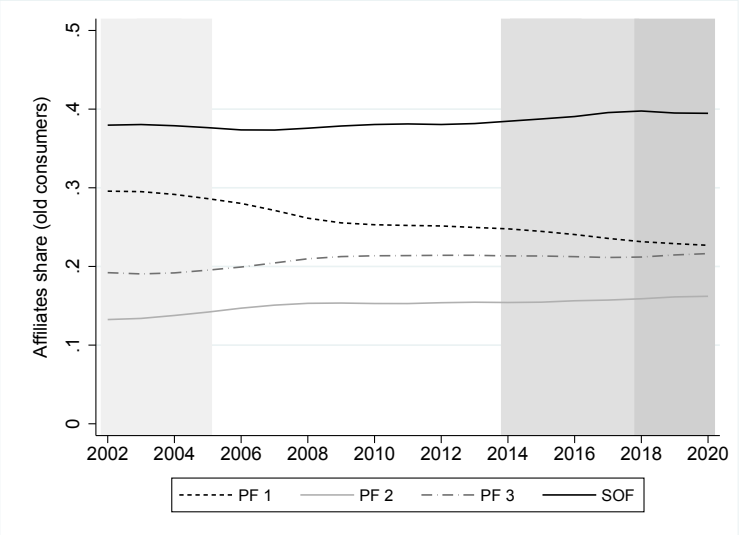

Sales Force

Table 3.6: Sales Force Agents - Share by Firm and Period

Period	Firm			
	PF 1	PF 2	SOE	PF 3
2002-2005	0.21	0.24	0.35	0.21
2014-2017	0.22	0.15	0.36	0.27
2018-2020	0.24	0.15	0.36	0.25

Notes. Average share of total.

Entry-Enrollment Spell

Figure 3.11: Labor Market Entry and PFA Enrollment

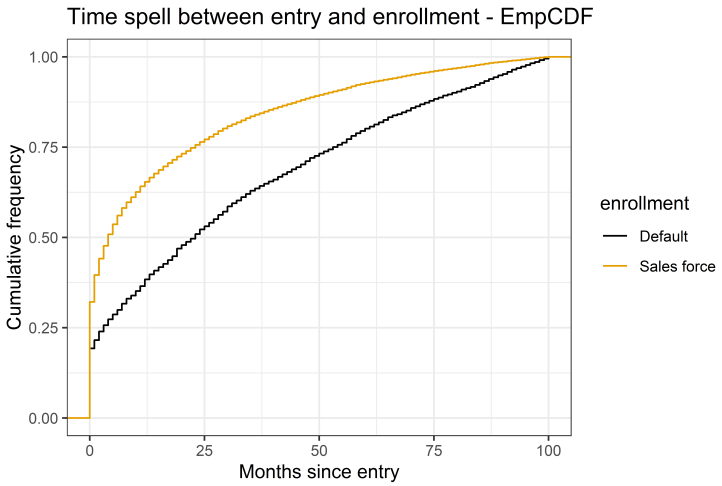

Note. Yellow curve: empirical cumulative distribution function of time spell between entry and enrollment for individuals voluntarily enrollment. Black curve: cumulative distribution function of time spell between entry to the labor market and enrollment for individuals enrolled by default (mandatory enrollment).

Switching Rates

Figure 3.12: Switching Patterns

(a) Equilibrium 2014-17, Before cap

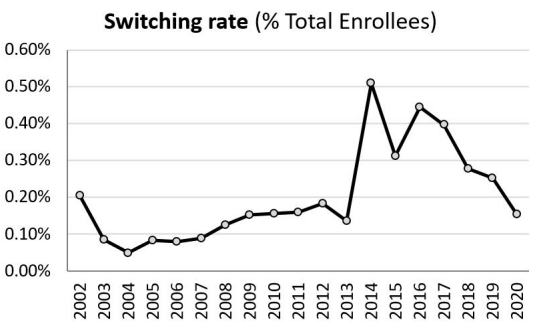

(b) Equilibrium 2020, After cap

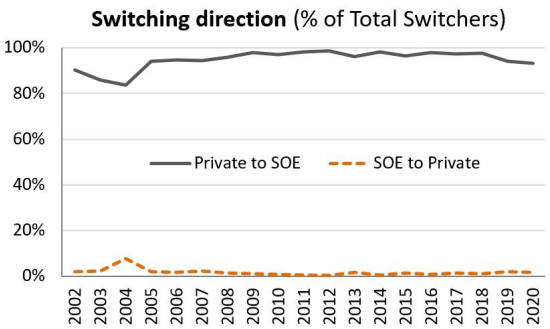

Expected Savings

Based on (1) the low switching rates, (2) the almost unanimous switching direction from private firms to the public option, (3) the lower fees charged by the public option relative to private firms, and (4) the superior investment performance of some private firms, we study the expected level of savings under the two most common enrollment strategies: (1) single-firm enrollment, where the worker remains with the same PFA until retirement, and (2) switch-once, where the worker initially enrolls in the low-fee public option and later switches to a higher-return private firm.

