ESSAYS ON HEALTH INSURANCE AND INDUSTRIAL ORGANIZATION

Cristián Figueroa Berríos

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 2023

ESSAYS ON HEALTH INSURANCE AND INDUSTRIAL ORGANIZATION

Dissertation Committee:

Michael D. Grubb (Chair)

Richard L. Sweeney

Julie Holland Mortimer

Abstract

This dissertation addresses questions in the health insurance and industrial organization fields. In the first chapter, I investigate how gender-based pricing bans affect health insurance markets offering long-term contracts. In the second chapter, I examine lapsing, and its implications, in a health insurance market offering long-term contracts. In the third chapter, I study the long-term implications of product unavailability in the beer market.

Chapter 1: In theory, guaranteed renewable (GR) insurance contracts can efficiently insure against reclassification risk without causing adverse selection on pre-existing conditions. In practice, however, adverse selection can still arise on other dimensions. In 2020, in response to protests demanding gender equality, Chile banned gender-based pricing in its private health insurance market. I investigate how this policy impacts Chile's health care system, which consists of a low-quality public option and a private market characterized by the use of GR contracts. I find that, if the ban is implemented, prices in the private market would increase as low-cost men switch to the public option and high-cost women

enter. Overall, the regulation causes a shift of surplus from men to women. The ban is regressive, as high-income groups benefit more than low-income groups, creating a trade-off between gender-based equity and income-based equity. Subsidies that induce low-cost enrollees to remain in the private market are the most effective mitigation strategy to contain higher premiums. Finally, relative to non-GR contracts, the number of individuals choosing the private market is lower under guaranteed renewability.

Chapter 2: Guaranteed renewable (GR) insurance contracts have the potential to efficiently protect individuals against reclassification risk without the negative side effects of price regulation, such as adverse selection. For these contracts to work properly, consumers must pay front-loaded premiums when healthy and stick with their contracts for many years in order to subsidize their future high-risk selves. This paper studies lapsing in the Chilean private health insurance markets, a system characterized by the offering of GR contracts. I find that most policyholders lapse their insurance plans just a few years after signing their contracts. I show that policies and lapse patterns predicted by standard theoretical models of long-term contracts are the opposite of those observed empirically. Finally, premiums increasing over time, and consumers lapsing their contracts because of those price changes, are a key determinant of insurers' profits.

Chapter 3: The marketing literature has investigated the processes potentially leading to brand building and the benefits these brands may enjoy over time. One of those possible benefits is resilience in the face of a reputational challenge or a crisis. This chapter focuses on the long-term implications of product unavailability. We leverage a quasi-natural experiment that exogenously re-

moved the top leading beer brands from retail stores for several weeks. We test whether these prolonged stockouts can erode market shares beyond the current or subsequent purchase occasions and study the potential mechanisms at play. Using panel data of consumer purchases before and after the product shortage, we observe that the top brands only partially recovered their pre-stockout market shares, especially among their most frequent buyers. We identify a sizable portion of consumers who tried small brands for the first time during the stockout period and remained to buy those products persistently. To control for prices, state dependence, and product availability, we estimate a choice model with heterogeneous preferences and find that exposure to stockouts has long-run effects on purchase behavior. We interpret our estimates as evidence that consumers facing a restricted choice set may learn or become aware of competing products with long-lasting consequences on preferences.

Contents

List of Figures

List of Tables

Acknowledgements

I am extremely grateful to my advisors Michael Grubb, Richard Sweeney, and Julie Mortimer. They did not only help me with my dissertation but they shaped my whole Ph.D. journey in many positive ways. I thank Michael because it was in one of his excellent classes that I decided I wanted to work on the intersection between IO and health insurance, and he was always supportive in that journey. My research ideas developed into papers thanks to the generous time he spent with me brainstorming about them. I thank Richard for taking the time to get to know my research, even if it was not in his main field. His comments were crucial to my dissertation, but also his support throughout the whole process, always asking about how I was doing before talking about research. I thank Julie for her excellent IO course, for her constant help during my Ph.D., and for being an inspiring role model in academia. It was an honor being her TA and RA for many years. Apart from my advisors, I am also indebted to Fernando Luco, from TAMU, whose help, in many ways, has been essential to this dissertation.

I thank all my friends in the department. I thank Kritika Goel for her help during the Ph.D. She has helped me since the first year studying for the comps, to the last year answering my questions about the job market. I thank Paul Sarkis and Jenny Cui as well for discussing research ideas with me during the third year. I thank Juan Luis Fuentes and Rodrigo González Valdenegro for being there when I needed to rant in Spanish. Outside the department, I thank Pinjas Albagli from LSE for being an outstanding and supportive friend, who was always willing to help in any way. I thank my siblings Francisco and Puchi for their ongoing support, both when I was in Chile and now in Boston. I thank my best friend Nicole, who made my transition to Boston much easier and helped me innumerable times. I thank my dog Mika for giving me calm and peace when I was feeling anxious during the job market.

Finally, I am grateful to the three most important people in my life. First, my dad, David Figueroa, for giving me everything in life and more. He is the harderst working and brightest person I know and I am constantly inspired by him. Second, my mom, María Loreto Berríos, for all the love she gave me for 34 years and counting. She is the strongest and the most loving person I know and I am where I am today because of her. Lastly, but not least, my wife Adriana Claudio-Hernández, for being the best and most supportive partner I could have asked for. She makes me feel at home even far away from my own country. I will always be grateful for her help and patience during the job market. I dedicate this dissertation to them.

1 Chapter 1: Adverse Selection and Equity in Markets with Guaranteed Renewable Contracts: Evidence from Chile [1]

1.1 Introduction

Guaranteed renewable (GR) insurance contracts guarantee that the terms of a policy will not be cancelled or modified, even if the policyholder develops a medical condition. They are popular in insurance markets such as term life insurance and long-term care insurance. In theory, such contracts can mitigate reclassification risk—the exposure of individuals to substantial premium increases due to changes in health status—without causing adverse selection. The intuition is that consumers pay front-loaded premiums to guarantee affordable coverage in the future, regardless of any possible negative health shocks.

The theoretical literature tipically rules out adverse selection in long-term insurance contracts by allowing for screening on pre-existing conditions at enrollment (Ghili et al., 2022). In practice, however, adverse selection can arise on other dimensions.[2] For instance, policies designed to promote gender equality in insurance markets can introduce adverse selection on gender. A common example of this is gender-based pricing bans, which prevent insurers from charging higher prices to women than men, despite the fact that woman have higher health care costs on average. Gender-based

[1]I thank Marlene Sanchez for patiently helping me with questions about the Chilean private health insurance market, and to Matías Stäger from *QuePlan.cl* for providing me access to their internal data of contracts. Finally, this research would not have been possible without the help of Mercedes Jeria from the Superintendencia de Salud of Chile that provided me with access to the administrative data. This work received support from the Institute for Humane Studies under grant no. IHS016632

[2]As another dimension of adverse selection, consumers might be forward looking and select high-quality plans before periods of costly elective medical spending. This is usually labeled as *selection on moral hazard* (*e.g.* Einav et al., 2013, Cabral, 2017, Diamond et al., 2021 and Shepard, 2022).

pricing bans have been been implemented in health insurance markets in the U.S. and Europe, and, since 2020, in Chile.

This paper investigates how gender-based pricing bans impact insurance markets characterized by the use of GR contracts. To do this, I estimate a discrete-choice model of plan choice using detailed data from the Chilean private health insurance system between 2013 to 2016. I find that, under a conterfactual ban, prices in the market increase by more than 30% as high-cost women enter the system and low-cost young men leave. Overall, the regulation causes annual consumer surplus to increase $373 per insured woman and decrease $279 per insured man, which is around 2% and 1% of average annual per capita income, respectively. The ban is regressive, as high-income groups benefit more than low-income groups, creating a trade-off between gender-based equity and income-based equity. Subsidies that induce low-cost enrollees (young men) to remain in the private market are the most effective mitigation strategy to contain higher premiums. Finally, relative to non-GR contracts, the number of individuals choosing the private sector is lower under guaranteed renewability.

The Chilean health care system is ideal for this type of study for at least three reasons. First, Chile has one of the very few health insurance markets, in addition to Germany, featuring GR contracts. By law, Chilean workers must choose between a public option (generally considered of low quality) or a private market with contracts that offer guaranteed renewability. In contrast to the typical short-term health insurance contracts available in the U.S., premium changes in GR contracts have to be community rated; that is, premium changes over the lifecycle of a contract are independent of changes in policyholders' health status. This imples that the higher costs of adverse selection, induced by the ban, will be spread evenly across contracts.

Second, in response to protests that started by the end of 2019 demanding gender equality in health care markets, Chile banned gender-based pricing in its private sector in 2020. In particular, before 2020, women in their 30s had to pay around 3 times more than men for the same plan. Third, unlike other insurance markets with GR contracts, the Chilean regulator gives researchers access to unusually rich individual-level data, thus allowing for a detailed study of the overall impact of the policy under examination.

I begin by providing several stylized facts regarding how a health insurance market offering GR contracts without a ban of gender-based pricing works. First, enrollees in the Chilean private sector are more likely to be high-income (3 times higher wages than enrollees in the public option), men (63%) and young (49% below age 35). The first two points are a consequence of a higher quality private system with gender-based pricing. The last point is contrary to what the theory of long-term contracts predicts (*i.e.* long tenure in a contract). Second, women pay higher premiums than men in the private market, conditional on plan quality, but they also spend more than men in health care. The latter implies that, if women have higher willingness-to-pay for insurance, adverse selection will emerge after the ban is implemented. Third, policyholders do not stick with their plans for long periods. Annual switching rates in the system are high (over 20%) and these rates are partially explained by changes in plan premiums. Thus, in spite of contracts' guaranteed renewability, policyholders are price sensitive and make frequent active choices in the market.

To quantify the impact of the ban on market outcomes (*e.g.* health-insurance premiums, market composition, and consumer surplus), and motivated by contracts' guaranteed-renewability and enrollees' response to premium changes, I estimate a two-stage discrete-choice demand model for plans in the private market using detailed administrative data from 2013 to 2016. The first stage determines the probability a policyholder makes an active choice of insurance plan in a year. This probability depends on announcements of premium increases, changes in personal income, and changes in family size. The second stage determines the choice of plan for the policyholder, which is a function of household characteristics, premiums, switching costs and the expected value of the hospital network. This last component is estimated using a hospital discrete-choice model and hospital admissions data, which also allow me to calculate expected costs per enrollee in each plan. The results of the second stage of the model show that women are relatively less price sensitive and have higher willingness-to-pay for coverage. These two facts together, in addition to women having higher health care costs, will drive adverse selection after gender-based pricing is banned. Importantly, including the first stage in the estimation matters. Switching costs are noticeably higher if the model is estimated without the first stage because in that case all the inertia in the market is attributed to switching costs. The two stages are separately identified by relying on exclusion restrictions.

I then solve for the equilibrium premiums during the same time period but under a counterfactual ban of gender-based pricing. In practice, the ban only applies to policyholders that decide to switch to a new plan. For that reason, the regulator spent considerable resources in promoting this policy, especially to women. The two-stage model allows me to force women to make active choices without removing switching costs. I show that, after the ban is implemented, premiums for private insurance plans increase as a result of adverse selection. The latter can be decomposed into intensive margin of adverse selection (movements of enrollees within the private market) and extensive margin of adverse selection (movements of enrollees in and out of the private market). In this setting, I find that extensive margin is the main driver of higher premiums. Without the ban, women below age 45 account for 28% of the private market, and that share jumps to 39% under the ban. In the case of men below age 45, the share goes from 47% to 37%. From a (remedial) policy perspective, subsidies that induce low-cost enrollees (young males) to remain in the private market are the most effective mitigation strategy to contain premiums. Monetary transfers between companies that balance expected costs are not as effective because they do not stop the movements of enrollees between the private and the public system.

The way GR contracts protect consumers from reclassification risk is by preventing insurers from changing prices of a given plan in response to changes in enrollees' health care costs. In contrast, short-term contracts without this kind of reclassification risk protection allow premiums of each plan to respond to changes in enrollment composition. This feature of non-GR contracts implies that, as a response to the higher costs of adverse selection, changes in prices are independent across plans.[3] Thus, premiums of individuals in low-quality plans will be higher under GR contracts (relative to a market with only non-GR contracts) because they must cross-subsidize selection on high-quality plans. Pairing this fact with the presence of a low-quality public option should lead to a lower number of enrollees choosing the private market than under non-GR contracts. In line with this hypothesis, I find that 60% of potential consumers select the private market after the ban of gender-based pricing is implemented under non-GR contracts compared to 52% who do so in a market with

[3]As noted by Handel et al. (2015) and Azevedo and Gottlieb (2017), among many others, in a competitive market with short-term health insurance contracts, companies will set prices independently across plans; that is, high-coverage policies that become more adversely selected over time will experience a higher rate of increase in prices compared to low-coverage policies.

GR contracts.[4]

The interaction between adverse selection on dimensions beyond pre-existing conditions and GR contracts in health insurance markets is an important topic not only for Chilean regulators, but for academics and policymakers in general. A recent literature studies the potential benefits of GR contracts and asks whether they should be implemented in health insurance markets in the U.S.[5] Ghili et al. (2022) charactize optimal long-term insurance contracts with one-sided commitment, as in Harris and Holmstrom (1982) and Hendel and Lizzeri (2003), and find that in certain scenarios, these contracts can achieve higher consumer welfare than ACA-like contracts. Similarly, Atal et al. (2020) show that GR contracts in Germany, despite not being optimally designed, obtain similar welfare outcomes as those in Ghili et al. (2022). These studies' favorable evaluation of GR contracts are based on assumptions that rule out the possibility of adverse selection. In practice, however, GR contracts can face adverse selection due to policy choice. Policymakers around the world are increasingly restricting insurers' pricing in favor of gender-based equity. This is because charging women higher premiums than men for the same level of coverage due to conditions such as pregnancy is considered unfair. As such, gender-based pricing is banned in the ACA Marketplaces in the U.S., and in all insurance industries in

[4]In this particular analysis, non-GR contracts represent GR contracts with a temporary waiver that allows insurers to respond to compositional changes in each plan independently (until a new equilibrium is reached). This is how the short-term contracts available in the U.S. would respond to adverse selection. The reason I cannot make a direct comparison between GR contracts and short-term contracts is that, in the data, I do not observe individuals choosing between these two contracts. Thus, in the simulation of the ban under non-GR contracts, I keep consumer preferences for plans fixed.

[5]In the case of policy research in the U.S., to name some examples, Cochrane (2017) and Pope (2020) advocate for long-term contracts to replace the current short-term contracts in the individual market, and Duffy et al. (2017) from RAND posit the question of whether the individual market could perform better under long-term contracts.

the European Union.[6] Similar restrictions exist in employer-sponsored health plans and in Medicare Advantage and Medicare Part D in the U.S., and in private health insurance markets in Australia, Colombia and South Africa.

Empirical studies of health insurance markets with long-term contracts are rare because few health insurance markets offer these contracts. Pauly and Herring (2006) show evidence of front-loaded prices in GR contracts in the individual market in the pre-ACA period. In the context of the small group market pre-ACA, Fleitas et al. (2020) document limited dynamic pass through of expected medical costs into premiums, and provide evidence that GR contracts indeed give protection against reclassification risk. Browne and Hoffmann (2013) study the German private health insurance market and find that front-loading in premiums generates lock-in of consumers. Furthermore, they document that consumers that lapse (*i.e.* switch contracts) are healthier than those who do not. Closest to this paper is Huang and Salm (2019), who consider the impact of banning gender-based pricing in German's private health insurance market. Using survey data on enrollment composition in each system, they find the mandate increases the probability of switching from the public sector to the private system for women relative to men, which implies a worsening of the private sector risk pool. All of these papers use data from specific employers, insurance companies or surveys, whereas I use detailed data on both the supply and the demand side of the whole Chilean private health insurance system.[7]

The remainder of the paper is organized as follows. Section 1.2 illustrates the

[6]Plans in the Affordable Care Act (ACA) exchanges in the U.S. cannot price discriminate based on sex or medical history, and there are binding restrictions on how age can enter pricing. In the case of Europe, the European Union's high court ruled in 2011 that sex cannot enter premium determination in health insurance, life insurance, or annuities.

[7]A few papers have used the same Chilean data to look at contract lock-in (Atal, 2019) and vertical integration (Cuesta et al., 2019).

intuition behind the effects of banning gender-based pricing in a two-tiered system with GR contracts. Section 1.3 describes the main institutional details of the Chilean health insurance system and the regulation that is the focus of this paper. Section 1.4 introduces the data and provides stylized facts of the Chilean private sector before the ban was implemented. Section 1.5 presents a two-stage empirical demand model of plan choice. Section 1.6 discusses the parameter estimates and simulates Chile's gender-based pricing ban. Section 1.7 concludes.

1.2 Conceptual Model

In this section, I provide a conceptual model that illustrates the effects of banning gender-based pricing in health insurance markets, and how these effects vary under GR contracts and non-GR contracts. In Section A in the Appendix, I develop a stylized theoretical model based on Einav et al. (2010) and Geruso et al. (2021) supporting the predictions of this section.