The expected savings at retirement under the first strategy are shown in Figure 3.13, comparing outcomes before the introduction of the cap on management fees (3.13a) and after (3.13b). In both cases, we compare enrollment in the public option to enrollment in the private firm with the highest observed returns in our sample.

Figure 3.13: Expected Savings by PFA and Period, Public Option vs Best Private Firm

(a) Equilibrium 2014-17, Before cap (b) Equilibrium 2020, After cap

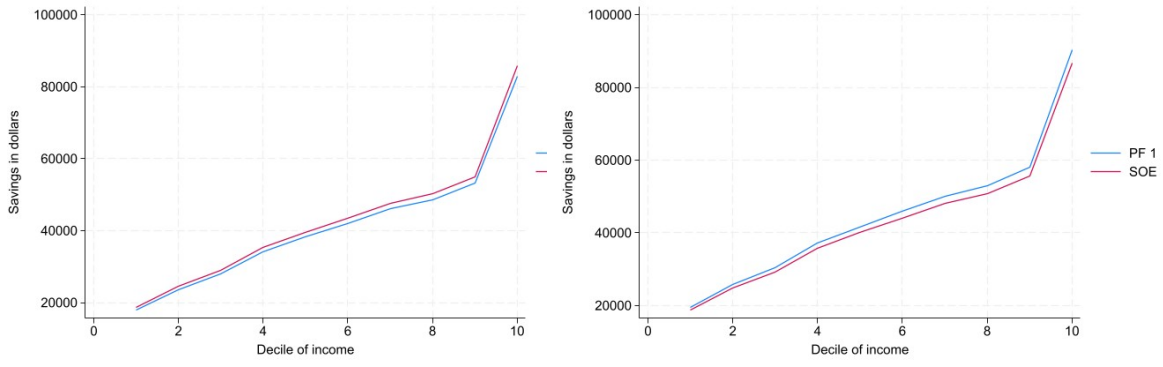

Importantly, before the introduction of the cap, differences in returns were

not sufficient to offset the fee advantage of the public option, regardless of the worker's gross wage level. These results suggest that enrolling in a private firm was a "mistake" under the 2014–2017 configuration. However, after the cap was introduced, this result changes: the best-performing private firm is now able to compensate for the higher fee it charges by delivering superior returns.

Second, given that most workers switch firms at most once, we explore an alternative switch-once strategy in which the worker begins with the low-fee public option and switches later to a high-return private firm. The intuition behind this strategy is that early in the contribution period—when the accumulated savings are still low—differences in fees matter more than differences in returns for growing the savings stock. As the savings accumulate, the return on investment becomes more influential, making a switch to a higher-return private firm potentially beneficial.

We find that, before the introduction of the cap on management fees, the optimal time to switch was approximately 28 years after entering the system. After the cap was introduced, this optimal switching point dropped to 9 years for the average worker. These results are illustrated in Figure 3.14.

Figure 3.14: Optimal Switching Timing: Public Option to Best Private Firm

(a) Equilibrium 2014-17, Before cap

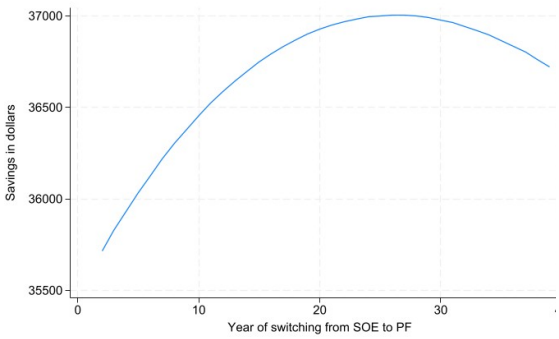

(b) Equilibrium 2020, After cap

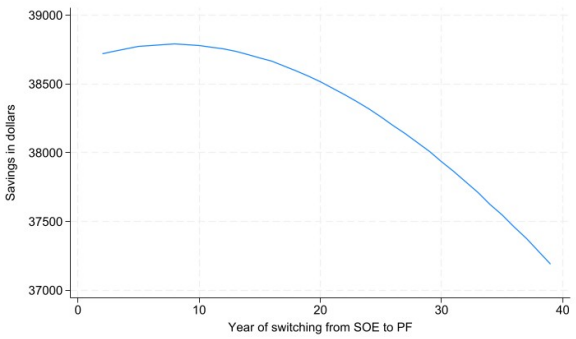

3.2.3 Model

Labor Market

We use a stationary model of the labor market where the mass of wages of new workers is a fixed fraction α of the aggregate wage mass M_t, and the probability of retirement is deterministic and constant. We make the following assumptions:

1 Retirement probability: Workers remain active in the labor market for forty years and then retire. Accordingly, we assume a constant retirement probability of 0 for the first forty periods and 1 thereafter.

2 Workers and firms consider the average gross wage over the life cycle, w_{it}, when solving their optimization problems[4].

3 Each year t, a new cohort of equal size, preferences, and gross wages enters the labor market.

4 The wage mass M_t is constant each period $M_t = M$. Therefore, the wage mass of the new cohort is equal to the wage mass of the cohort who retires.

Revenues

In each period there are new and old workers who earn gross wages w_{it}^n and w_{it}^o, respectively. The total wage mass M_t aggregates over the distribution of each group:

$$M_t = \int_{i_n} w_{it}^n dF_{i_n} + \int_{i_o} w_{it}^o dF_{i_o} = M_{t,n} + M_{t,o} \tag{3.2}$$

As we assume cohorts are of equal size, we write the wage mass of new and old cohorts of workers as $M_{t,n} = \alpha M_t$ and $M_{t,o} = (1-\alpha)M_t$. Then, we can express

[4]See Appendix 3.2.4 for details on the estimation procedure and results.

firm expected revenues (before switchers are taken into account) as:

$$\mathbb{E}[Y_{jt}] = \int_{i_n} pr_{ijt}(w^n_{it}, \mathrm{f}_t, \boldsymbol{\mu}_t)\; w^n_{it}\; dF_{i_n} + \int_{i_o} pr_{ijt}(w^o_{it}, \mathrm{f}_t, \boldsymbol{\mu}_t)\; w^o_{it}\; dF_{i_o} \tag{3.3}$$

$$= \alpha M_t \underbrace{\int_{i_n} pr_{ijt}(w^n_{it}, \mathrm{f}_t, \boldsymbol{\mu}_t) \frac{w^n_{it}}{\alpha M_t} dF_{i_n}}_{s^n_{jt}(\mathrm{f}_t, \boldsymbol{\mu}_t)} + (1-\alpha) M_t \underbrace{\int_{i_o} pr_{ijt}(w^o_{it}, \mathrm{f}_t, \boldsymbol{\mu}_t) \frac{w^o_{it}}{(1-\alpha) M_t} dF_{i_o}}_{s^o_{jt}(\mathrm{f}_t, \boldsymbol{\mu}_t)} \tag{3.4}$$

This equation is consistent with the revenue equation 2.6 and show the re-weighting of choice probabilities. Under our labor market assumptions, we re-express these probabilities as revenue shares, with higher-wage enrollees being relatively more valuable to firms than lower-wage ones. Note also that we use the same probability for new and existing enrollees, consistent with our assumption of equal preferences across cohorts (conditional on a life-cycle wage level).

SOE Revenues

In the baseline model, the SOE receives revenues not only from its original enrollees (new and old) but also from old enrollees switching from private PFAs. For this reason, we adjust Equation 2.6 as follows:

$$\mathbb{E}\big[Y_{SOE,t}\big] = f_{SOE,t} \times \Bigg[\underbrace{\Big(\underbrace{\alpha\, M_t\, s^n_{SOE,t}(\mathrm{f}_t, \boldsymbol{\mu}_t)}_{\text{New Workers' Wages}} + \underbrace{(1-\alpha)\, M_t\, s^o_{SOE,t}(\mathrm{f}_t, \boldsymbol{\mu}_t)}_{\text{Old Workers' Wages}} \Big)}_{\text{Original Enrollees}} \tag{3.5}$$

$$+ \underbrace{\sum_{j_{PF}} (1-\alpha)\, M_t\, s^o_{jt}(\mathrm{f}_t, \boldsymbol{\mu}_t)\, aw_{jt}(\mathrm{f}_t, \boldsymbol{\mu}_t)}_{\text{Incoming Switchers}} \Bigg]$$

Importantly, notice that we assume that no one is leaving the SOE by removing the awareness probability from the old-workers' term, and that it captures the all

enrollees leaving the group of private firms $\mathcal{J}_{PF}$.

Expected Capitalization Cost

For period t, the expected capitalization cost of firm j depends on the Pension Savings Fund (PSF_{jt})[5], and on the realization of firm-specific and average market returns R_{jt} and $\bar{R}_t$. Even though the PSF_jt also is a function of R_{jt}, for simplicity we assume that the capitalization cost is based on the stock before R_{jt} realizes.