As described in Section 1.3, the Chilean health insurance market is a two-tiered system with a private sector and a safety net public option. For simplicity, here I assume that within the private market there is perfect competition with homogenous firms and zero administrative costs. Each firm from the private market sells two insurance plans, a high-coverage plan and a low-coverage plan. The public option offers a single plan that can be thought as lower coverage than any plan from the private market. The are only four consumer types in the market: "high-income males", "low-income males", "high-income females" and "low-income females". Gender determines health care costs, with female enrollees spending more in health care than male enrollees.[8] Income and gender together determine price sensitivity, with high-income groups and females being less price sensitive. Finally, I assume that females pay X times more for the same plan as their males counterpart (with $X > 1$). In Chile, before banning gender-based pricing, $X \approx 3$. At those prices, high-income males select high-coverage plans, low-income males and high-income females select low-coverage plans, and low-income females select the public option. Figure 1.1 shows that these choices are similar to the choices empirically observed in the Chilean health insurance market before the ban was implemented.

Figure 1.1: Gender-rated prices and selection

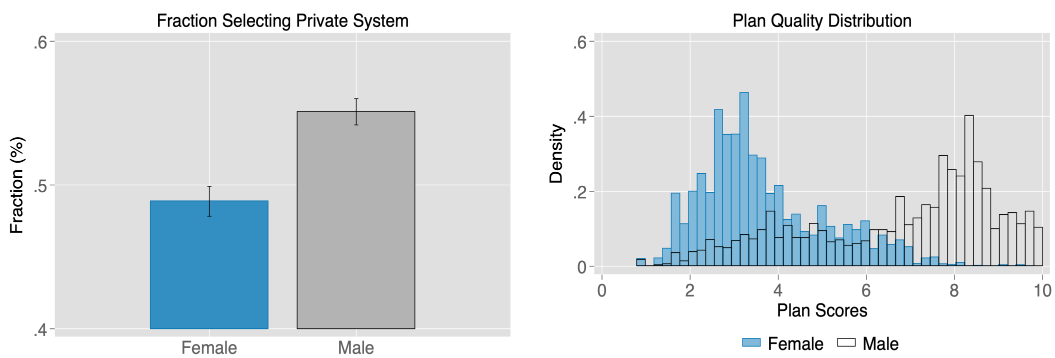

Notes: The left figure shows the fraction selecting the private system from a sample of single individuals that are realistically on the margin between the two systems (see Section D for details regarding sample construction). The right figure shows the distribution of plan scores (a measure of plan quality) from a sample of single individuals at the top income tercile (see Section 1.4.1 for details regarding plan scores).

Under a ban of gender-based pricing, firms are forced to charge females the same prices they were previously charging to males.[9] If high-coverage plans are more expensive, conditional on enrollment, than low-coverage plans, then the difference in

[8]See Figure 1.6 in Section 1.4.

[9]This is how the regulation was implemented in Chile in 2020. Another possibility would be to allow firms to increase prices of males such that they pay the same prices females are paying.

premiums between the two plans will go down. This implies that, in this scenario, high-income females can afford to buy high-coverage plans instead of low-coverage plans, which is called *adverse selection on the intensive margin*. Furthermore, the difference between the price of low-coverage plans and the public option will go down as well. This implies that low-income females are able to enter the private market instead of choosing the public option, which is called *adverse selection on the extensive margin*.

Since females have higher health care costs than males, firms will face pressure to increase premiums after these movements due to the zero profits condition. In response to higher prices in the private market, high-income males, that are less costly and more price sensitive than high-income females, switch from high-coverage plans to low-coverage plans, exacerbating *adverse selection on the intensive margin*. Conversely, low-income males move from the private market to the public option, exacerbating *adverse selection on the extensive margin*. These changes in enrollment composition put further pressure on firms to increase premiums again. This process of consumers switching and prices responding accordingly will continue until an equilibrium is reached. Overall, the private market is more adversely selected and with higher premiums after the ban is implemented. These market predictions are summarized in Figure 1.2.

Figure 1.2: Private market predictions: banning gender-based pricing

	Share Policyholders Private System		**Premiums Private System**
	Female	**Male**	
Baseline	Lower	Higher	Lower
Under ban	Higher	Lower	Higher

How would the outcomes of banning gender-based pricing change in a health insurance market offering non-GR contracts? The way GR contracts protect consumers from reclassification risk is by preventing insurers from changing prices of a given plan in response to changes in enrollees' health care costs. Therefore, the higher costs of adverse selection will be spread evenly across contracts and premiums of both high-coverage plans and low-coverage plans must increase at similar rates. This implies that low-coverage plans cross-subsidize selection in high-coverage plans.

In contrast, non-GR contracts without this kind of reclassification risk protection allow premiums of each plan to respond to changes in enrollment composition such that profits in each plan are zero (Azevedo and Gottlieb, 2017). Therefore, in response to the higher costs of adverse selection, the price of high-coverage plans will increase at a higher rate than under GR contracts to reflect the higher health care costs of high-income females. The opposite will hold for the price of low-coverage plans because it does not have to cross-subsidize selection on high-coverage plans. This implies that, after gender-based pricing is banned, the number of enrollees choosing the private market should be higher under non-GR contracts as more price sensitive consumers are able to buy low-coverage plans. Figure 1.3 below summarizes the differences between the two cases.[10]

To summarize, there are two main channels that explain higher prices in the private market after a ban of gender-based pricing. First, an increase in the *intensive margin of adverse selection* (movements within private market). Second, an increase in the *extensive margin of adverse selection* (movements between systems). To quantify the importance of each channel, an empirical model is needed. The effectiveness

[10] In Azevedo and Gottlieb (2017) the setting is a market with short-term contracts in which firms set plan prices equal to their annual average costs. In this setting, I am assuming that the same will

Figure 1.3: Private market predictions: GR versus non-GR contracts

	Premiums		Share Policyholders Private System
	Low Quality Plan	High Quality Plan	
GR contracts	Higher	Lower	Lower
Non-GR contracts	Lower	Higher	Higher

of additional policies to mitigate higher premiums, such as subsidies to low-income groups (targeting the extensive margin) or monetary transfers between companies (targeting the intensive margin), depends on which channel matters the most. Furthermore, under non-GR contracts, prices of high-coverage plans are higher, prices of low-coverage plans are lower, and a higher number of enrollees choose the private sector.

The actual outcome in the Chilean health insurance market after implementing the ban is an empirical question that depends on the observed preferences of both males and females. For example, if males do not actually lapse or switch their plans when premiums increase, then the degree of adverse selection will be lower. Estimating these preferences, and extending this simple framework to a more realistic setting with multiple demographic and income groups, and with multiple heterogenous firms and health plans, is the objective of Section 1.5. Importantly, an empirical model allows for consumer welfare assessments of the ban.

hold even though this is a market with long-term contracts. A way to think about this assumption is that firms will set plan prices equal to their long-term average costs. Thus, when defining expected costs, I actually mean long-term expected costs, and profits actually mean long-term expected profits.

1.3 Institutional Framework

The insurance system in Chile combines public and private provision.[11] The safety net public option, FONASA, is a pay-as-you-go system financed by the contributions of affiliates and public resources. The private sector—operated by a group of insurance companies—is a regulated health insurance market. In 2015, FONASA covered 77.3% of the population and the private system covered 15.1%. The remainder of the population is presumed to be affiliated with special healthcare systems such as those of the Armed Forces or to not have any coverage at all.[12]

Workers and retirees have the obligation to contribute 7% of their wages to the public system, or to buy a plan that costs at least 7% of their wages in the private system, with a cap of $207 per month.[13] The two systems differ in many respects, including provider access, premiums, coinsurance structure, exclusions, and quality. Unlike the private sector, in FONASA there are no exclusions based on pre-existing conditions, nor pricing based on age or gender, and there is no additional contribution for dependents. As a consequence, the private sector serves the richer, healthier, and younger portion of the population (Pardo and Schott, 2013).

The private health insurance market is comprised of 13 insurance companies, which are classified into two groups: six *open* (available to all workers) and seven *closed* (available only to workers in certain industries). This paper focuses only on *open* insurers, which account for 96% of the private market. Contracts in the private sector are, for the most part, individual arrangements between the insured and the

[11] The details of the Chilean health care system have already been described elsewhere, in particular Duarte (2012), Atal (2019), Cuesta et al. (2019) and Pardo (2019). I draw from those papers heavily in this section.

[12] See Figure D.3 for historical market shares in each segment of the health insurance market in Chile.

[13] All monetary amounts are measured in U.S. dollars using the exchange rate on December 2016.

insurance company. A key feature of these contracts is that they offer guaranteed renewability, meaning that enrollees can stay in their health insurance plans as long as they wish. Furthermore, insurers cannot change the characteristics of these plans over time. Only the price can change but in a limited way in order to protect consumers from reclassification risk (see details below). Once a policyholder has been in a contract for one year, she may lapse her contract and switch to another company. Switching plans within an insurer is allowed at any time.

The monthly premium for individual i under plan j in year t, P_{ijt}, is a combination of a base premium P^B_{jt} and a risk-rating factor r_i so that:

$$P_{ijt} = P^B_{jt} \times r(enroll\ age_i, gender_i) \tag{1}$$

where $r(enroll\ age_i, gender_i)$ is the risk-rating factor, which is a function of age at enrollment and gender. These factors are fixed over time as long as enrollees stay in their plans. For dependents, there is a similar $r(enroll\ age_i, gender_i)$ function and the full premium of the plan in that case is the base price P^B_{jt} multiplied by the sum of the risk-rating factors r_i of each member of the family. In the empirical model in Section 1.5, this is the premium policyholders observe in each plan. A couple of features of the market restrict the extent to which private firms can risk-rate their plans when individuals enroll. First, base premiums are set at the plan (and not the individual) level. Second, the r function is not individual-specific: each firm can have at most two r functions.

Several features of the plan determine the base premium P^B_{jt}. A plan has two main coinsurance rates, one for inpatient care and another for outpatient care. Unlike in the U.S., plans do not include deductibles and out-of-pocket maximums. Additionally, plans offer either unrestricted open networks or tiered networks.[14] Hospitals in Chile cannot deny health care to patients, and therefore all consumers have access to all hospitals, although they may have zero coverage from their plan.[15]

Base premiums are indexed to inflation, and adjustments to the base premium in real terms can be made once a year. In March of each year, companies must inform the regulator of their projected premium increases for the year. Each company must also inform their clients (through letters) about these increases, justify their reasons for the changes, and offer alternative contracts to their clients that keep monthly premiums more or less constant but that often imply lower coverage.

Reclassification risk

Reclassification could occur if firms could adjust the base premium P_{jt}^B of any given plan j based on the health status of the pool of enrollees in j. However, the market regulation involves also a restriction that limits the extent of reclassification of individuals already in a contract: the increase in P_{jt}^B of any particular plan j of insurer k cannot be higher than 1.3 times the weighted average price increase of all plans of insurer k. Formally

$$\frac{P_{jt+1}^B - P_{jt}^B}{P_{jt}^B} \leq \frac{1.3}{|J_k|} \sum_{j' \in J_k} \frac{P_{j't+1}^B - P_{j't}^B}{P_{j't}^B} \tag{2}$$

where J_k is the set of plans of company k.

Figure D.4 in the Appendix suggests that this regulation works in limiting the

[14]Unrestricted network plans provide the same coverage for all hospitals. Tiered networks offer differentiated coverage across sets of private hospitals, as PPO plans in the U.S.. Few plans offer restricted networks, as HMO plans in the U.S., and they are rarely observed in the data and not offered publicly. I do not consider them in my analysis.

[15]Other important characteristics of the plans are: a) Capitation scheme: Plans can either be capitated or not, b) Maternity-related expenses: Some plans do not have coverage for maternity-related expenses (in 2019 the regulator prohibited companies from selling these plans anymore). As these two characteristics also contribute to the determination of the base premium of the plan, I control for them in the demand model of Section 1.5.

extent of reclassification risk. For the season 2013/2014, which is representative of the pattern for all years in the sample, 5 out of 6 companies applied the same percentage price increase to all their plans, and the sixth firm increased its prices within a narrow window of 2.2% and 2.6%. Moreover, in Figure D.5 I plot the evolution of base prices by plan quality, showing that they all increase at similar rates within a company. This practice limits the correlation between individual health status and individual price increases, which implies limited reclassification.

Pre-existing conditions

Each new potential insured has to fill a "Health Declaration" before signing a new contract with a private firm. The companies are allowed to deny coverage of any pre-existing condition during the first 18 months of enrollment, or even to reject the prospective enrollee altogether. Although there is no available data on the extent to which insurers deny coverage, anecdotal evidence and conversations with industry actors suggests that this is a regular practice.

Hospitals in Chile

The health care system combines public and private provision. The public hospital network is broader than the private one, with 191 public hospitals compared to 83 private hospitals in 2016 (Chile, 2016). The private and public sectors are mostly segmented. Private insurers primarily cover admissions to private hospitals, whereas the public option mostly covers admissions to public hospitals. In fact, 97% of private insurer payments are to private hospitals, whereas only 3% are to public hospitals (Galetovic and Sanhueza, 2013). An important feature of this market is price transparency, as consumers are often able to obtain price quotes before choosing a hospital.

In the analysis of the network of health insurance plans, I focus on a particular geographic segment of the market. Specifically, I focus on the private hospitals in the city of Santiago, which is the largest health care market in the country and where more than a third of private hospitals and around half of the capacity is located (Galetovic and Sanhueza, 2013). Additionally, I only consider inpatient care, which represents more than half of health care expenditure. This segment is comprised of remarkably fewer players than the outpatient care sector, with more pronounced differences on prices and quality, and therefore strategic concerns associated with choosing the right network are more relevant in this segment.[16]

1.3.1 New Price Regulation

In March of 2020, in response to a wave of protests that erupted by the end of 2019 demanding, among other things, more gender equality, the regulator of the private health insurance market, *Superintendencia de Salud*, implemented a new policy banning gender-based pricing and adding restrictions on age-based pricing. Specifically, now the r function cannot depend on gender, and it can only depend on enrollment age in a predetermined way common across companies, allowing for much less price differentation between age groups (in the paper I refer to this regulation as *the ban of gender-based pricing*). The change is shown in Figure 1.4, where I plot the old risk-rating factors for a representative company in 2016 (left figure), the new risk-rating factors implemented in 2020 (right figure) and health care costs by gender and age (relative to a 30-years old male) across enrollees in the private market in 2016.[17]

[16] Plans are advertised and differentiated mostly on the inpatient hospital network offered. This is because outpatient procedures are less expensive and more homogenous. A limitation of this is that I can only predict inpatient health care costs. However, these predictions are at the group level (*e.g.* low-income young males), and at that level the correlation of spending between inpatient health care and outpatient health care is quite high.

[17] Figure D.1 in the Appendix repeats this excercise for the six companies together. Figure D.2

Figure 1.4: Risk-rating factors for one firm

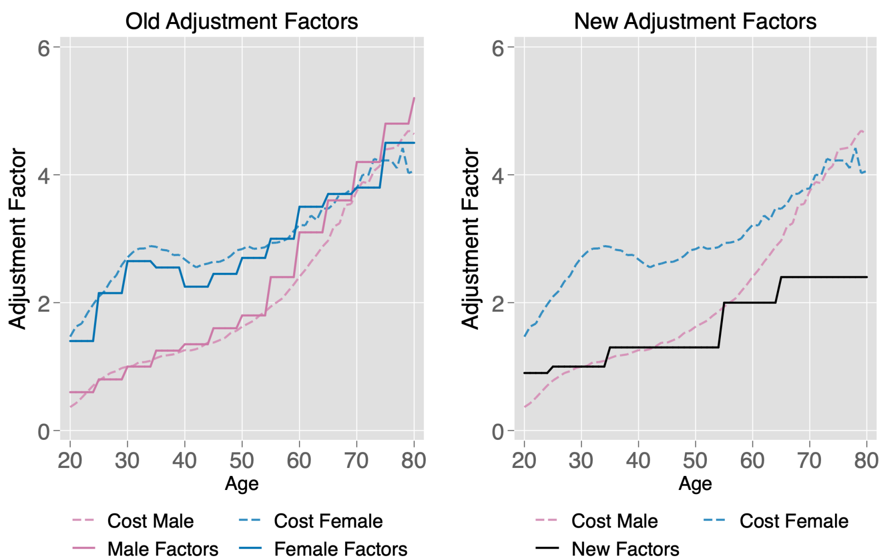

Notes: This figure shows the old risk-rating factors for a representative company in 2016 (left figure), the new risk-rating factors implemented in 2020 (right figure) and health care costs by gender and age (relative to a 30-years old male) across enrollees in the private market in 2016

Under the new risk-rating factors, as described in Section 1.2, firms from the private sector will face pressure to increase premiums due to the higher costs of adverse selection, both on the intensive margin (which plan to buy within private market) and extensive margin (private versus public system). The simulation in Section 1.6 allows for these two mechanisms to play a role in explaining higher prices in the system.

But what actually happened after the regulation was implemented in 2020? Even though a full empirical analysis is impossible due to the recency of the policy and a lag in the availability of the data, in Section E in the Appendix I provide anedoctal evidence indicating that costs in the private system have skyrocketed since the new risk-rating factors were implemented.[18]

Interestingly, a similar regulation was implemented in the German private health insur- ance market in 2011, the only other health insurance market in the world with GR contracts. Using survey data from both the private health insurance (PHI) sector and the social health insurance (SHI) sector, Huang and Salm (2019) find that women started switching from SHI to PHI at higher rates and more men were switching from PHI to SHI once the regulation was enacted, even though the latter result is not significant, likely because switching from PHI to SHI is highly restricted in Germany. In terms of prices, they show descriptive evidence that they increased at a much faster rate after the policy, in line with the results of this paper.

shows in detail the change of the r function for a particular firm. These are the changes in the factors of policyholders. A similar change was also implemented for dependents. In the simulation of the policy in Section 1.6, these two changes are taken into account.

[18]The simulation in Section 1.6 corresponds to the equilibrium in the market after implementing the ban, which, in practice, could take many years.

1.4 Data and Stylized Facts

1.4.1 Data

I exploit administrative data collected by the *Superintendencia de Salud* containing the universe of insureds in the private market for the period of 2013-2016 (Superintendencia, 2006).[19] Insurers must report data on individual claims to the regulatory agency. These data cover every health service provided to a private plan policyholder in 2013–2016, including financial and medical attributes along with consumer, plan and hospital identifiers. Additionally, I have data on all private plans offered during the period of analysis. This includes data on plans' company name, base prices, risk-rating function r, preferential networks, extra plan characteristics, availability in the market over time, and the date at which the plan was introduced in the market. Furthermore, I can match plans and their enrollees and observe basic demographics of policyholders and their dependents.[20]

I use administrative claims data to construct hospital admissions. I restrict the analysis to the 14 hospitals with highest market share, which account for 86% of the admissions in the data. The remaining hospitals are relatively small, and I group them into the outside option along with public hospitals. All these hospitals receive patients from all insurers in the market. Using claim dates and patient identifiers, I identify unique medical episodes of inpatient care which I label as admissions. The data contain detailed financial and medical information for each admission. Financial

[19] In practice, I have data from 2007 to 2016. The reason I focus most of the analysis in the period 2013-2016 is that these were stable years in this market (*e.g.* the regulator did not pass any important mandate during this period) and data from one of the companies is unreliable before 2013.

[20] One caveat is that I am not able to match spouses with different plans to the same household. That is, a low-income policyholder might actually be part of a high-income household and I cannot observe that.

information includes the hospital charges, insurer coverage, and consumer copayment. Medical information includes the diagnosis and the list of claims for different services provided by the hospital. I code admissions to diagnoses using ICD-10 codes, resulting in medical episodes that cover 16 diagnoses groups.[21] These diagnoses account for 90% of admissions and 92% of hospital revenue. Finally, I combine these data with plan attributes and consumer covariates, such as age, income, gender, and the number of dependents.

Even though plans in this market are differentiated by the coverage rate offered in each of the main hospitals of the capital, those rates are not available in the data. Instead, an online platform called *QuePlan.cl* provided me with access to their administrative plans database, allowing me to observe the actual contract of each plan. Thus, I can extract the actual coverage rate of each plan in each hospital. In Section C of the Appendix, I provide further details about how I construct these coverage rates. In addition, *QuePlan.cl* gave me access to "plan scores", a measure of plan quality. I use this variable in Section 2.3.[22]

Finally, in order to obtain information about households that are in the public option, I use the Chile National Socioeconomic Characterization Survey (CASEN) of 2017. The survey contains data on basic demographic and socieconomic characteristics of a representative sample of Chilean families, along with their health information (*e.g.* whether they have pre-existing conditions or not) and their choice of insurance system. I use this information in the simulation of the policy in Section 1.6.

[21] The list of diagnoses covers infections and parasites, neoplasms, blood diseases, endocrine diseases, nervous system diseases, ocular diseases, ear diseases, circulatory diseases, respiratory diseases, digestive diseases, skin diseases, musculoskeletal diseases, genitourinary diseases, pregnancy, perinatal treatments, and congenital malformation.

[22] The "plan score" is a standarized measure that goes from 0 to 10, where 10 represents a plan with almost perfect coverage for the most expensive private hospitals of Santiago. As the score goes

1.4.2 Stylized Facts

In this subsection, I present five stylized facts of the Chilean health insurance system under gender-based pricing. I focus on differences in enrollment composition between the public option and the private system, and differences in premiums paid for private plans between demographics groups. These differences are economically justified by comparing health care spending between these groups. Finally, I also document switching rates across consumers in the private market, and I study whether premium changes can explain these switching rates.[23]

First, Table 1.1 presents descriptive statistics of individuals in the public option and the private market in 2016 (see Section D for details regarding sample construction). Monthly wages of policyholders in the private market are, on average, 2.71 times higher than in the public option and they pay 3 times more for health insurance plans. Regarding market composition, females account for 45% of total enrollment in the public option, but only 37% in the private sector. Strikingly, while enrollees below 35 years old (26%) and above 55 years old (25%) account for similar shares in the public option, in the private system they represent 49% and 9% of the market, respectively. This last point is contrary to what the theory of long-term contracts predicts (*i.e.* long tenure in a contract). These differences between the two systems are in line with previous literature documenting that the private sector serves the richer, healthier, and younger portion of the population (Pardo and Schott, 2013).

Second, conditional on plan quality (plan scores), premiums paid differ greatly

down, the coverage rate for private hospitals goes down as well. The reason I do not use this variable in the empirical model of Section 1.5 is that I can only match scores to a subset of plans available in the market.

[23] Table D.1 in the Appendix presents descriptive statistics for the hospitals in the sample. Moreover, Table D.2 shows market shares for multiple demographic and socioeconomic groups.

Table 1.1: Descriptive statistics public and private market

	Wages (1000s US$)	Premium Paid (1000s US$)	Female %	Below 35 %	Above 55 %
Public	0.71	0.05	0.45	0.26	0.25
	(0.25)	(0.17)			
Private	1.93	0.15	0.37	0.49	0.09
	(0.86)	(0.09)			

Notes: This table shows descriptive statistics for my simulation sample. Only households without pre-existing conditions and with income high enough to consider the private system are in the public option sample (see Section D for details regarding sample construction). For wages and premiums paid, means are reported and standard deviations are in parentheses. Monetary values are measured in thousands of U.S. dollars for December, 2016.

between demographic groups. Specifically, Figure 1.5 shows that, for policyholders signing a new contract in 2016, young females pay much higher premiums for plans with similar quality than young males. For example, for a policy in the lowest plan score quartile, a young male pay around $71 dollars per month on average and a young female pay $105 dollars. Similar differences exist between young males and old males.[24] Third, female and old enrollees spend the most in health care. Figure 1.6 shows that these two groups are more likely to file a claim in a particular year and to spend more on health care, conditional on having positive spending. These two facts, differences by age and gender in premiums paid and health care spending, are in line with the adjustment factors and expected costs shown in Figure 1.4. Importantly, if women and older people also have higher willingness-to-pay for insurance, which is estimated in Section 1.5, then adverse selection will emerge after banning gender-based pricing.

Fourth, the switching rates in the private system are remarkable, especially con-

[24]In Table D.3 in the Appendix, I document that, beyond plan quality and firm fixed-effects, gender and age explain most of the variation in premiums paid across policyholders.

Figure 1.5: Premiums paid private market

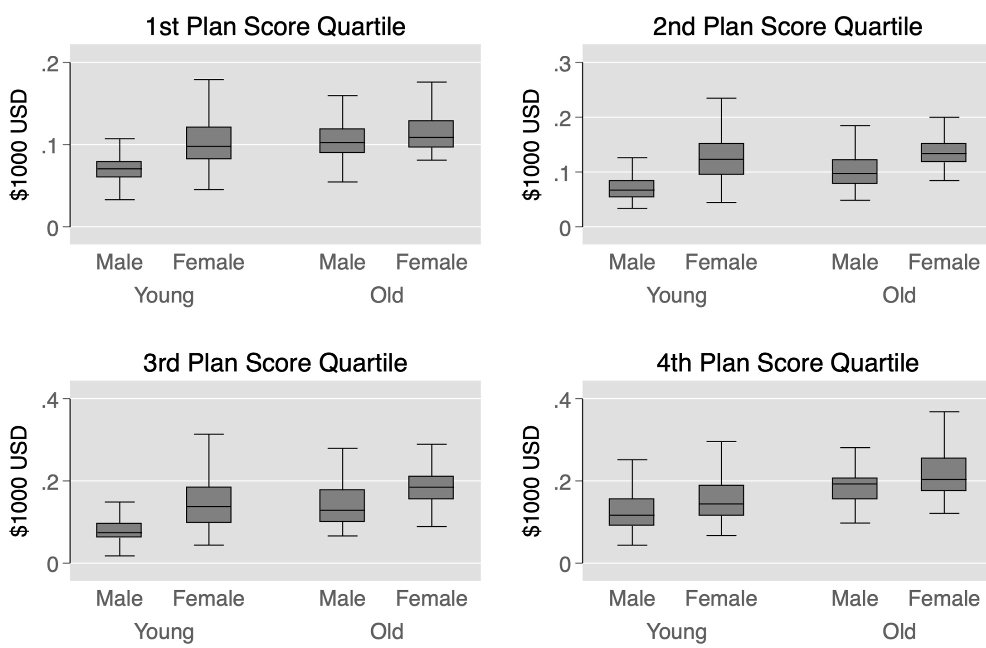

Notes: This figure shows a box plot of premiums paid by policyholders signing a new contract in the private market in 2016. Young is defined as an individual with 45 years old or less. Old is an individual with age above 45 years old. Prices are measured in thousands of U.S. dollars for December, 2016.

Figure 1.6: Health care costs by gender and age groups

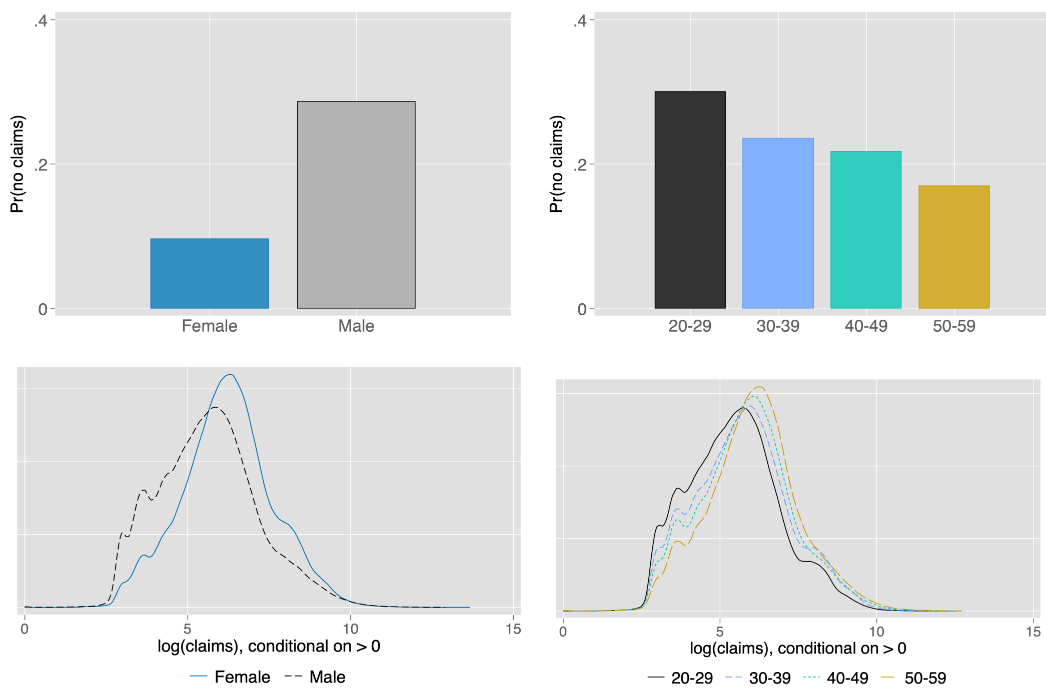

Notes: The upper figures show the probability of having zero claims in 2016 by gender (left upper figure) and age groups (right upper figure). The lower figures show the distribution of $log(claims)$ conditional on having positive claims during 2016 by gender (left lower figure) and age groups (right lower figure).

sidering that this is a market with GR contracts. Specifically, as documented in Table 2.1, almost 21% of policyholders lapse their plans during a year, with 11.7% switching plans within their companies, 6% switching companies, and 9.5% leaving the private market in order to go the public option. For comparison, the annual switching rate in Medicare Advantage is 8% and in Medicare Part D is 10%, which are similar private health insurance markets with a public option but with short-term contracts. Furthermore, in Figure D.6 in the Appendix, I show that only around 29% of enrollees stay in the same contract after 70 months, which undermines, in theory, the effectiveness of long-term contracts. In the next subsection, I explore one of the reasons why people lapse their plans at such a high rate: changes in premiums.

Table 1.2: Switching rates private market

Switching Rates (%)	
Within company switching	11.71
Company switching	5.95
Public option switching	9.54
Any switching	20.76

Notes: This table shows switching rates across policyholders in the data in 2016.

1.4.3 Evidence of Lapsing

It is expected that companies will increase prices after the ban of gender-based pricing is implemented because of the higher costs of adverse selection. Thus, it is important to identify consumers' price sensitivity in this market as it is possible that, due to the guaranteed renewability of their contracts, policyholders will stick with their plans even after their premiums increase. Therefore, in this subsection, I explore whether policyholders in this market lapse their insurance plans in response to premium changes.

To answer this question, I examine the relationship between premium increases and the probability of switching plans within a company.[25] Figure 2.3 shows the results of an event study regression where the dependent variable is a dummy equal to one if the consumer switches plans within an insurer, and the event is the month in which the contract was signed in the first place. In Chile, the signing month is the month in which premium changes are applied to each policyholder. Importantly, given the nature of the contracts, these premium changes are the sole reason for enrollees to switch plans in this month in particular (*i.e.* it is the only characteristic of the plan that is changing in the contracts). Furthermore, I also control for policyholder fixed effects and date (month-year) fixed effects, meaning that I am looking at the effect of the signing month on lapsing at the individual level, controlling for dates in which lapsing might be higher (or lower) than average. In order to have a clean panel of policyholders, I restrict the estimation sample to enrollees that do not switch insurance companies and that do not leave the private market and re-enter in later dates. This exercise is done on a 10% random sample of policyholders.

As documented in the figure, I find a large spike in the probability of switching plans in response to premium changes, and the effect appears consistent across income terciles. In terms of magnitude, the lapsing probability goes from 0.3% on average in any month of the year to 1.2% during the signing month. In Figure D.7, I look specifically at the plans to which policyholders lapse, finding that they switch to lower

[25] I do not examine switching companies because this type of switching can only be done after being one full year in your current company. Thus, a spike in switching one year after enrollment could be attributed to premium changes or to the fact that policyholders could not switch before that. In the structural model, I can incorporate any type of switching because a period is defined as a year.

Figure 1.7: Probability of switching plans due to price changes

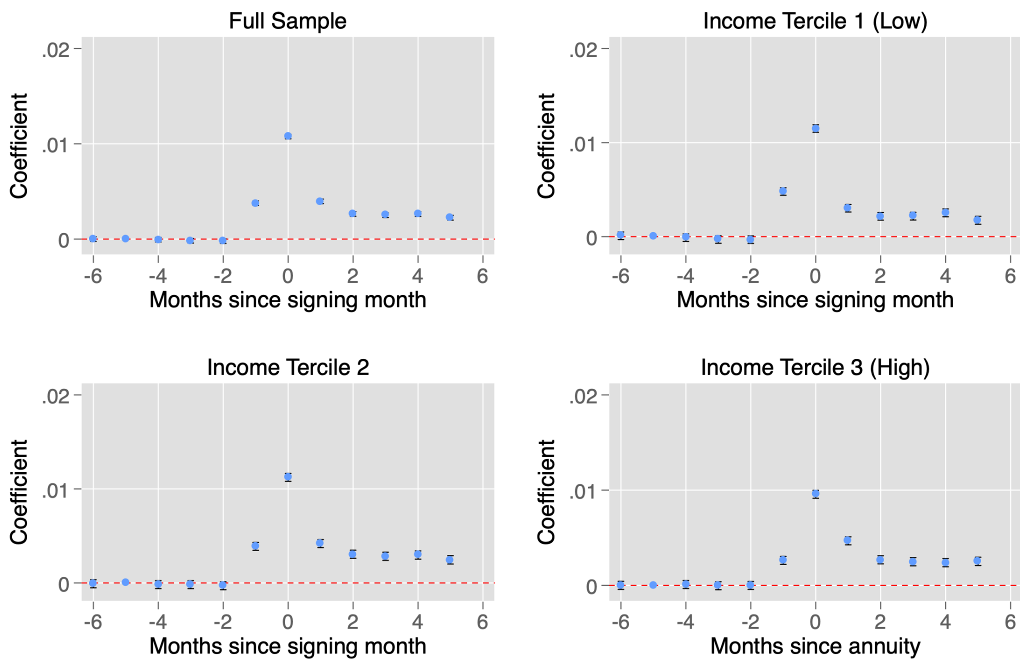

Notes: This figure shows an event study regression where the dependent variable is a dummy equal to one if a consumer switches plans within an insurer. The event is the month in which price changes are applied to health plans. Controls include individual fixed effects and date (month-year) fixed effects. I restrict the estimation sample to policyholders that do not switch insurance companies and that do not leave the private market and re-enter in later dates. Additionally, I drop individuals with zero or missing income at any month. This exercise is done on a 10% random sample of policyholders.

quality plans to keep their premiums paid roughly constant.[26] Finally, in Figure D.8, I repeat the same exercise but with the dependent variable now being switching from the private market to the public option. I show that policyholders on the lowest income tercile are prone to leave the private sector in response to premium changes.