By law, the minimum rate of return $r_{\min,t}$ that pension fund administrators must secure their enrollees is the minimum between: 1) 2%, and 2) $\bar{R}_t - 2\%$[6]. To calculate the expected cost we work with the support of $r_{\min}$:

$$\mathbb{E}[\text{Cap. Cost}]_{jt} = \int_{\text{supp}(r_{\min})} \left(\int_{-\inf}^{r_{\min,t}} \text{PSF}_{jt} \cdot (r_{\min,t} - R_{jt}) \cdot f(R_{jt}, \bar{R}_t) d\text{R} \right) d\text{F}(r_{\min}) \quad (3.6)$$

Notice that when $\bar{R} \geq 4\%$, the minimum rate is 2%, while when $\bar{R} < 4\%$, the minimum rate is $\bar{R} - 2\%$. Then, the expected capitalization cost of firm j in period

[5]This represents the value of the assets under management.

[6]$\bar{R}_t$ is the average of R_{jt} across PFAs', weighted by a vector ω that sum to 1 and depends on the PSF_{jt} of each firm: $\bar{R}_t = \sum_j \omega_{jt} R_{jt}$ with $\omega_{jt} = \frac{PSF_{jt}}{\sum_j PSF_{jt}}$

t as a function of the joint distribution of $\bar{R}$ and R_j is:

$$\begin{aligned}
\mathbb{E}[\text{CC}]_{jt} &= \Bigg(\underbrace{\int_{-\inf}^{4\%} \int_{-\inf}^{\bar{R}-2\%} \text{PSF}_{jt} \cdot (\bar{R}_t - 2\% - R_{jt}) \cdot f(R_j, \bar{R}) \text{dRd}\bar{R}}_{E(\text{Cap Cost}|\bar{R} \leq 4\%)_{jt}} \Bigg) \cdot F_{\bar{R}}(4\%) \\
&+ \Bigg(\underbrace{\int_{4\%}^{\inf} \int_{-\inf}^{2\%} \text{PSF}_{jt} \cdot (2\% - R_{jt}) \cdot f(R_j, \bar{R}) \text{dRd}\bar{R}}_{E(\text{Cap Cost}|\bar{R} \geq 4\%)_{jt}} \Bigg) \cdot (1 - F_{\bar{R}}(4\%)) \\
&= \Bigg(\int_{-\inf}^{4\%} \int_{-\inf}^{\bar{R}-2\%} \text{PSF}_{jt} \cdot (\bar{R}_t - 2\% - R_{jt}) \cdot f(R_j, \bar{R}) \text{dRd}\bar{R} \Bigg) \cdot F_{\bar{R}}(4\%) \\
&+ \Bigg(\int_{4\%}^{\inf} \int_{-\inf}^{2\%} \text{PSF}_{jt} \cdot (2\% - R_{jt}) \cdot f(R_j, \bar{R}) \text{dRd}\bar{R} \Bigg) \cdot (1 - F_{\bar{R}}(4\%))
\end{aligned}$$

To work with this term, we assume that R_j follows a Normal distribution with $R_j \sim \mathcal{N}(\mu_j, \sigma_j^2)$ and therefore, the average return also follows a normal distribution $\bar{R} \sim \mathcal{N}(\bar{\mu}, \bar{\sigma}^2)$ with $\bar{\mu} = \sum_j \omega_j \mu_j$ and $\bar{\sigma}^2 = \sum_j \sum_j \omega_j \omega_k \sigma_{jk} = \omega' \Sigma \omega$. Furthermore, given that both R_j and $\bar{R}$ are Normal random variables, the joint distribution $f(R_j, \bar{R})$ has the following expression:

$$\begin{aligned}
f(R_j, \bar{R}) &= \frac{1}{2\pi \sigma_{R_j} \sigma_{\bar{R}} \sqrt{1-\rho^2}} \\
&\cdot \exp\left\{ -\frac{1}{2(1-\rho^2)} \left[\left(\frac{R_{jt} - \mu_j}{\sigma_{R_j}} \right)^2 + \left(\frac{\bar{R}_t - \mu_{\bar{R}}}{\sigma_{\bar{R}}} \right)^2 - 2\rho \frac{\left(R_{jt} - \mu_{R_j}\right)\left(\bar{R}_t - \mu_{\bar{R}}\right)}{\sigma_{R_j} \sigma_{\bar{R}}} \right] \right\}
\end{aligned} \tag{3.7}$$

3.2.4 Estimation

Life Cycle Gross Wage

Our database contains payroll data between 1996 and 2020 for a random sample of 125,000 workers, drawn from the universe of active workers. We observe each

worker's monthly wage during this period. To construct an estimate of workers' wages over their life cycle, we proceed as follows:

1 We calculate an earnings curve over the working life cycle for the economy by regressing annual wages on age and age squared. Because our sample includes workers at different career stages, it captures individuals nearing retirement in 1996 as well as new entrants to the labor market.

2 We apply this earnings curve to each worker used to estimate demand — specifically, those who entered the labor market for the first time between 2002 and 2020— using their observed salary in the year of entry.