These results are not only relevant for this particular setting, but also for the literature of insurance with long-term contracts in general. In theoretical models, lapsing only occurs by healthy consumers trying to find a lower price in the spot market. This holds by assuming that income paths are flat or that consumers have perfect foresight of their income paths.[27] Nonetheless, recent papers in other industries show that there might be additional reasons why individuals lapse their long-term contracts (*e.g.* Gottlieb and Smetters, 2021). The lapsing evidence presented above motivates the choice of a two-stage empirical model, estimated in Section 1.5. That is, due to front-loading and guaranteed renewability (Atal, 2019), policyholders have incentives to remain passive in their contracts, but they will make active choices in the market if they face premium changes in their plans (in the model I also add changes in personal income and changes in family size as explanatory variables). Furthermore, decomposing inertia into switching costs and remaining passive matters for the simulation in Section 1.6 because it allows me to force women to make active choices after implementing the ban without removing switching costs.

In summary, the stylized facts documented in this section provide a clear picture of the Chilean private health insurance market before the ban of gender-based pricing was implemented. First, the market is mainly composed by rich, young and male

[26] Additionally, in Figure D.10 I look at switching by age and gender groups, finding that they all lapse in response to premium changes.

[27] Theoretical models with long-term contracts treat policies as securing a certain level of income (or consumption), which is feasible by assuming *ex-ante* known income paths. However, if income

policyholders. Second, females and older enrollees pay higher prices than young males for the same plans, but they also spend the most in health care. Third, for a market offering GR contracts, switching rates are non-trivial and they are explained, at least partially, by changes in premiums. Thus, if prices increase after the ban is in place, lapsing will be triggered in the market. These facts will be part of the empirical model and the simulation of the ban in the following sections.

1.5 Empirical Model

The evidence provided in Section 1.4 does not allow me to quantify the welfare effects among consumers of banning gender-based pricing in the private health insurance market. To allow for such quantification, I formulate a two-stage econometric model of how individuals choose which contract to enroll in. The first stage determines whether a policyholder makes an active choice of insurance plan in a year. The second stage, conditional on making an active choice, determines the choice of plan for the policyholder. The choice of plans is a function of household characteristics, premiums, the expected value of the hospital network in the upcoming year, and switching costs. The two stages are estimated simultaneously.[28] For the calculation of the expected value of the hospital network, I separately estimate a hospital discrete-choice model (Capps et al., 2003 and Ho, 2006). Importantly, this model allows me to calculate expected costs per enrollee in each plan.

paths are uncertain, then insurance against these income shocks is needed as well, which is not provided by real applications of long-term insurance policies (*e.g.* the GR health insurance contracts offered in Chile and Germany).

[28]Abaluck and Adams (2021) and Heiss et al. (2021) estimate similar models to explain inertia in Medicare Part D. In my setting, this model is even more compelling because what they call *inattention* might be rational behavior in an insurance market with GR contracts due to front-loading. That is also the reason why in this paper I use the terms *active* versus *passive* instead of *attention* versus *inattention*.

I note that, besides the expected value of the hospital network in the upcoming year, the demand model abstracts away from forward-looking behavior. Forward-looking generates an option value that may affect current choices, since they affect the set of feasible future choices. Specifying a dynamic demand model would require to specify individual's perceptions about the distribution of their future preference shocks, supply-side behavior, and discount rate.[29] Nevertheless, in Section B of the Appendix I study the effects of a regulation implemented in 2011 that changed the incentives of consumers to be forward-looking. I find no structural change, on aggregate, on how policyholders choose plans after this regulation. Thus, this behavior does not seem to be a first order issue in this market.

1.5.1 Discrete-Choice Demand Model

First-Stage

The first stage of the demand model determines the probability, s^a_{ft}, that an incumbent household f makes an active choice of insurance plan in year t ($a_{ft} = 1$). This probability depends on announcements of premium increases, changes in personal income, and changes in family size. It takes the following logit form:

$$s^a_{ft} \equiv Pr(a_{ft} = 1|\mathbf{x_{0t}}) = \frac{exp(\mathbf{x_{0t}}\tau^{\mu})}{1 + exp(\mathbf{x_{0t}}\tau^{\mu})} \tag{3}$$

where $\mathbf{x_{0t}}$ is a vector containing the percentage increase in the premium of the household's own plan, an indicator for changes above 30%, in either direction, in personal

[29]The complexity of choice in health-insurance as well as the evidence showing choice inconsistencies in this market is arguably a main reason why most recent papers estimating health insurance demand in dynamic settings do not incorporate forward-looking behavior. A recent literature uses Medicare part D dynamic pricing incentives to estimate discount factors and myopia in drug purchases, finding strong levels of myopia (*e.g.* Abaluck et al., 2018 or Dalton et al., 2020).

income, and an indicator for changes in family size (typically indicates birth of a new child). τ^{μ} are the corresponding first stage parameters to be estimated.

Second-Stage

Conditional on making an active choice, an incumbent household f considers all health insurance plans in their choice set $\mathcal{Y}_{ft}$. To change plans, however, they must incurr in a switching cost γ_f. Each year after signing a GR contract, the incumbent household f's utility from choosing plan j in period t takes the following form:

$$u_{fjt} = \alpha_f p_{fjt} + \beta_f \sum_{i \in f} EU_{ijt} + \phi X_{fj} + \gamma_f y_{fjt-1} + \epsilon_{ijt} \tag{4}$$

where ϵ_{ijt} is distributed Type 1 EV. Utility thus depends on the plan premium p_{fjt}, EU_{ijt} and X_{fj}, explained below, switching cost γ_f, and an indicator for remaining in the same plan y_{fjt-1}.[30] EU_{ijt} is the expected utility of consumer i of household f from the hospital network of plan j at time t, which depends on how sick consumers expect to be and the plan coverage in the hospitals they expect to go. I separately estimate EU_{ijt} using a hospital discrete-choice model and hospital admissions data, a standard procedure in the health insurance literature (Capps et al., 2003 and Ho, 2006). Importantly, this model also allows me to calculate expected costs for each household in each plan of the market, which I then use in the simulation of the ban to calculate profits. See Section C of the Appendix for details about the estimation of this model and how EU_{ijt} is calculated. Finally, X_{fj} includes additional plan characteristics,

[30]In this particular case, switching costs could be paperwork or enrollment costs, acclimation costs, and even reclassification risk. The latter could occur if families with sick members find it harder to switch plans, maybe due to the possibility of being denied coverage. Importantly, the coefficients themselves may pick up both true switching costs and persistent unobserved heterogeneity. For my purposes, it is not clear that is important to distinguish these factors. Doing so would matter primarily for dynamic price competition, which I do not model.

such as whether the plan is capitated and whether the plan covers maternity-related expenses, and firm fixed effects interacted with household characteristics in order to control for segmentation at the company level.

Under these assumptions, the conditional probability that household f chooses plan j from choice set $\mathcal{Y}_{ft}$ in year t given they make an active choice is:

$$s^{y}_{ftj} \equiv Pr(y_{ft} = j|a_{ft} = 1; \gamma_f, .) = \frac{\alpha_f p_{fjt} + \beta_f \sum_{i \in f} EU_{ijt} + \phi X_{fj} + \gamma_f y_{fjt-1}}{\sum_{k \in \mathcal{Y}_{ft}} \alpha_f p_{fkt} + \beta_f \sum_{i \in f} EU_{ikt} + \phi X_{fk} + \gamma_f y_{fkt-1}} \quad (5)$$

For new households entering the private market, the main difference is that they do not have a default plan, and, hence, they do not have switching costs. Policyholders are assumed to choose the plan that gives them the highest utility. Notice that this formulation assumes that beneficiaries choose the option with the highest "perceived" utility, which may not necessarily correspond to the "best" plan for them from an actuarial risk-protection perspective (Abaluck and Gruber, 2011, Abaluck and Gruber, 2016). For the analysis of risk allocation and choices in the simulation of the ban, however, this "perceived" utility is exactly the object of interest.

Household preferences for health insurance may reflect both health risks, as well as horizontal tastes and risk aversion. Importantly, these differences in willingness-to-pay for insurance between demographic groups will determine the degree of adverse selection in the market after the ban is implemented. To capture these features of insurance demand in the model, I include rich observed heterogeneity in the specification of marginal utility. Specifically, I allow preferences for plan characteristics, α_f and β_f,to vary by two age groups, gender, two income groups (low- and high-income) and household size. Additionally, in the case of switching costs γ_f, I allow them to depend on pre-existing conditions in order to capture lock-in (Atal, 2019).

1.5.2 Identification and Choice Sets

In this subsection, I go through the sources of identification for all the components of the discrete-choice model outlined above. Additionally, I explain how I deal with the problem that choice sets are unobserved in the data.

The simplest explanation for why a policyholder decides to stay in the default plan is that she carefully compared plans and came to the conclusion that her default plan is the most attractive one available for the next year. Given the incentives for policyholders to remain passive in their GR contracts (*e.g.* front-loaded premiums), this alone is unlikely to be the only reason that explains choosing the default plan. In the second stage of the model, I account for a large part of the plan characteristics that should be relevant to a policyholder's decision of what plan to choose. Allowing for preference heterogeneity provides a comprehensive model of deliberate plan choices. The remaining persistence that cannot be explained by plan features and preference heterogeneity is what I call inertia and attribute to switching costs and remaining passive.

For identification, I exploit the fact that, each year, there are new cohorts of consumers entering the private market and that are forced to make an active choice. This implies that I observe policyholders choosing with and without frictions from the same menu of contracts. On the one hand, the preference parameters that enter in the second stage utility are identified from the choices of consumers who enter the market under the assumption that unobservables are uncorrelated with premiums and plan characteristics. Given the rich set of observables included in the model, this assumption should be reasonable. On the other hand, the difference in behavior between households with default plans and without default plans is attributed to inertia (Handel, 2013).

First stage parameters τ^{μ} are identified by the fact that the explanatory variables driving the probabilities of making an active choice (first stage) are quite different than the ones explaining the plan choices (second stage), creating exclusion restrictions that add to the identification of the model. For example, if the premium of the $t-1$ plan increases in t—a change that is made salient by letters sent from insurers—my model implies an increase in the probability of becoming active, so the policyholder starts comparing alternatives. Conversely, if the level of the premium of the $t-1$ plan is high for year t relative to the alternatives, a fact that is only apparent after comparing plans, this might contribute to overcoming switching costs and changing the plan (Heiss et al., 2021).[31]

Regarding contract characteristics, these are identified by variation in premiums and offered hospital networks across plans within an insurer and household type. Still, endogeneity of premiums might threaten identification. In particular, insurers could provide coverage for additional services that are not captured by offered hospital networks. If these extra services affect premium setting, it would cause an endogeneity concern. However, this is not a big problem in this setting because I observe in detail the characteristics of each plan, which allows me to add them to X_{fj} in the demand model.

Finally, in the case of the choice set of policyholders, it is unrealistic to believe

[31] Another identification issue arises in dynamic panel data models with lagged dependent variables if unobserved initial conditions and unobserved heterogeneity are correlated. However, this is not a problem in my setting since I am using data from 2013-2016 to estimate the model, while still having data back from 2007.

that consumers observe all the plans for each company, especially given that there are thousands of plans available each year. Due to the mandatory rule of spending at least 7% of your wage in a health insurance plan, the choice set that each individual faces is restricted to only plans with a premium equal or higher than 7% of their wages. Additionally, according to the function r, the price of plans depends also on the age and gender of the policyholder. In practice, each firm offers betweeen 2 to 4 plans to consumers, conditional on a year, preferences for plan type (*i.e.* tiered network or unrestricted choice), hospital network offered, and their income and demographics. Therefore, for estimation purposes, I split the sample of consumers into different groups depending on the year, age group, gender, income group, and household size. Within each group, I find the most popular plans purchased in a year, conditional on plan type, network of providers offered and premium deciles. These are the plans, along with the default plan, that enter the choice set of each household.[32][33]

1.5.3 Maximum Likelihood Estimation

Given my parametric specification, I can estimate all model parameters simultaneously using maximum likelihood. For notational convenience, I collect all parameters in the vector θ. Since I assume independence across households, the likelihood function is the product of the individual likelihood contributions $\mathcal{L}_f(\theta)$. The observed outcome for household f is the sequence of plan choices made in years $t = 1, \ldots, T$ (making active choices is unobserved). Let j_{ft} denote the observed plan choice in year

[32] Companies usually offer a plan at your 7% income level and the other plan(s) offered are priced above that 7% (either around 10% of your income level or even above that in some cases). In practice, then, conditional on plan type and hospital network, I include three plans in the choice set of each policyholder, one at 7% of their income, one between 7% and 10%, and one above 10%.

[33] Cuesta et al. (2019) and Robles-Garcia (2022)) follow similar strategies in their settings with unobserved choice sets.

t. Thus, the likelihood contribution of each household is given by:

$$\mathcal{L}_f(\theta) = Pr(y_{f1} = j_{f1}, ..., y_{fT} = j_{fT}|\theta, d_f) \tag{6}$$

where d_f collects the histories of all the observed covariates. Further, conditional on these covariates, the time-constant switching cost component γ_f, and whether the policyholder is active or not a_{ft}, the household choices are independent over time. I can therefore write:

$$Pr(y_{f1} = j_{f1}, ..., y_{fT} = j_{fT}|\theta, d_f) = \prod_{t=1}^{T} Pr(y_{ft} = j_{ft}|\theta, d_f, a_{ft}, y_{ft-1}) \tag{7}$$

These probabilities are readily available given my assumptions on the first and second stages. First, I obtain the choice probabilities for period t:

$$Pr(y_{ft} = j_{ft}|\theta, d_f, a_{ft}, y_{ft-1}) = s^a_{ft} Pr(y_{ft} = j|a_{ft} = 1, d_f, y_{ft-1}) \tag{8}$$

$$+(1 - s^a_{ft}) Pr(y_{ft} = j|a_{ft} = 0, d_f, y_{ft-1}) \tag{9}$$

This expression involves probabilities that condition on making active choices ($a_{ft} = 1$), which is not observable. By plugging in s^a_{ft} and s^y_{ftj} from equations (3) and (5), respectively, it is easy to see that the choice probabilities have different forms depending on whether the household has a default plan or not and whether the household stay in the default plan or switch to a different plan, which are both observable events. If I_{ft} is an indicator for a household having a default plan, then for households without a default plan ($I_{ft} = 0$), the probability of making an active choice is $s^a_{ft} = 1$, thus the probability of choosing plan j is equal to s^y_{ftj}. For a household with a default plan ($I_{ft} = 1$), I either observe that they switch plans (and thus choose a plan j with $y_{ft-1} \neq j$) or that they stay in their default plan (and choose plan j with $y_{ft-1} = j$). Switching (choice of j for $I_{ft} = 1 \wedge y_{ft-1} \neq j$) can only occur if the household is active. The probability of choosing a plan j that is not the default plan is therefore $s^a_{ft} s^y_{ftj}$. Staying in the default plan (choice of j for $I_{ft} = 1 \wedge y_{ft-1} = j$) can occur either if the household is not active (with probability $(1 - s^a_{ft})$) or if they are active (with probability s^a_{ft}) and choose the default plan j as the preferred option (with probability s^y_{ftj} for $y_{ft-1} = j$).

Putting everything together, for any plan j:

$$Pr(y_{ft} = j|.) = \begin{cases} s^y_{ftj} & \text{if } I_{ft} = 0 \\ s^a_{ft} s^y_{ftj} & \text{if } I_{ft} = 1 \wedge y_{ft-1} \neq j \\ s^a_{ft} s^y_{ftj} + (1 - s^a_{ft}) & \text{if } I_{ft} = 1 \wedge y_{ft-1} = j \end{cases} \tag{10}$$

To obtain the probability of household f's entire choice sequence, I plug this expression into equation (7).

1.6 Estimation Results and Simulation

1.6.1 Parameters Estimates

Table 1.3 displays the results for the first stage of the demand model.[34] The estimates suggest that policyholders start making active choices in the market in response to changes in premiums, changes in personal income and changes in family size. Moreover, the average household is active in a particular year with a probability of about

[34]Section D in the Appendix provides details about the construction of the sample used in the estimation of the model.

30%, and hitting all enrollees with the average premium change (6.4%) would increase their probability of making an active choice by 15.5%.

Table 1.3: Parameters estimates - First stage

Variable	Coeff	S.E.	Average Marginal Effect
%Δ Premium	0.037	(0.002)	15.52%
$1\{\|\Delta Income\| \geq 30\%\}$	1.020	(0.029)	82.75%
$1\{\|\Delta Family\| > 0\}$	1.896	(0.058)	158.88%
Constant	-1.316	(0.023)	
Mean Active		29.82%	
Observations		3,706,628	

Notes: This table shows the first stage logit estimates of the discrete-choice model. Average marginal effects for the price are computed by simulating the increase in the probability of making an active choice after imposing that all enrollees are hit by the average premium change on in the sample. For the income change (or the family size change), the average marginal effect is computed by simulating the increase in the probability of making an active choice in the case where no one has an income change (or family size change) versus when every enrollee face an income change (or family size change). Standard errors are in parenthesis.

Table 1.4 documents the results of the second stage of the demand model. Specifically, the first panel of the table highlights the premium coeffients,[35] the second panel shows the expected utility of health care coefficients, and the third panel displays the switching costs coefficients. Each row represents a different estimate for a different household group. Different columns indicate a different specification. Columns (1) and (3) consider insurer fixed effects, and columns (2) and (4) consider insurer fixed effects interacted with household groups. Additionally, columns (3) and (4) include the first stage in the maximum likelihood estimation. Column (4) is my preferred specification.[36]

[35] See Figure D.11 for estimated premium elasticities across my sample.

[36] Table D.4 in the Appendix assesses which model has a better fit to the data. The full model that includes the first stage and insurer-household group fixed effects provides the best fit according

The main conclusions from the table are: (i) females and older enrollees are relatively less price sensitive; (ii) females and older enrollees have a higher relative preference for higher "quality" plans (in terms of coverage and hospital networks); (iii) including the first stage in the estimation matters, especially in the case of the switching costs estimates. Points (i) and (ii) together will drive adverse selection after gender-based pricing is banned, that is, the two groups with higher health care costs also have the higher willingness to pay for better insurance plans.

To highlight the effect of including the first stage in the model, in Figure 1.8 I plot the distribution of annual switching costs (in dollars) across enrollees for specifications (2) and (4) from the table. These switching costs are calculated by dividing the switching costs parameters with the premium parameters (and adjusting for annual costs). The differences are striking. In the case of the demand model including the first stage, the average annual switching costs ($376) account for only 21% of average premiums paid. On the contrary, when the demand model is estimated without the first stage, the average annual switching costs ($890) account for almost 50% of average premiums paid. The intuition behind this result is that estimating a discrete-choice model with a dummy variable for the default plan, but without a first stage, implies that each policyholder compares the available plans in each year and deliberately makes a choice—which seems unrealistic when the availability of a default contract with guaranteed renewability invites *passiveness*.

1.6.2 Simulation

Description of Simulation

With the parameter estimates from the demand model, I can now simulate a

Table 1.4: Parameters estimates - Second stage

Variable	(1) Coeff	(1) S.E.	(2) Coeff	(2) S.E.	(3) Coeff	(3) S.E.	(4) Coeff	(4) S.E.
α_f - Plan premium p_{fjt}								
age ≤ 45	-21.928	(0.283)	-22.579	(0.300)	-27.422	(0.402)	-28.096	(0.421)
age > 45	-16.466	(0.315)	-17.625	(0.338)	-22.917	(0.476)	-23.885	(0.488)
Single female	2.807	(0.218)	3.319	(0.231)	2.223	(0.306)	3.102	(0.309)
Family	8.373	(0.238)	8.661	(0.260)	9.851	(0.333)	9.694	(0.350)
High-Income	5.659	(0.255)	7.733	(0.278)	8.697	(0.406)	11.350	(0.433)
β_f - WTP network EU_{ijt}								
age ≤ 45	-0.651	(0.209)	2.295	(0.276)	-1.646	(0.269)	2.170	(0.335)
age > 45	-0.753	(0.195)	3.603	(0.291)	-1.747	(0.277)	2.886	(0.388)
Single female	1.775	(0.163)	1.119	(0.278)	2.713	(0.202)	1.895	(0.338)
Family	0.999	(0.170)	-0.470	(0.298)	1.621	(0.206)	0.347	(0.368)
High-Income	6.144	(0.164)	-1.649	(0.253)	6.265	(0.197)	-2.155	(0.300)
γ_f - Lagged plan y_{fjt-1}								
age ≤ 45	4.555	(0.017)	4.626	(0.020)	1.168	(0.111)	1.509	(0.083)
age > 45	5.283	(0.029)	5.405	(0.047)	2.435	(0.097)	2.679	(0.078)
Single female	-0.078	(0.021)	-0.141	(0.038)	0.234	(0.041)	0.189	(0.051)
Family	-0.737	(0.022)	-0.708	(0.040)	-0.405	(0.044)	-0.389	(0.041)
High-Income	1.257	(0.022)	1.053	(0.000)	2.368	(0.080)	2.085	(0.059)
Pre-existing condition	0.341	(0.024)	0.330	(0.024)	0.271	(0.046)	0.287	(0.042)
Observations	3,706,628		3,706,628		3,706,628		3,706,628	
First Stage	N		N		Y		Y	
Plan Characteristics	Y		Y		Y		Y	
Insurer FE	Y		N		Y		N	
Insurer-Demographics FE	N		Y		N		Y	

Notes: This table shows the logit estimates of the discrete-choice model. The first panel displays the premium coeffients, the second panel shows the expected utility of health care coefficients, and the third panel shows the switching costs coefficients. Estimates vary across age groups, household composition, and income. Different columns indicate a different specification. Columns (1) and (3) consider insurer fixed effects, and columns (2) and (4) consider insurer fixed effects interacted with household groups. Additionally, columns (3) and (4) include the first stage in the maximum likelihood estimation. Standard errors are in parenthesis.

Figure 1.8: Switching costs

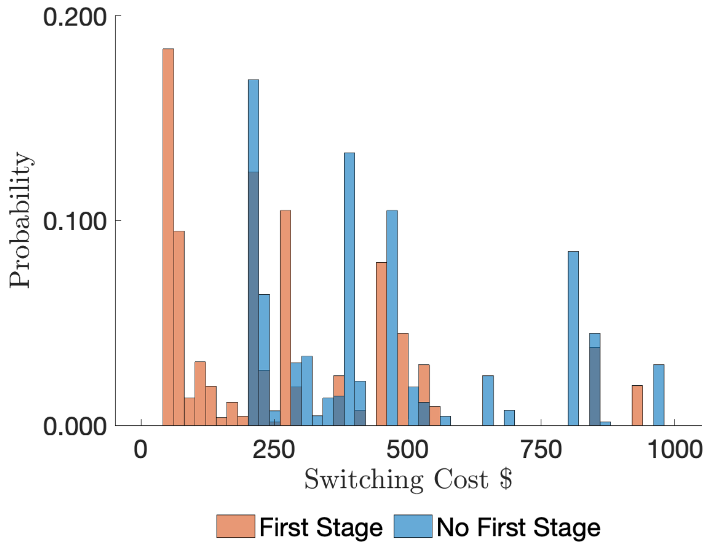

This figure shows the distribution of annual switching costs (in dollars) across enrollees in the estimation sample. These switching costs are calculated by dividing the switching cost parameters with the premium parameters (and adjusting for annual costs). Orange bars represent switching costs for the demand model including the first stage. Blue bars represent switching costs for the demand model without the first stage. Prices are measured in U.S. dollars for December, 2016.

gender-based pricing ban in the Chilean private health insurance market. Specifically, for the year 2016, I change the old risk-rating factors by the new risk-rating factors that will determine the final premium of each plan. In this subsection, I describe the steps taken in order to perform this simulation.

First, in practice, the regulation only implements new risk-rating factors to enrollees if they switch to a new plan, but for incumbents that decide to stay in their old plans, they keep their old risk-rating factors. This means that the ban does not change premiums of default plans, which implies that, in the simulation, few policyholders

to multiple tests.

will make active choices. That is not realistic as the regulator spent considerable resources in promoting this policy, using targeted advertising to women and the elderly population (see Figure D.12 for an ad example). The two-stage model allows me to force women and older people (above 55 years old) to make active choices without removing switching costs.

Second, in the case of households from the public option, I only select consumers that are realistically on the margin between the two systems, and I include them in the simulation. *Realistically on the margin* means that their income is not sufficiently low such that they would never enter the private sector even if prices drop by a high amount, nor do they have pre-existing conditions such that they could be denied coverage by private insurers. Additionally, I also assume that potential consumers from the public option are actively making plan choices in every iteration of the simulation. In Section F in the Appendix, I study how the results change if instead only 50% of these households are active in each iteration.

As noted in Section 1.2, these movements from female and older enrollees will put pressure on private insurers to increase premiums. There are two main reasons for this: (1) adverse selection on the intensive margin, and (2) adverse selection on the extensive margin. Fully modelling the supply side of a health insurance market with GR contracts is a complicated task that remains an open challenge in the literature (*e.g.* Atal, 2019) and is beyond the scope of this paper. Instead, I allow firms to adjust prices following two simple rules. First, following existing regulation of GR contracts, the rate at which they increase premiums has to be the same for all their plans (*i.e.* protection against reclassification risk). Second, they must keep their profits constant relative to the sample period (2013-2016).[37] That is, in each iteration, insurers will raise prices, in a restricted way, until profits are back to sample period levels. Profits are calculated using prices, expected costs per enrollee in each plan (obtained from the hospital discrete-choice model in Appendix Section C) and enrollment. In Section F in the Appendix, I report results from a sensitivity analysis where I run multiple simulations with different profit levels.[38]

The moment companies increase their prices, consumers might start making active choices in response to these changes. Once active in the market, they decide whether to remain in their default plans or to switch to a new, likely cheaper, plan. Furthermore, some of them, especially low-income enrollees, might decide to even leave the private market altogether. If these switchers also have relatively lower health care spending, costs in the private market will further raise, putting further pressure on firms to increase premiums again. I repeat this process until I reach convergence, which in this case means that companies do not have incentives to change prices. This is what I call an *equilibrium* in the market.

Mitigation policies

I perform two additional simulations of the private sector under the ban, but with mitigation policies that address the higher costs of adverse selection. Particularly, I simulate two strategies that are commonly used in health insurance markets in other

[37]Specifically, I calculate profits for each firm in each year from 2013 to 2016. Then, I calculate the median profit for each firm across those years. These are the profit levels that insurers must keep constant. Importantly, when updating prices, I am only considering *static* changes in costs in each iteration. If, instead, firms take into account the *long-term* changes in costs in order to set premiums, then prices might increase even further.

[38]I keep the set of contracts fixed throughout the simulation, meaning that insurers are not allowed to create new health insurance plans after the ban is implemented. In practice, however, companies can create new plans if they wish. This assumption will bias my results if companies use this channel as a way to practice "cream-skimming" in the market. In Section E of the Appendix, I discuss and

countries but that have not been implemented in Chile yet. The first strategy is risk adjustment transfers (or monetary transfers between companies). The objective of this policy is to address adverse selection on the intensive margin by balancing the expected costs across insurers such that, the most adversely selected firms, in relative terms, do not have more incentives to increase premiums than the least adversely selected firms. In short, firms with a pool of high-cost enrollees, relative to the market, receive money from firms with a pool of low-cost enrollees, until all companies have the same expected costs.

The second mitigation policy is subsidies to low-income enrollees in the private market. The purpose of this policy is that a higher share of those enrollees remain in the private system because they are the ones that are likely to leave after premiums increase. Importantly, they are also the consumers with the lowest expected health care costs (young men). Thus, by keeping them in the private sector, costs will not increase as much as they would otherwise. In other words, this strategy is targeting adverse selection on the extensive margin.

Notice that the effectiveness of each policy depends on which margin of selection is more important once the ban is implemented. If most of the increase in costs is due to adverse selection on the intensive margin, then risk adjustment transfers will be more effective. If adverse selection on the extensive margin explains most of the rising costs, then subsidies will be more effective. If both margins are important, then it is probably a good idea to use both policies at the same time.[39]

Non-GR contracts

show evidence that this was not the case in the Chilean private sector in the years following the ban.

[39] One caveat of this analysis is that subsidies are politically much harder to implement because, unlike risk adjustment transfers, they are not a budget neutral policy.

Finally, I simulate an scenario in which the ban is implemented and premiums of each plan are allowed to reflect enrollment composition in that plan (*i.e.* non-GR contracts). According to Section 1.2, the number of enrollees in the private market should be higher under these kind of contracts as low-quality plans do not have to cross-subsidize selection in high-quality plans, which is exactly the case under GR contracts.

To conduct these counterfactual simulations, I need a method to find new equilibrium premiums in each plan as the enrolled population changes. Under the assumption of a perfectly competitive individual insurance market, Azevedo and Gottlieb (2017) provide an algorithm to compute this equilibrium.[40] In brief, I augment the pool of households with a mass of "behavioral consumers" who incur zero covered health costs and choose each available contract with equal probability; the inclusion of these behavioral types ensures that all contracts are traded. I then apply a fixed point algorithm in which, in each iteration, consumers choose contracts according to their preferences, taking prices as given. Prices are adjusted up for unprofitable contracts and down for profitable contracts until an equilibrium is reached.[41]

In my setting, this method is an imperfect approximation because it assumes perfect competition. Therefore, and following Dickstein et al. (2021), to determine equilibrium premiums in this counterfactual environment, I apply a modified version of the algorithm from Azevedo and Gottlieb (2017) that includes a fixed markup by

[40]The model in Azevedo and Gottlieb (2017) features an insurance market with short-term contracts. As in the main simulation of the ban under GR contracts, I assume that insurers respond to *static* changes in profits. If that is the case, I can apply then the algorithm provided by Azevedo and Gottlieb (2017) to find equilibrium prices.

[41]Given the large number of plans available in the market, to simplify the algorithm, within each firm I pool similar plans into "plan groups", according to their plan scores and plan characteristics. Therefore, firms adjust prices independently across "plan groups", but they apply the same price change for plans within each "plan group".

plan (the markup is fixed in that it does not vary with the equilibrium outcome). In addition, this approach of fixed markups per plan does not allow cross-subsidization of plans within an insurer, that is, an insurer cannot subsidize an unprofitable high-quality plan with a profitable low-quality plan. As discussed in Azevedo and Gottlieb (2017), one can micro-found this restriction with a strategic model with differentiated products. If an insurer taxes one plan to subsidize another, it risks being undercut on the taxed plan and only selling the money-losing option.

1.6.3 Simulation Results

In Figure 1.9, I decompose, approximately, how much of the rise in premiums after banning gender-based pricing is due to adverse selection on the intensive margin or adverse selection on the extensive margin. First, I simulate an escenario in which I implement the new risk-rating factors only for female and older enrollees without allowing for consumer switching ("Ban without switching").[42] Second, I simulate the implementation of the new risk-rating factors and allowing for switching, but only within the private market ("Ban with within market switching"). Third, the last scenario ("Ban with full switching"), is the full simulation of the ban. As shown in the figure, adverse selection on the extensive margin is the most important channel in explaining the rise in prices compared to intensive margin adverse selection. Specifically, extensive margin adverse selection accounts for 54% of the price increase and adverse selection on the intensive margin accounts for only 25%. This implies that subsidies should be a more effective tool than risk adjustment transfers in containing the rise in premiums.

[42] Prices increase, on average, in the full population in this scenario because of a base effect (the percentage price increase for men is, in absolute terms, higher than the percentage price decrease

Figure 1.9: %Δ in prices after ban: decomposition

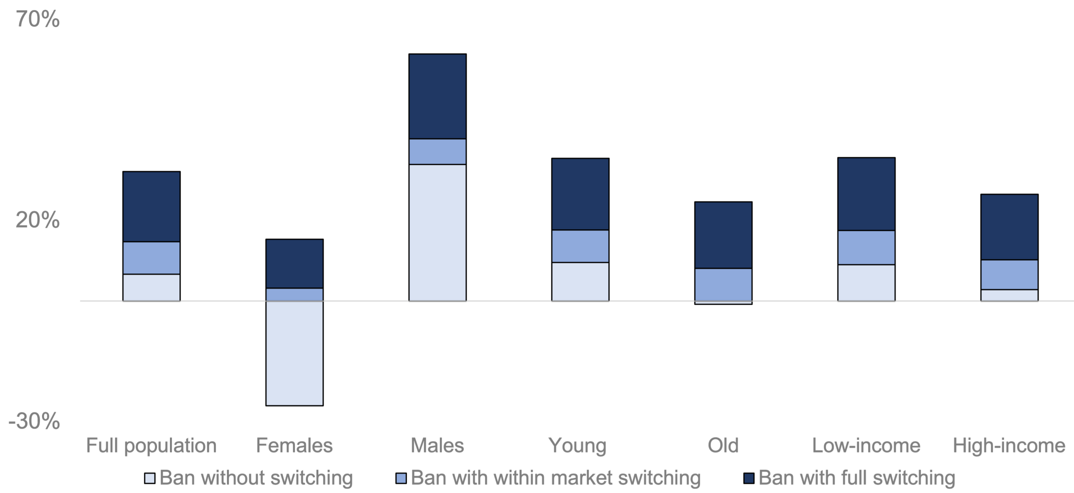

This figure shows how much of the premium increases after simulating the ban of gender-based pricing are due to intensive margin adverse selection and extensive margin adverse selection. To perform this, first I simulate an escenario implementing the new price factors only for female and older enrollees without allowing for consumer switching ("Ban without switching"). Second, I simulate an escenario in which I implement the new price factors and allowing for switching but only within the private market ("Ban with within market switching"). The last scenario ("Ban with full switching") is the full simulation of the ban.