3 We compute the average wage for each worker over their working life cycle.

1^{st} Stage - Estimation Results

Table 3.7: Switching Probability

	Y_{jt} = Switchers (% Enrollees)
$\hat{\beta}_{\Delta\mu}$	-5.893***
	(0.473)
$\hat{\beta}_{\Delta\text{fee}}$	52.02***
	(2.671)
Post_{2013}	0.374***
	(0.00515)
FE PF 3	0.0649***
	(0.00620)
FE PF 1	-0.900***
	(0.0334)
FE PF2	0.0528***
	(0.00581)
Constant	-3.046***
	(0.00786)

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

2^{nd} Stage - Additional Results

Figure 3.15: Median Management Cost Elasticity by Bracket and Period

(a) Period 2002-05 (b) Period 2014-17

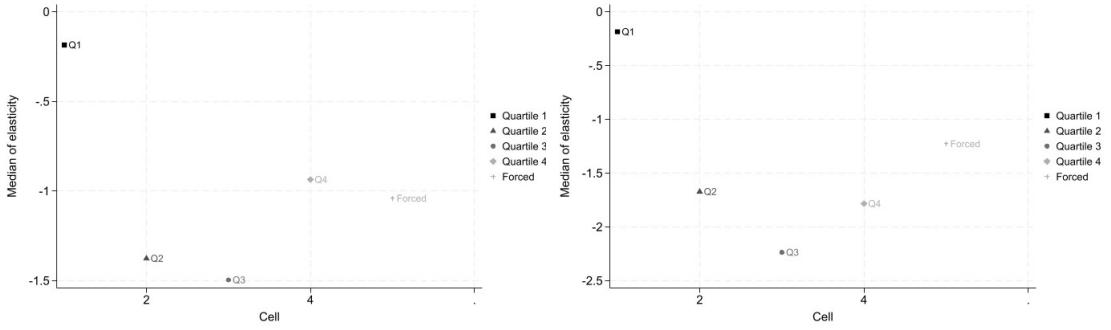

Note. Within-bracket median elasticity with respect to the management fee by period. Elasticities are calculated using observed fees, mean returns, and gross wages.

Enrollment Probability and Income

The results described in Section 2.5.2 are consistent with the fact that higher-wage workers enroll with higher probability in the public option. This explains why despite having a market share of enrollees of approximately 40%, the SOE receives close to 55% of monthly gross contributions in the subsystem. In Figure 3.16 we display the predicted enrollment probability to the public option by gross wage decile, as implied by the model.

Figure 3.16: SOE Enrollment Probability by Wage Decile

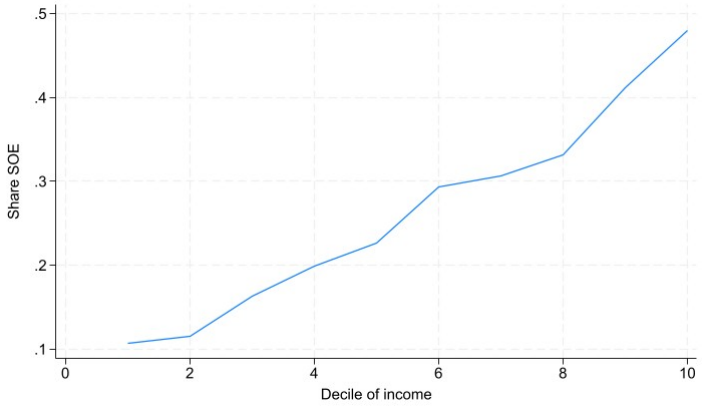

Note. Median probability by gross wage decile as predicted by the model for the period 2014-17.

3.2.5 Counterfactual

(Lack of) Entry Incentives

A new pension fund administrator entering the market under conditions similar to those observed historically would take approximately 12 years to achieve a positive return on equity (ROE), and around 19 years to reach an ROE of 50%—a level comparable to that of incumbent firms. This slow trajectory is primarily driven by high fixed operating costs. Investment management requires a similar organizational scale whether the fund is small or large, while sales force and marketing expenses are aimed at acquiring new customers and are largely unaffected by the number of current enrollees.

Unlike the early days of the system —when all firms competed to enroll a large, untapped pool of existing workers— new entrants today can only target new cohorts of labor market entrants, making the growth of a sufficiently large

customer base slower and more costly than during the system's inception.

Figure 3.17: Entry

(a) Revenues, Costs and Equity

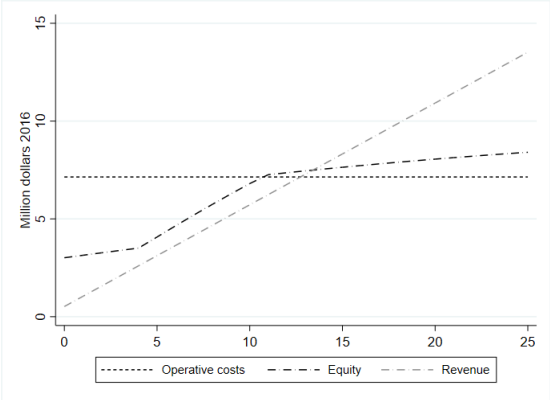

(b) Return-Over-Equity

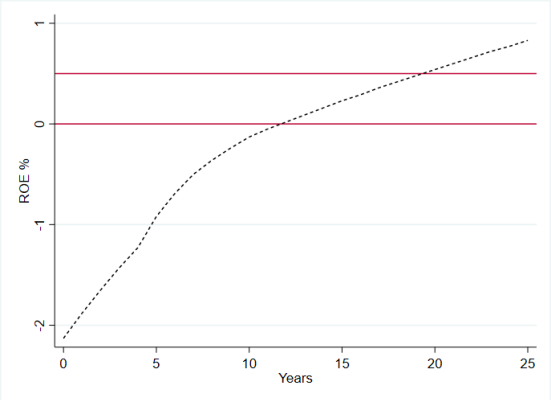

Bibliography

Aaronson, Daniel, Eric French, Isaac Sorkin, and Ted To (2018). "Industry dynamics and the minimum wage: a putty-clay approach". In: International Economic Review 59.1, pp. 51–84.

Ackerberg, Daniel A, Kevin Caves, and Garth Frazer (2015). "Identification properties of recent production function estimators". In: Econometrica 83.6, pp. 2411–2451.

Amodio, Francesco and Nicolas De Roux (2021). "Labor Market Power in Developing Countries: Evidence from Colombian Plants". In.

Atal, Juan Pablo, José Ignacio Cuesta, Felipe González, and Cristóbal Otero (2021). "The Economics of the Public Option: Evidence from Local Pharmaceutical Markets". In: Available at SSRN.

Avignon, Rémi and Etienne Guigue (2022). Markups and markdowns in the french dairy market. Tech. rep. Technical report, mimeo.

Beggs, Alan and Paul Klemperer (1992). "Multi-period competition with switching costs". In: Econometrica: Journal of the Econometric Society, pp. 651–666.

Bello, Salvatore Lo and Lorenzo Pesaresi (2022). "Equilibrium Effects of the Minimum Wage: The Role of Product Market Power". In.

Berger, David, Kyle Herkenhoff, and Simon Mongey (2022a). "Labor market power". In: American Economic Review 112.4, pp. 1147–93.

Berger, David, Kyle Herkenhoff, and Simon Mongey (2022b). Minimum wages, efficiency and welfare. Tech. rep. National Bureau of Economic Research.

Berry, Steven, James Levinsohn, and Ariel Pakes (1995). "Automobile prices in market equilibrium". In: Econometrica: Journal of the Econometric Society, pp. 841–890.

Berry, Steven T (1994). "Estimating discrete-choice models of product differentiation". In: The RAND Journal of Economics, pp. 242–262.

Braun, Sebastian (2011). "Unionisation structures, productivity and firm performance: New insights from a heterogeneous firm model". In: Labour Economics 18.1, pp. 120–129.

Brenkers, Randy and Frank Verboven (2006). "Liberalizing a distribution system: The European car market". In: Journal of the European Economic Association 4.1, pp. 216–251.

Bulow, Jeremy I, John D Geanakoplos, and Paul D Klemperer (1985). "Multimarket oligopoly: Strategic substitutes and complements". In: Journal of Political economy 93.3, pp. 488–511.

Busso, Matias and Sebastian Galiani (2019). "The causal effect of competition on prices and quality: Evidence from a field experiment". In: American Economic Journal: Applied Economics 11.1, pp. 33–56.

Card, David and Ana Cardoso (2022). "Wage flexibility under sectoral bargaining". In: Journal of the European Economic Association 20.5, pp. 2013–2061.

Card, David and Alan B Krueger (1993). Minimum wages and employment: A case study of the fast food industry in New Jersey and Pennsylvania.

Card, David, Ana Cardoso, Joerg Heining, and Patrick Kline (2018). "Firms and labor market inequality: Evidence and some theory". In: Journal of Labor Economics 36.S1, S13–S70.

Cardell, N Scott (1997). “Variance components structures for the extreme-value and logistic distributions with application to models of heterogeneity”. In: Econometric Theory 13.2, pp. 185–213.

Casacuberta, Carlos and Néstor Gandelman (2023). “Wage councils, product markups and wage markdowns: Evidence from Uruguay”. In: International Journal of Industrial Organization 87, p. 102916. issn: 0167-7187.

Chen, Yongmin and Michael H Riordan (2008). “Price-increasing competition”. In: The RAND Journal of Economics 39.4, pp. 1042–1058.