Table 1.5 shows the full results of the simulation. The first panel of the table displays the share of the simulation sample that is in the public option,[43] the second panel shows the percentage change in overall prices, and, the last panel documents the (annual) change in consumer surplus.[44] Each row in each panel is a different scenario (*i.e.* a different simulation) and each column is a different demographic or income group. All changes are relative to the baseline scenario under the old risk-rating factors in 2016.

The ban causes an influx of women from the public option to the private market—the share of women choosing the public option drops from 47 to 38 percent. At the same time, it causes men to move in the opposite direction—the share of men choosing the public option increases from 39 to 55 percent. The net effect is an increase in the total share choosing the public option from 43 to 48 percent. Most of these movements are from young women and young men. Particularly, in terms of composition, without the ban, young men account for 47% of the private market and young women for 28%. After the ban, these shares change to 37% and 39% respectively. These two factors—women entering the private market and males leaving—are

for women) and because median profits across 2013-2016 are higher than profits in 2016.

[43]For the simulation, I impose that high-income policyholders stay within the private market. The reason is that, for them, it is not a rational decision to go to the public option, even if prices are higher in the private system. My preferred specification of the empirical model supports this prediction as well, but because of logit shocks, a small share of policyholders would choose the public option in the simulation. In the data, high-income policyholders rarely move to the public option, and when they do is mostly due to large income shocks.

[44]For consumer surplus I use the compensating variation metric of change in expected consumer surplus as derived in Small and Rosen (1981). In general, the change in expected consumer surplus for household f from a change in the characteristics of the choice set from J to J' is (conditional on the household-specific parameters of the utility function θ_f):

$$\Delta E[CS_f|\theta_f] = \frac{1}{\alpha_f}\left[ln\big(\sum_{j'\in J'} exp(v_{fj'})\big) - ln\big(\sum_{j\in J} exp(v_{fj})\big)\right]$$

where v_{fj} is the part of the utility function in Equation (4) without the unobserved portion ϵ.

Table 1.5: Simulation results

	Total	Women	Men	Young	Old	Low-income	High-income
Share in Public Option (%)							
Baseline	42.67	47.39	39.44	40.09	49.20	69.05	-
Ban	47.88	37.71	54.81	44.92	55.36	77.63	-
Ban + Risk Adjustment	47.61	36.79	54.98	44.51	55.43	77.19	-
Ban + Subsidies	44.83	36.05	50.81	41.97	52.05	72.68	-
Ban + RA + Subsidies	42.23	36.16	46.37	38.60	51.41	68.47	-
%Δ in Overall Prices - Relative to Baseline							
Ban	32.23	-10.60	61.45	35.51	23.94	35.72	26.60
Ban + Risk Adjustment	32.22	-10.58	61.42	35.38	24.22	36.02	26.10
Ban + Subsidies	21.41	-17.83	48.18	24.45	13.71	24.42	16.56
Ban + RA + Subsidies	13.16	-23.30	38.04	15.91	6.22	16.19	8.30
Δ in Consumer Surplus ($) - Per enrollee per year - Relative to Baseline							
Ban	-15	373	-279	-73	132	-68	71
Ban + Risk Adjustment	0	396	-269	-57	144	-59	97
Ban + Subsidies	66	481	-217	-14	268	-16	198
Ban + RA + Subsidies	126	556	-167	30	370	21	296

Notes: This table shows the results of the simulation implementing the ban of gender-based pricing. The first panel of the table displays the share of the simulation sample that are in the public option, the second panel shows the percentage change in overall prices relative to the baseline and the last panel documents the change in consumer surplus relative to the baseline. Each row in each panel is a different scenario (*i.e.* a different simulation), and each column is a different demographic or socioeconomic group.

the main reason prices go up more than 30% in the private sector (the second panel). Figure 1.10 plots the distribution of premiums for single females and single males, before and after implementing the ban. In the baseline scenario, female premiums are higher than male premiums, which is in line with the factors displayed in Figure 1.4. Once the ban is implemented, as expected, premium distributions are much more homogenous between the two groups. Importantly, as noted by the vertical dashed lines in the figure, the average premium shifts to the right due to the higher costs of adverse selection.

The pattern of consumer welfare changes mirrors those for premium changes: there is little average change for the whole population, but surplus shifts sharply from men to women. The regulation causes average annual consumer surplus to increase \$373 per insured woman and decrease \$279 per insured man, which is around 2% and 1% of average annual per capita income, respectively. Splitting the sample by income, surplus shifts from relatively low-income to relatively high-income enrollees. This is because the women who benefit from price cuts are, on average, higher income than the men who face price increases, and the women who gain private market coverage are on average higher income than the men who lose this coverage. Relatedly, high-income women are one of the groups with the highest willingness-to-pay for insurance while the opposite is true for low-income males.[45]

Mitigation policies, risk adjustment transfers and subsidies, have a positive effect on surplus. However, as expected given the results in Figure 1.9, subsidies are much more effective. The intuition behind this is that, by keeping young males in the private market, costs do not increase as much as they would otherwise, mitigating then

[45] In terms of government revenue, I find that, with the implementation of the ban, the government would gain more than 10 million dollars (in 2016 dollars). This is between 0.1% and 0.2% of what

Figure 1.10: Histogram of premiums by gender

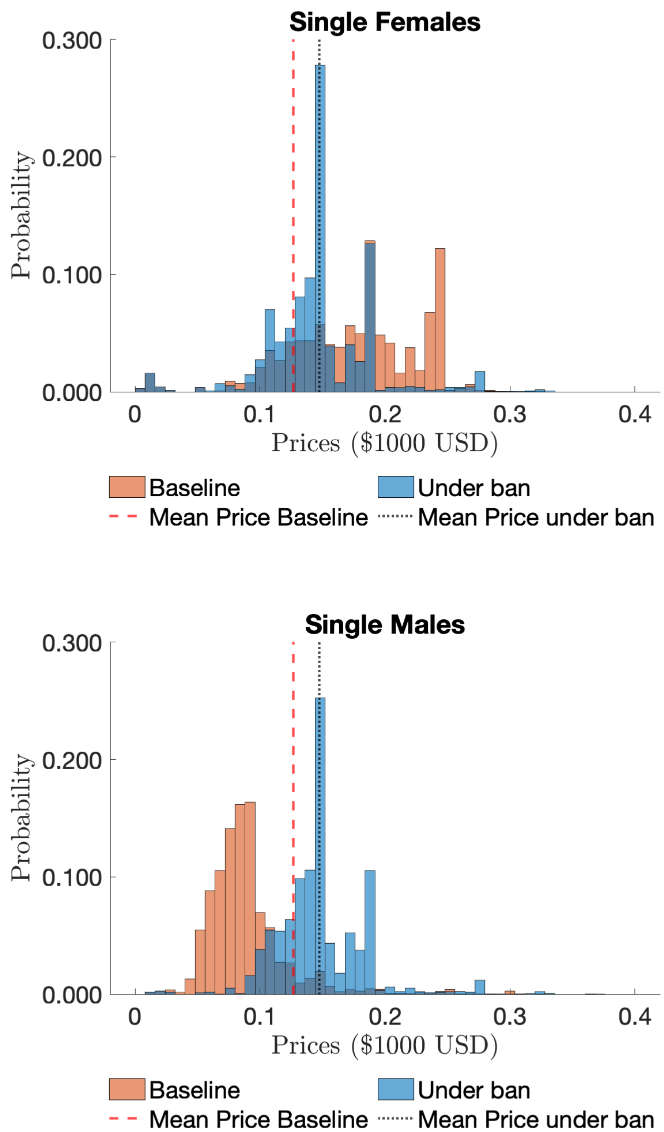

This figure shows the distribution of premiums for both single females (top plot) and single males (bottom plot) in the baseline (red bars) and under the ban of gender-based pricing (blue bars). Red dashed vertical lines despict the average premium across the two groups in the baseline and black dashed lines despict the average premium under the ban. Prices are in thousands of dollars of 2016.

adverse selection on the extensive margin. Importantly, the gain in consumer surplus from the subsidies (over \$3 millions) outweighs the costs of them (\$1.5 millions). Implementing both strategies at the same time delivers the best outcome for consumers. In this scenario, low-income groups have a positive change in consumer surplus, which makes the ban less controversial. This result is in line with the two margin problem described in Geruso et al. (2021); addressing one margin of selection (*e.g.* intensive margin) can exacerbate the other margin (*e.g.* extensive margin), therefore the best solution is to address both margins at the same time.

Finally, notice that the average consumer surplus results in Table 1.5 do not account for inequality, that is, this is an unweighted welfare measure. In contrast, in Figure 1.11 I show the consumer surplus impact of the ban if I introduce different levels of aversion to inequality. More specifically, to assess the equity implications of banning gender-based pricing, I rely on income as the measure of inequality and consider alternative welfare weights for different income groups. Following Handel et al. (2021), the consumer surplus of an individual in income group y_δ, where $y_\delta = 1$ for low-income and $y_\delta = 2$ for high-income, is weighted by $y_\delta^{-\epsilon}/(\sum y_\delta^{-\epsilon}/2)$ for $\epsilon = 0.5$, $\epsilon = 1.0$ and $\epsilon = 1.5$.[46] As shown in the figure, consumer surplus is lower the higher the aversion to inequality. This is intuitive because low-income groups have a lower change in surplus than high-income groups. Thus, when the regulator cares more about income inequality, banning gender-based pricing looks less appealing.

A different way to assess the implementation of the ban is instead to assume that the regulator has aversion to health care spending inequality. To do this, I split enrollees into four groups depending on their age (young and old) and gender (female

the public option spent on enrollees' health care in 2016 (FONASA, 2019).

[46]This is called the Atkinson index of inequality (Atkinson, 1970), which uses a social welfare

and male), and rank them according to their health care costs. That is, the consumer surplus of an individual in group y_δ is weighted by $y_\delta^{-\epsilon}/(\sum y_\delta^{-\epsilon}/4)$ for $\epsilon = 0.5$, $\epsilon = 1.0$ and $\epsilon = 1.5$. In Figure 1.12 I show the results using this variant of the Atkinson index of inequality. In contrast to the results in Figure 1.11, now the (weighted) average consumer surplus is positive for all scenarios and is quite large for high aversion to inequality. This follows because the groups that benefit the most from the ban are also the ones with highest health care spending (*i.e.* female and older enrollees).[47]

Figure 1.11: Consumer surplus impact and income inequality aversion

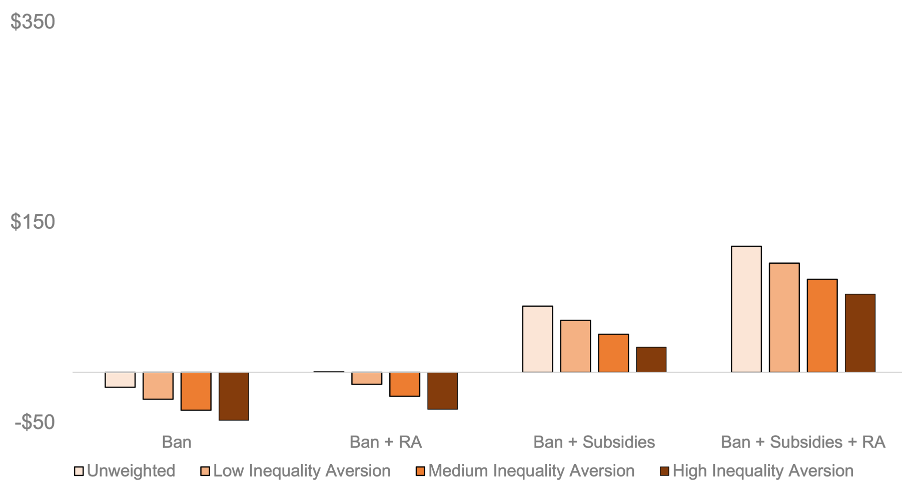

This figure shows the average consumer surplus change (in annual US per household) of implementing the ban of gender-based pricing with and without mitigation policies. The consumer change impact is calculated with equal weights for both income groups, low inequality aversion, medium inequality aversion and high inequality aversion. Weights y_δ are computed as $y_\delta^{-\epsilon}/(\sum y_\delta^{-\epsilon}/2)$ for $\epsilon = 0.5$, $\epsilon = 1.0$ and $\epsilon = 1.5$.

function of the form $y_i^{1-\epsilon}$ with $\epsilon \geq 0$ as measure of inequality aversion. Here, I follow Handel et al. (2021) and I weigh the welfare gain for each household depending on the income group they are in by $y_\delta^{-\epsilon}/(\sum y_\delta^{-\epsilon}/2)$, which ensures comparability with the unweighted case.

[47] In unreported results, I also use income net of health care spending as a measure of inequality.

Figure 1.12: Consumer surplus impact and health spending inequality aversion

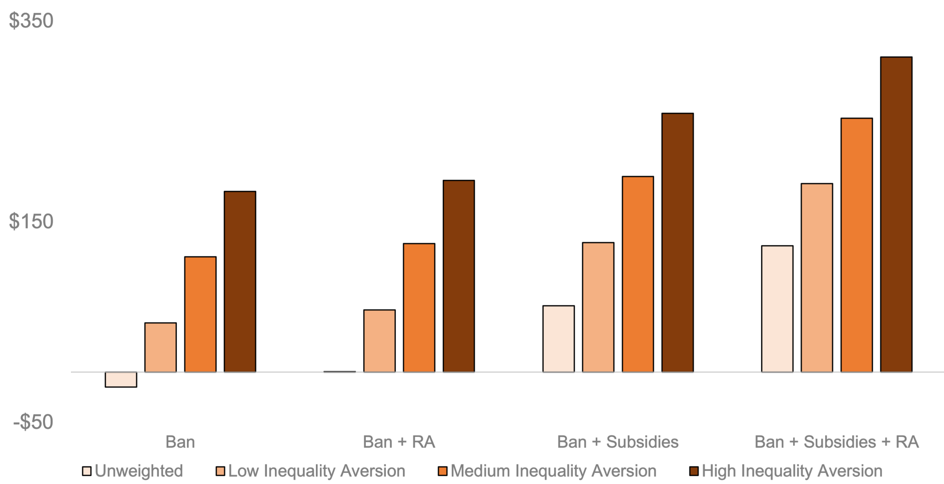

This figure shows the average consumer surplus change (in annual US per household) of implementing the ban of gender-based pricing with and without mitigation policies. The consumer welfare change is calculated with equal weights for the four health spending groups, low health spending inequality aversion, medium health spending inequality aversion and high health spending inequality aversion. Weights y_δ are computed as $y_\delta^{-\epsilon}/(\sum y_\delta^{-\epsilon}/4)$ for $\epsilon = 0.5$, $\epsilon = 1.0$ and $\epsilon = 1.5$.

1.6.4 Banning gender-based pricing under non-GR contracts

One of the predictions of the conceptual model in Section 1.2 is that the number of enrollees in the private market would be lower under GR contracts with protection against reclassification risk compared to non-GR contracts without that protection. Figure 1.13 below confirms this prediction by comparing the share of the simulation sample that is in the public option in (i) the baseline with the old risk-rating factors and under GR contracts, (ii) after the new risk-rating factors are implemented and under GR contracts, and, (iii) after the new risk-rating factors are implemented and under non-GR contracts. Particularly, as shown by the left bars of the figure, the share of enrollees in the public option goes from 43% to 48% after implementing the new risk-rating factors under GR contracts, but it goes down to 40% under non-GR contracts. These results are driven mostly by consumers that are on the margin between the two systems, such as low-income females and low-income males, who are more likely to either remain in the private market or to enter this system from the public option when firms are allowed to price plans independently.

What explains this result? The reason price sensitive consumers are able to choose the private sector at higher rates under non-GR contracts is that in this scenario prices of low-quality plans are lower as they do not have to cross-subsidize selection in high-quality plans. This can be seen in the upper panel of Figure 1.14, where I plot the distribution of prices for the cheapest plans in the market (less than $100 USD per month) for single young enrollees after implementing the ban under GR contracts and under non-GR contracts. By allowing firms to price plans independently, premiums of low-quality plans can better reflect the lower costs of

The results in this case are quite similar to the ones in Figure 1.11 because accounting for health care spending creates almost no change in the ranking of consumers in terms of income.