Conlon, Christopher and Nirupama Rao (2020). “Discrete prices and the incidence and efficiency of excise taxes”. In: American Economic Journal: Economic Policy 12.4, pp. 111–143.

Conlon, Christopher, Rodrigo González Valdenegro, Julie Holland Mortimer, and Paul Sarkis (2024). “Effects of Product Availability: Experimental Evidence”. In: Work in Progress.

Conlon, Christopher and Jeff Gortmaker (2020). “Best practices for differentiated products demand estimation with pyblp”. In: The RAND Journal of Economics 51.4, pp. 1108–1161.

Cunha, Jesse M, Giacomo De Giorgi, and Seema Jayachandran (2019). “The price effects of cash versus in-kind transfers”. In: The Review of Economic Studies 86.1, pp. 240–281.

Curto, Vilsa, Liran Einav, Amy Finkelstein, Jonathan Levin, and Jay Bhattacharya (2019). “Health care spending and utilization in public and private Medicare”. In: American Economic Journal: Applied Economics 11.2, pp. 302–32.

De Loecker, Jan and Paul Scott (2022). “Markup estimation using production and demand data. an application to the us brewing industry”. In: Unpublished paper.

DeLoecker, Jan and Chad Syverson (2021). “An industrial organization perspective on productivity”. In: Handbook of industrial organization. Vol. 4. Elsevier, pp. 141–223.

Dodini, Samuel, Anna Stansbury, and Alexander Willén (2023). How Do Firms Respond to Unions? Tech. rep. Norwegian School of Economics, Department of Economics.

Dube, Arindrajit (2019). “Impacts of minimum wages: review of the international evidence”. In: Independent Report. UK Government Publication, pp. 268–304.

Duggan, Mark and Fiona M Scott Morton (2006). “The distortionary effects of government procurement: evidence from Medicaid prescription drug purchasing”. In: The Quarterly Journal of Economics 121.1, pp. 1–30.

Dustmann, Christian, Attila Lindner, Uta Schönberg, Matthias Umkehrer, and Philipp Vom Berge (2022). “Reallocation effects of the minimum wage”. In: The Quarterly Journal of Economics 137.1, pp. 267–328.

Fan, Ying and Chenyu Yang (2020). Merger, Product Variety and Firm Entry: the Retail Craft Beer Market in California. Available at SSRN 3681556.

Farrell, Joseph and Paul Klemperer (2007). “Coordination and lock-in: Competition with switching costs and network effects”. In: Handbook of industrial organization 3, pp. 1967–2072.

Fonseca, Julia and Adrien Matray (2022). The real effects of banking the poor: Evidence from brazil. Tech. rep. National Bureau of Economic Research.

Friedman, Milton (1962). “Capitalism and Freedom: The Relation Between Economic Freedom and Political Freedom,„”. In: Capitalism and Freedom, pp. 7–17.

Grieco, Paul, Joris Pinkse, and Margaret Slade (2018). "Brewed in North America: Mergers, marginal costs, and efficiency". In: International Journal of Industrial Organization 59, pp. 24–65.

Handbury, Jessie and Sarah Moshary (2021). School food policy affects everyone: Retail responses to the national school lunch program. Tech. rep. National Bureau of Economic Research.

Harasztosi, Péter and Attila Lindner (2019). "Who pays for the minimum wage?" In: American Economic Review 109.8, pp. 2693–2727.

Hastings, Justine and Lydia Ashton (Jan. 2008). "Financial Literacy, Information, and Demand Elasticity: Survey and Experimental Evidence from Mexico". In: National Bureau of Economic Research, Inc, NBER Working Papers. doi: 10.3386/w14538.

Hastings, Justine, Ali Hortaçsu, and Chad Syverson (2017). "Sales force and competition in financial product markets: the case of Mexico's social security privatization". In: Econometrica 85.6, pp. 1723–1761.

Hastings, Justine S and Lydia Tejeda-Ashton (2008). Financial literacy, information, and demand elasticity: Survey and experimental evidence from Mexico. Tech. rep. National Bureau of Economic Research.

Haucap, Justus, Uwe Pauly, and Christian Wey (2001). "Collective wage setting when wages are generally binding An antitrust perspective". In: International Review of Law and Economics 21.3, pp. 287–307.

Hausman, Jerry, Gregory Leonard, and J Douglas Zona (1994). "Competitive analysis with differenciated products". In: Annales d'Economie et de Statistique, pp. 159–180.

Hermo, Santiago (2023). "Collective Bargaining Networks, Rent-sharing, and the Propagation of Shocks". In.

Hidalgo, Julian (2023). “This Craft’s For You! Entry and Market (ing) Competition in the US Beer Industry”. In: Job Market Paper.

Hijzen, Alexander and Pedro S Martins (2020). “No extension without representation? Evidence from a natural experiment in collective bargaining”. In: IZA Journal of Labor Economics 9.1.