Figure 1.13: Share in public option with and without ban

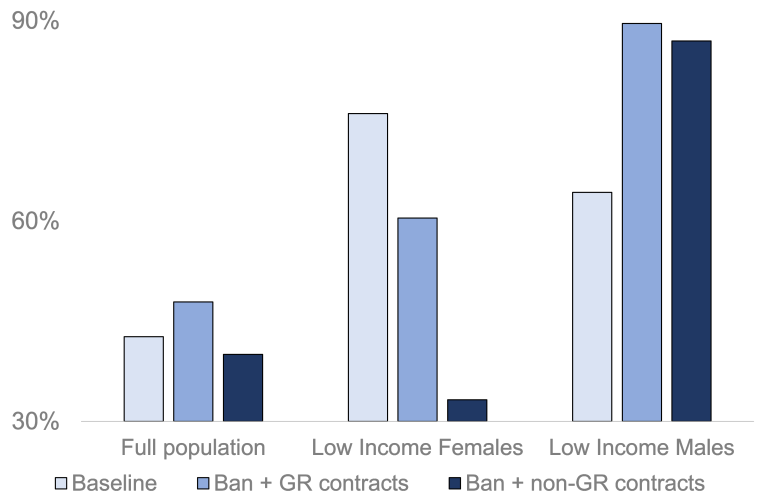

This figure shows the share of policyholders in the public option with and without banning gender-based pricing in Chile. Lighter bars represent shares before implementing the ban under GR contracts. Medium-light bars represent shares after implementing the ban under GR contracts that provide protection against reclassification risk. Darker bars represent shares after implementing the ban under non-GR contracts that do not provide protection against reclassification risk.

their enrollees. On the contrary, and as predicted in Section 1.2 and shown in the lower panel of Figure 1.14, premiums of high-quality plans are much higher in this scenario, with some of them even unravelling (*i.e.* no enrollee is willing to purchase them). Finally, in terms of consumer surplus, Figure D.13 in the Appendix shows that enrollees are better-off when the ban is implemented under non-GR contracts as more affordable plans enter the choice set of most consumers.[48]

It is important to emphasize that this result is not stating that protection against reclassification risk is not desirable. This analysis is not taking into account the long-term benefits that this protection provides to enrollees (Handel et al., 2015). What this shows is simply that, in the case of banning gender-based pricing, contracts that allow premiums to reflect enrollment composition in each plan are better equipped to limit the number of consumers choosing the public option than contracts that do not allow for this.[49]

1.7 Conclusion and Further Comments

Guaranteed renewable insurance contracts have the potential to mitigate reclassification risk without causing adverse selection on pre-existing conditions. However, adverse selection on gender will arise in these contracts if policymakers restrict insurers' pricing in favor of gender-based equity. This paper studies how a regulation banning gender-based pricing impacts the Chilean health insurance system, which is characterized by a private market offering GR contracts and a safety net public option.

[48] Unlike the result stating that the number of consumers in the private market should be higher under non-GR, the consumer surplus result is likely to be highly dependent on this particular empirical setting.

[49] For similar reasons, an important caveat of this analysis is that, if consumers understand and

Figure 1.14: Distribution of plan prices

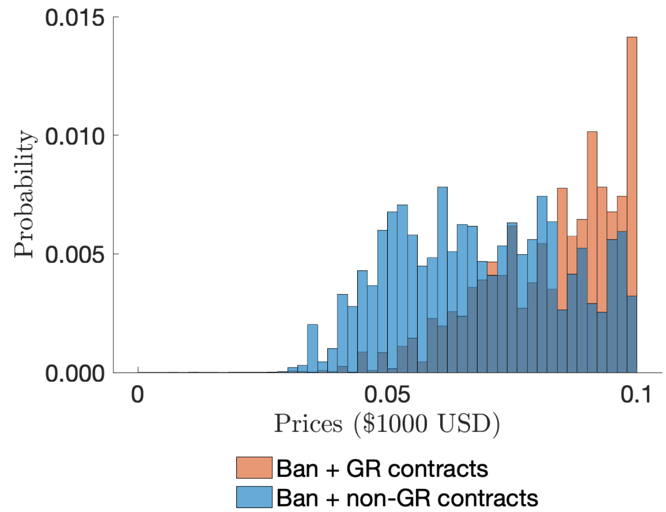

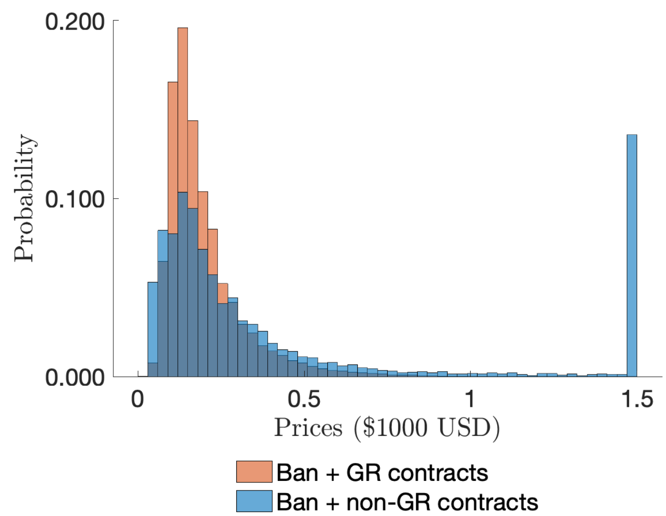

This figure shows the distribution of prices after banning gender-based pricing in Chile, only for cheap plans (less than $100 USD per month) in the upper plot and for all prices in the lower plot (plans above $1,500 USD per month are clustered at that price). Premiums are restricted for single young enrollees only. Red bars represent prices under GR contracts that provide protection against reclassification risk. Blue bars represent prices under non-GR contracts that do not provide protection against reclassification risk.

First, I present several stylized facts of the Chilean health insurance system under gender-based pricing. Female enrollees pay higher prices than males for the same plans in the private system, but they also spend more in health care. If women have higher willigness-to-pay for insurance, this will drive adverse selection after implementing the ban. In addition, for a market offering GR contracts, switching rates are non-trivial and they are explained, at least partially, by changes in premiums. Motivated by these facts, I estimate a two-stage discrete-choice demand model for plans in the private market using detailed administrative data from 2013 to 2016. The first stage of the model determines whether a policyholder will make an active choice in the market, and the second stage determines, conditional on making an active choice, which plan the policyholder will choose. After estimating the model, I simulate a ban of gender-based pricing in the private system.

Overall, I find that women benefit from the regulation and men are negatively impacted by it. Since in my sample gender is also correlated with income, I find that low-income enrollees benefit less than high-income enrollees. Subsidies that induce young men to remain in the private market are the most effective mitigation strategy to contain higher premiums as extensive margin adverse selection is the main driver of adverse selection. Lastly, under the ban, a market offering non-GR contracts would increase the number of consumers choosing the private market instead of the public option.

The findings of this paper are relevant to policymakers as recent research is considering whether GR contracts should be implemented in the health insurance individual market in the U.S. A limitation of this literature is that, for tractability, they ignore the possibility of adverse selection on dimensions beyond pre-existing conditions. However, insurance markets around the world are increasingly implementing restrictions on insurers' pricing in favor of gender-based equity. Therefore, to seriously consider whether to implement these contracts in the U.S., a better understanding of the interaction between these contracts and adverse selection on gender is required. This paper helps to fill that gap.

Regarding the lapsing evidence presented in Section 2.3.2, an interesting topic for further research is to explore in detail whether lapses in markets with GR contracts are explained by reasons beyond health shocks. The literature of insurance with long-term contracts assumes that lapsing only occurs by healthy consumers trying to find a lower price in the spot market. If, in practice, lapses are also explained by premium changes or income changes, researchers might question then the effectiveness of long-term contracts. Importantly, if insurers foresee this, they could potentially based their pricing strategies on these lapses, an issue that has been studied before in the term life insurance industry (Gottlieb and Smetters, 2021).

Finally, one important caveat in the analysis of this paper is that I do not observe consumers choosing between GR contracts and short-term contracts in the data. Instead, I simulate how the Chilean private market will respond to the ban under contracts with a temporary waiver that allows insurers to respond to compositional changes in each plan independently (until a new equilibrium is reached). This is how the short-term contracts available in the U.S. would respond to adverse selection. Quantifying willingness-to-pay for long-term contracts (on top of short-term contracts) is a challenging but important avenue for future research.

---

value the protection against reclassification risk provided by GR contracts, then their preferences might change as well under non-GR contracts without that protection. Therefore, here I am assuming that this change in consumer preferences would not affect plan choices drastically.

2 Chapter 2: Lapsing in Health Insurance Markets with Guaranteed Renewable Contracts

2.1 Introduction

Guaranteed renewable (GR) insurance contracts guarantee that the terms of a policy will not be cancelled or modified, even if the policyholder develops a medical condition. They are popular in insurance markets such as term life insurance and long-term care insurance. In theory, such contracts can mitigate reclassification risk—the exposure of individuals to substantial premium increases due to changes in health status—without causing adverse selection. The intuition is that consumers pay front-loaded premiums to guarantee affordable coverage in the future, regardless of any possible negative health shocks.

In standard models of perfectly competitive long-term health insurance markets with insurer commitment to a smooth price schedule and reclassification risk, individuals purchase a contract when young and healthy and rationally decide to stick with their plans in order to subsidize their future old and high-risk selves. In this environment, lapsing only occurs by healthy consumers trying to find a lower price in the spot market (*i.e.* reclassification risk, see Ghili et al., 2022 among many others). In imperfectly competitive markets without insurer commitment, however, there are multiple reasons that could lead to policyholders lapsing their contracts beyond reclassification risk. For example, consumers might be sensitive to income or premium fluctuations and lapse their contracts in response to them.[50] If insurers internalize these lapses, this might lead to lapse-supported pricing, in which premiums chang-

[50]Another reason for consumers to lapse their contracts is that consumers' preferences for plans might evolve over time. For example, a female policyholder might want to switch plans during her childbearing age in order to get access to hospitals with better maternity care (Atal, 2019).

ing, and consumers leaving their contracts, are a main component of insurers' profits (Gottlieb and Smetters, 2021).

This paper studies lapsing in the Chilean private health insurance market, a system characterized by the offering of GR contracts. I use detailed claim-level data from 2013 to 2016 to show that the annual switching rate in the market is high, especially compared to switching rates in health insurance markets offering short-term contracts, and that most policyholders lapse their insurance plans just a few years after signing their contracts, with less than 30% of them staying in the same plan after 70 months. Furthermore, I find that policies and lapse patterns predicted by standard theoretical models of long-term contracts are the opposite of those observed empirically in Chile. Finally, premiums increasing over time, and consumers lapsing their contracts because of those price changes, are a key determinant of insurers' profits, leading to a lapse-based insurance equilibrium.

The Chilean health care system is ideal for this type of study for at least two reasons. First, Chile has one of the very few health insurance markets, in addition to Germany, featuring GR contracts. By law, Chilean workers must choose between a public option (generally considered of low quality) or a private market with contracts that offer guaranteed renewability. In contrast to the typical annual short-term health insurance contracts available in the U.S., individuals in GR contracts can stay in their plans as long as they wish (*i.e.* one-sided commitment) and non-price characteristics of the contract are fixed over time. Prices can change but in a limited way. In particular, premium changes are community rated; that is, price adjustments over the lifecycle of a contract are independent of changes in policyholders' health status. Second, unlike other insurance markets with GR contracts, such as the German private health insurance market or the U.S. life insurance market, the Chilean regulator gives researchers access to unusually rich individual-level data, thus allowing for a detailed characterization of lapsing patterns in a health insurance market offering long-term contracts.

I begin by providing reduced-form evidence of lapsing in the Chilean private health insurance system. In particular, the annual lapsing rate in the market is almost 21%, a high switching rate considering that this is a market offering GR contracts. For example, the annual switching rate in U.S. private health insurance markets offering short-term contracts, and with a public option, is only around 10%. Moreover, I find that less than 30% of policyholders signing a new contract stay in the same contract after 70 months. This is a potential contradiction to the goal of long-term contracts, in which individuals are supposed to pay front-loaded premiums when young and healthy and stick with their policies for many years in order to subsidize their future old and high-risk selves.

These lapses could undermine the effectiveness of long-term contracts in health insurance markets, especially if the reasons policyholders are lapsing their plans are unrelated to positive health shocks. In line with this, I show that, among other possible reasons, premium changes and income changes are strong predictors of lapsing in the Chilean system. Specifically, in the case of premium changes, I find that they increase the probability of policyholders lapsing their contracts tenfold, and that, in response to them, consumers move to lower quality plans in order to avoid paying higher prices. Furthermore, these lapsers are more likely to be higher risk than stayers, which is in direct contrast to the predictions from the standard models of long-term insurance contracts.

Finally, I test the predictions of the lapse-supported pricing theory introduced by Gottlieb and Smetters (2021) and tested in the life insurance industry in the U.S. They argue that lapses in this industry are not explained by the rational models of long-term contracts (*i.e.* reclassification risk). Instead, these lapses are induced by the pricing strategies used by insurers, and they are big component of companies' profits. In line with their model, in Chile I find that most policies sold would lose money without premium changes and the corresponding lapses. In practice, the majority of policies are indeed profitable, with premiums increasing over time being a main determinant of insurers' profits.

In terms of policy implications, anecdotal evidence from local news suggests that insurers changing their prices over time is a common practice and that consumers are highly affected by these changes. Specifically, these premium changes, and the lapses induced by them, have led to a market in which consumers do not feel certain nor secure about their health insurance coverage in the future, which is the main purpose of designing a health insurance market offering long-term contracts. Possible solutions to get closer to the efficient and welfare improving contracts designed by Ghili et al. (2022) are for regulators to be more involved in how premiums are set initially, in order to help insurers set optimal front-loaded prices, and to make long-term contracts insure against income changes as well.

This paper contributes to the broad literature of long-term contracts in insurance markets. Recent studies highlight the potential benefits of GR contracts and ask whether they should be implemented in health insurance markets in the U.S.[51]

[51]In the case of policy research in the U.S., to name some examples, Cochrane (2017) and Pope (2020) advocate for long-term contracts to replace the current short-term contracts in the individual market, and Duffy et al. (2017) from RAND posit the question of whether the individual market could perform better under long-term contracts.

Ghili et al. (2022) charactize optimal long-term insurance contracts with one-sided commitment, as in Harris and Holmstrom (1982) and Hendel and Lizzeri (2003), and find that in certain scenarios, these contracts can achieve higher consumer welfare than ACA-like contracts. Similarly, Atal et al. (2020) show that GR contracts in Germany, despite not being optimally designed, obtain similar welfare outcomes as those in Ghili et al. (2022). These studies' favorable evaluation of long-term contracts are based on insurer commitment to an optimally designed smooth premium schedule and on assumptions such as the absence of income uncertainty, because income paths are flat or consumers have perfect foresight of their income paths. In the Chilean setting, however, consumers face income fluctuations and, because insurers adjust prices constantly, premium fluctuations.

Empirical studies of health insurance markets with long-term contracts are rare because few health insurance markets offer these contracts. Pauly and Herring (2006) show evidence of front-loaded prices in GR contracts in the individual market in the pre-ACA period. In the context of the small group market pre-ACA, Fleitas et al. (2020) document limited dynamic pass through of expected medical costs into premiums, and provide evidence that GR contracts indeed give protection against reclassification risk. Browne and Hoffmann (2013) study the German private health insurance market and find that front-loading in premiums generates lock-in of consumers. Furthermore, they document that consumers that lapse (*i.e.* switch contracts) are healthier than those who do not. Closest to this paper is Atal (2019), who considers contract lock-in in the Chilean private system. Existing studies assume perfectly competitive health insurance markets with insurer commitment to a smooth price schedule. This paper complements the literature by studying an imperfectly competitive health insurance system with frequent price adjustments, and the corresponding lapses.[52]

The remainder of the paper is organized as follows. Section 2.2 describes the main institutional details of the Chilean health insurance system and introduces the data. Section 2.3 documents lapsing rates in the private market and studies why policyholders lapse their plans. Section 2.4 provides evidence of lapse-supported pricing and section 2.5 discusses the implications of this pricing strategy in the market and potential solutions. Section 1.7 concludes.

2.2 Institutional Framework and Data

The insurance system in Chile combines public and private provision.[53] The safety net public option, FONASA, is a pay-as-you-go system financed by the contributions of affiliates and public resources. The private sector—operated by a group of insurance companies—is a regulated health insurance market. In 2015, FONASA covered 77.3% of the population and the private system covered 15.1%. The remainder of the population is presumed to be affiliated with special healthcare systems such as those of the Armed Forces or to not have any coverage at all.[54]

Workers and retirees have the obligation to contribute 7% of their wages to the public system, or to buy a plan that costs at least 7% of their wages in the private

[52]There is also a growing number of studies that examine why consumers lapse contracts in the life insurance industry (*e.g.* Fang and Kung, 2021 and Gottlieb and Smetters, 2021) and the consequences of premium adjustments over time in the long-term care insurance market (*e.g.* Aizawa and Ko, 2023).

[53]The details of the Chilean health care system have already been described elsewhere, in particular Duarte (2012), Atal (2019), Cuesta et al. (2019) and Pardo (2019). I draw from those papers heavily in this section.

[54]See Figure D.3 for historical market shares in each segment of the health insurance market in Chile.

system, with a cap of \$207 per month.[55] The two systems differ in many respects, including provider access, premiums, coinsurance structure, exclusions, and quality. Unlike the private sector, in FONASA there are no exclusions based on pre-existing conditions, nor pricing based on age or gender, and there is no additional contribution for dependents. As a consequence, the private sector serves the richer, healthier, and younger portion of the population (Pardo and Schott, 2013).

The private health insurance market is comprised of 13 insurance companies, which are classified into two groups: six *open* (available to all workers) and seven *closed* (available only to workers in certain industries). This paper focuses only on *open* insurers, which account for 96% of the private market. Contracts in the private sector are, for the most part, individual arrangements between the insured and the insurance company. A key feature of these contracts is that they offer guaranteed renewability, meaning that enrollees can stay in their health insurance plans as long as they wish. Furthermore, insurers cannot change the characteristics of these plans over time. Only the price can change but in a limited way in order to protect consumers from reclassification risk (see details below). Once a policyholder has been in a contract for one year, she may lapse her contract and switch to another company. Switching plans within an insurer is allowed at any time.