Ho, Kate, Joseph Hogan, and Fiona Scott Morton (2017). “The impact of consumer inattention on insurer pricing in the Medicare Part D program”. In: The RAND Journal of Economics 48.4, pp. 877–905.

Illanes, Gastón (2016). “Switching Costs in Pension Plan Choice”. In: PhD Dissertation. Massachusetts Institute of Technology.

Jiménez-Hernández, Diego and Enrique Seira (2021). Should the government sell you goods? Evidence from the milk market in Mexico. Tech. rep. Stanford University Working Paper.

Kang, Zi Yang (2022). The Public Option and Optimal Redistribution. Tech. rep. Working paper.

Kroft, Kory, Yao Luo, Magne Mogstad, and Bradley Setzler (June 2021). Imperfect Competition and Rents in Labor and Product Markets: The Case of the Construction Industry. Working Paper w27325. National Bureau of Economic Research. doi: 10.3386/w27325. url: http://www.nber.org/papers/w27325.

Luco, Fernando (May 2019). “Switching Costs and Competition in Retirement Investment”. en. In: American Economic Journal: Microeconomics 11.2, pp. 26–54. issn: 1945-7669. doi: 10.1257/mic.20160332.

Luduvice, Andre, Tomas Martinez, and Alexandre Sollaci (2024). “Minimum Wage, Business Dynamism, and the Life Cycle of Firms”. In.

Lusardi, Annamaria (2008). Financial literacy: an essential tool for informed consumer choice? Tech. rep. National Bureau of Economic Research.

MaCurdy, Thomas (2015). “How effective is the minimum wage at supporting the poor?” In: Journal of Political Economy 123.2, pp. 497–545.

Manning, Alan (2003). “The real thin theory: monopsony in modern labour markets”. In: Labour economics 10.2, pp. 105–131.

— (2021). “The elusive employment effect of the minimum wage”. In: Journal of Economic Perspectives 35.1, pp. 3–26.

Markowitz, Harry (1952). “Modern portfolio theory”. In: Journal of Finance 7.11, pp. 77–91.

Martins, Pedro S (2021). “30,000 minimum wages: The economic effects of collective bargaining extensions”. In: British Journal of Industrial Relations 59.2, pp. 335–369.

McFadden, D (1981). Econometric Models of Probabilistic Choice”.

Melgar, Natalia and Laura Rovengo (2004). “El monopolio cervecero y las normas de defensa de la competencia en Uruguay”. In: Documento de Trabajo/FCS-DE; 4/04.

Miller, Nathan H and Matthew C Weinberg (2017). “Understanding the price effects of the MillerCoors joint venture”. In: Econometrica 85.6, pp. 1763–1791.

Miravete, Eugenio J, Katja Seim, and Jeff Thurk (2023). “Pass-through and tax incidence in differentiated product markets”. In: International Journal of Industrial Organization 90, p. 102985.

Montag, Felix et al. (2022). Mergers, Foreign Entry, and Jobs: Evidence from the US Appliance Industry. Tech. rep. CEPREMAP.

Neumark, David and Peter Shirley (2022). “Myth or measurement: What does the new minimum wage research say about minimum wages and job loss in the United States?” In: Industrial Relations: A Journal of Economy and Society 61.4, pp. 384–417.

Olley, Steven and Ariel Pakes (1996). The dynamics of productivity in the telecommunications equipment industry.

Ramirez, Lucia and Mijail Yapor (2022). "Salarios mínimos sectoriales e informalidad laboral en el Uruguay del Siglo XXI". In: 5ta. Edicion, concurso Juan Pablo Terra 2022 5.1.

Robinson, Joan (1933). The economics of imperfect competition. Springer.

Rosen, Sherwin (1986). "The theory of equalizing differences". In: Handbook of labor economics 1, pp. 641–692.

Rubens, Michael (2023). "Market Structure, Oligopsony Power, and Productivity". In: American Economic Review 113.9, pp. 2382–2410. doi: 10.1257/aer.20210383.

Salop, Steven C. and David T. Scheffman (1983). "Raising Rivals' Costs". In: The American Economic Review 73.2, pp. 267–271. issn: 00028282. url: http://www.jstor.org/stable/1816853 (visited on 09/22/2022).

Weyl, E Glen and Michal Fabinger (2013). "Pass-through as an economic tool: Principles of incidence under imperfect competition". In: Journal of Political Economy 121.3, pp. 528–583.

Williamson, Oliver E (1968). "Wage rates as a barrier to entry: the Pennington case in perspective". In: The Quarterly Journal of Economics 82.1, pp. 85–116.

Yeh, Chen, Claudia Macaluso, and Brad Hershbein (2022). "Monopsony in the US Labor Market". In: Available at SSRN 4049993.