The monthly premium for individual i under plan j in year t, P_{ijt}, is a combination of a base premium P^B_{jt} and a risk-rating factor r_i so that:

$$P_{ijt} = P^B_{jt} \times r(enroll\ age_i, gender_i) \tag{11}$$

where $r(enroll\ age_i, gender_i)$ is the risk-rating factor, which is a function of age at

[55]All monetary amounts are measured in U.S. dollars using the exchange rate on December 2016

enrollment and gender. These factors are fixed over time as long as enrollees stay in their plans. For dependents, there is a similar $r(enroll\ age_i, gender_i)$ function and the full premium of the plan in that case is the base price P^B_{jt} multiplied by the sum of the risk-rating factors r_i of each member of the family. A couple of features of the market restrict the extent to which private firms can risk-rate their plans when individuals enroll. First, base premiums are set at the plan (and not the individual) level. Second, the r function is not individual-specific: each firm can have at most two r functions.[56]

Several features of the plan determine the base premium P^B_{jt}. A plan has two main coinsurance rates, one for inpatient care and another for outpatient care. Unlike in the U.S., plans do not include deductibles and out-of-pocket maximums. Additionally, plans offer either unrestricted open networks or tiered networks.[57] Hospitals in Chile cannot deny health care to patients, and therefore all consumers have access to all hospitals, although they may have zero coverage from their plan.[58]

Base premiums are indexed to inflation, and adjustments to the base premium in real terms can be made once a year. In March of each year, companies must inform the regulator of their projected premium increases for the year. Each company must also inform their clients (through letters) about these increases, justify their reasons

unless noted otherwise.

[56]These factors were mainly used to adjust premiums based on expected health care costs differences between demographic groups. In Figure D.2 in the Appendix I plot risk-rating factors for a company and expected costs by demographic groups in 2016. For more details regarding risk-rating factors, see Figueroa (2023).

[57]Unrestricted network plans provide the same coverage for all hospitals. Tiered networks offer differentiated coverage across sets of private hospitals, as PPO plans in the U.S.. Few plans offer restricted networks, as HMO plans in the U.S., and they are rarely observed in the data and not offered publicly. I do not consider them in my analysis.

[58]Other important characteristics of the plans are: a) Capitation scheme: Plans can either be capitated or not, b) Maternity-related expenses: Some plans do not have coverage for maternity-related expenses (in 2019 the regulator prohibited companies from selling these plans anymore).

for the changes, and offer alternative contracts to their clients that keep monthly premiums more or less constant but that often imply lower coverage.

Front-loaded premiums

Front-loaded premiums—premiums higher than concurrent expected costs—is one of the main features of insurance markets with long-term contracts. For example, Hendel and Lizzeri (2003) and Ghili et al. (2022) propose contracts that lock-in individuals with front-loaded premiums, but where the level of front-loading and insurance against reclassification depends on income paths to properly balance reclasssification risk and consumption-smoothing. Empirically, evidence of front-loaded premiums has been found in health insurance markets in the pre-ACA period (Pauly and Herring, 2006), in the German private system (Browne and Hoffmann, 2013), and in the Chilean private market between 2007-2009 (Atal, 2019). Importantly, lapses that are not explained by reclassification risk—consumers realizing that are healthier than expected and finding cheaper premiums conditional on plan quality—are against the rationale of front-loaded premiums and long-term contracts. This is because the purpose of these policies is that individuals pay high prices when young and healthy in order to subsidize their future old and high-risk selves.

I follow Atal (2019) and show evidence of front-loading in the Chilean private market by looking at the evolution of premiums relative to expected medical spending for single policyholders who stay in the same contract for 6 years (from 2012 to 2017). Let h_{it} be the expected claims in year t for individual i, and P_{it} the corresponding premium. I show that the ratio $ratio_{it} = \frac{P_{it}}{h_{it}}$ is decreasing in t.[59][60]

[59] Expected medical spending is calculated separately by age, gender and year using claim data. In unreported results, I find similar findings if I use instead realized medical spending.

[60] As Pauly and Herring (2006) and Atal (2019) argue, front-loading does not necessarily imply a decreasing premium schedule. Premiums can increase only to reflect the increase in the spot price

Still, decreasing markups is a strong test of front-loading, as even in its presence, markups could increase over time if individuals display enough inertia (as is often the case in health insurance markets, see Handel, 2013). In markets with consumer inertia, firms are expected to use an "invest-then-harvest" pattern for prices (*i.e.* start charging a low price and increase it over time, see Ericson, 2014). In the context of one-sided commitment, GR contracts combined with an "invest-then-harvest" strategy do not imply unambiguous price patterns. Intuitively, inertia relaxes the no-lapsing constraint that is needed to incentivize the healthier to stay. Therefore, firms can charge in period two a price that is above the actuarially fair premium for the healthy type. This increased revenue in period two is passed on to the first period in the form of lower premiums. Moreover, the evidence I provide is limited to the first 6 years of enrollment.

I test the hypothesis of decreasing markups using a panel of sampled individuals enrolled in the same contract from January 2012 to December 2017. Figure 2.1 displays the results. Overall, even though I only use 6 years of data, the figure shows that markups decrease over time as individuals stay enrolled in the same plan. Therefore, consumers lapsing their plans only a few years after enrollment are leaving money on the table as they are contributing front-loaded premiums to a plan that they will likely never use.[61]

Reclassification risk

Reclassification could occur if firms could adjust the base premium P^B_{jt} of any given

of the healthy individuals. Instead, front-loading means that the (expected) markup decreases as individuals stay in the contract. Since the theory predicts full insurance, there is no distinction between total cost and insurer cost. However, since individuals that stay in the same contract keep their coverage rates, the distinction is not relevant for testing the dynamics of either one relative to premiums.

[61] In Figure D.14 in the Appendix I repeat the same exercise but splitting policyholders by gender,

Figure 2.1: Front-loading in GR contracts

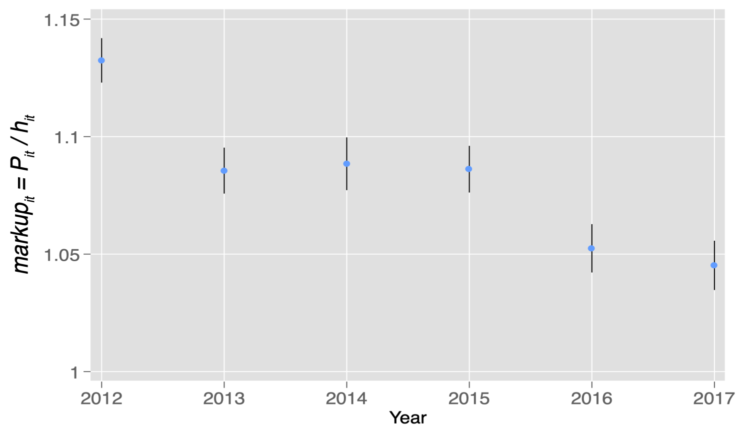

Notes: This figure shows a scatter plot of $ratio_{it} = \frac{P_{it}}{h_{it}}$, where h_{it} is the expected claims in year t by individual i, and P_{it} the corresponding premium. I use a panel of the sampled single policyholders enrolled in January 2012 and followed until December 2017.

plan j based on the health status of the pool of enrollees in j. However, the market regulation involves also a restriction that limits the extent of reclassification of individuals already in a contract: the increase in P_{jt}^B of any particular plan j of insurer k cannot be higher than 1.3 times the weighted average price increase of all plans of insurer k. Formally

$$\frac{P_{jt+1}^B - P_{jt}^B}{P_{jt}^B} \leq \frac{1.3}{|J_k|} \sum_{j' \in J_k} \frac{P_{j't+1}^B - P_{j't}^B}{P_{j't}^B} \tag{12}$$

where J_k is the set of plans of company k.

Figure D.4 in the Appendix suggests that this regulation works in limiting the extent of reclassification risk. For the season 2013/2014, which is representative of the pattern for all years in the sample, 5 out of 6 companies applied the same percentage price increase to all their plans, and the sixth firm increased its prices within a narrow window of 2.2% and 2.6%. Moreover, in Figure D.5 I plot the evolution of base prices by plan quality, showing that they all increase at similar rates within a company. This practice limits the correlation between individual health status and individual price increases, which implies limited reclassification.

Pre-existing conditions

Each new potential insured has to fill a "Health Declaration" before signing a new contract with a private firm. The companies are allowed to deny coverage of any pre-existing condition during the first 18 months of enrollment, or even to reject the prospective enrollee altogether. Enrollees are not usually required to fill a new declaration if they are switching to a lower quality plan within a company. Although there is no available data on the extent to which insurers deny coverage, anecdotal evidence and conversations with industry actors suggests that this is a regular practice.

Hospitals in Chile

The health care system combines public and private provision. The public hospital network is broader than the private one, with 191 public hospitals compared to 83 private hospitals in 2016 (Chile, 2016). The private and public sectors are mostly segmented. Private insurers primarily cover admissions to private hospitals, whereas the public option mostly covers admissions to public hospitals. In fact, 97% of private insurer payments are to private hospitals, whereas only 3% are to public hospitals (Galetovic and Sanhueza, 2013). An important feature of this market is price transparency, as consumers are often able to obtain price quotes before choosing a hospital.

In the calculation of coverage rates across plans, I focus on a particular geographic segment of the market. Specifically, I focus on the private hospitals in the city of Santiago, which is the largest health care market in the country and where more than a third of private hospitals and around half of the capacity is located (Galetovic and

Sanhueza, 2013).

2.2.1 Data

I exploit administrative data collected by the *Superintendencia de Salud* containing the universe of insureds in the private market for the period of 2013-2016 (Superintendencia, 2006).[62] Insurers must report data on individual claims to the regulatory agency. These data cover every health service provided to a private plan policyholder in 2013–2016, including financial and medical attributes along with consumer, plan and hospital identifiers. Additionally, I have data on all private plans offered during the period of analysis. This includes data on plans' company name, base prices, risk-rating function r, preferential networks, extra plan characteristics, availability in the market over time, and the date at which the plan was introduced in the market. Furthermore, I can match plans and their enrollees and observe basic demographics of policyholders and their dependents.

Even though plans in this market are differentiated by the coverage rate offered in each of the main private hospitals, those rates are not available in the data. Instead, an online platform called *QuePlan.cl* provided me with access to their administrative plans database, allowing me to observe the actual contract of each plan. Thus, I can extract the actual coverage rate of each plan in each hospital. In addition, *QuePlan.cl* gave me access to "plan scores", a measure of plan quality. I use this variable in the reduced-form evidence of section 2.3.[63] For more details regarding data construction,

---

finding similar results.

[62] In practice, I have data from 2007 to 2019. The reason I focus most of the analysis in the period 2013-2016 is that these were stable years in the market. In particular, the regulator did not pass any important mandate during this period, there were no mergers or bankruptcies among insurance companies, and data from one of the insurers is unreliable before 2013.

[63] The "plan score" is a standarized measure that goes from 0 to 10, where 10 represents a plan

see Figueroa (2023).

2.3 Reduced-form Evidence of Lapsing

In this section, I describe some important features of the Chilean private health insurance market. First, individuals in Chile lapse their contracts at a high rate, especially compared to switching rates in health insurance markets offering short-term contracts. Second, less than 30% of policyholders remain in the same contract after 70 months, a fact that is against the rationale of long-term contracts. Third, premium changes and income changes are an important reason of why consumers leave their plans. In the case of prices, lapsers tend to have higher medical risk than stayers. Thus, rational expectations models of reclassification risk and insurer commmitment to a smooth premium schedule face several challenges for being the primary explanation for these patterns of lapsing observed in the data.

2.3.1 Substantial Lapsing

The switching rates in the Chilean private system are remarkable, especially considering that this is a market with GR contracts. Specifically, as documented in Table 2.1, in 2016 almost 21% of policyholders lapsed their plans during the year, with 11.7% switching plans within their companies, 6% switching companies, and 9.5% leaving the private market in order to go the public option. The annual switching rates in 2015 and 2014 were 18.8% and 17.9% respectively. For comparison, the annual switching rate in Medicare Advantage in 2020 was 10%, which is a similar private health insurance market with a public option but offering short-term contracts (KKF,

with almost perfect coverage for the most expensive private hospitals of Santiago. As the score goes down, the coverage rate for private hospitals goes down as well.

2022). Furthermore, in Figure 2.2 I document tenure rates of policyholders in Chile. In particular, I look at policyholders that signed a new contract in March of 2011 and I follow them until December of 2016. The upper panel only considers leaving the insurance company as switching while the lower panel considers leaving the insurance plan as switching. In the latter case, I find that only around 29% of enrollees stay in the same contract after 70 months, which undermines, potentially, the effectiveness of long-term contracts. This is because the purpose of these contracts is that policyholders subsidize themselves by signing a contract and paying front-loaded premiums when healthy, and then staying long enough until their contracts protect them once they become old and sick. In the next subsection, I explore some of the reasons why people lapse their plans at such a high rate.

Table 2.1: Switching rates private market

Switching Rates (%)	2014	2015	2016
Within company switching	10.59	11.15	11.71
Company switching	5.10	5.52	5.95
Public option switching	7.66	8.01	9.54
Any switching	17.90	18.76	20.76

Notes: This table shows annual switching rates across policyholders in the data in 2014, 2015 and 2016.

2.3.2 Why Policyholders Lapse?

A high lapsing rate in the system can affect the functioning of long-term contracts in health insurance markets. This is true especially if the reasons policyholders are lapsing their plans are unrelated to positive health shocks, as predicted by reclas-

Figure 2.2: Tenure in GR contracts

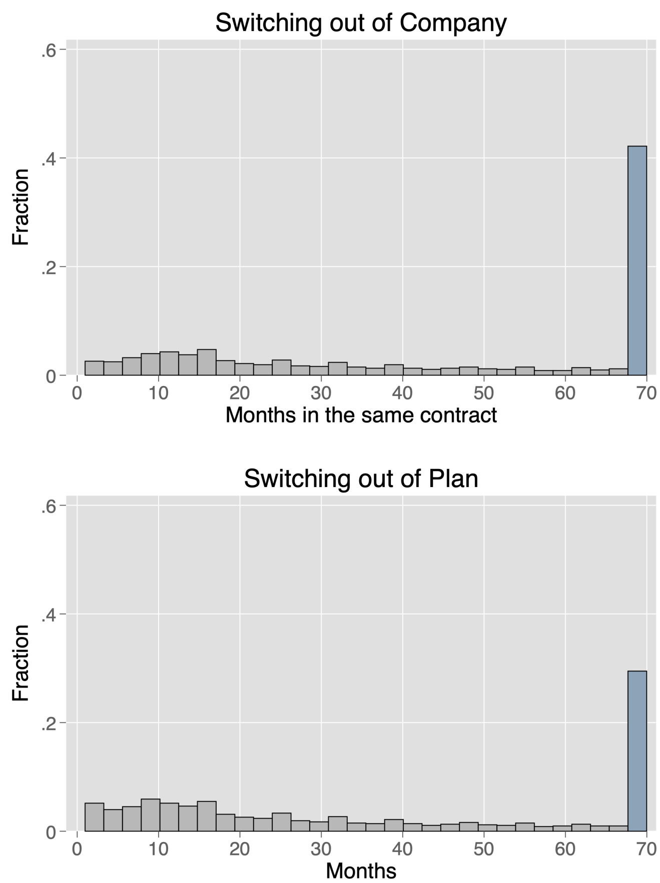

Notes: This figure shows histograms of how many months policyholders stay in their contracts after signing a new contract in March of 2011. The top figure shows this for the case where only switching out of he company is considered. The bottom figure shows the case in which any switching is considered.

sification risk theory, leading to a system in which individuals feel uncertain about their future coverage. Therefore, in this section, I explore empirically whether premium changes and income changes can explain lapsing in the Chilean private health insurance system.[64]

In the case of prices, I examine the relationship between premium increases and the probability of switching plans within a company.[65] Figure 2.3 shows the results of an event study regression where the dependent variable is a dummy equal to one if the consumer switches plans within an insurer, and the event is the month in which the contract was signed in the first place. In Chile, the signing month is the month in which premium changes are applied to each policyholder. Importantly, given the nature of the contracts, these premium changes are the sole reason for enrollees to switch plans in this month in particular (*i.e.* it is the only characteristic of the plan that is changing in the contracts). Furthermore, I also control for policyholder fixed effects and date (month-year) fixed effects, meaning that I am looking at the effect of the signing month on lapsing at the individual level, controlling for dates in which lapsing might be higher (or lower) than average. In order to have a clean panel of policyholders, I restrict the estimation sample to enrollees that do not switch insurance companies and that do not leave the private market and re-enter in later

[64] There are other reasons, besides premium changes and income changes, that could explain switching rates in this market and, hence, undermine the effectiveness of the system. For example, Atal (2019) documents that changes in preferences for hospital networks can induce policyholders to switch plans. Similarly, Figueroa (2023) finds that changes in family composition are also an important factor in explaining lapsing rates.

[65] I do not examine switching companies because this type of switching can only be done after being one full year in your current company. Thus, a spike in switching one year after enrollment could be attributed to premium changes or to the fact that policyholders could not switch before that. Furthermore, as highlighted by Table 2.1, within company switching is almost as twice as important as inter company switching. This is explained by the fact that to switch companies the policyholder must fill a new health declaration, risking then coverage denial. To switch plans within a company, a new health declaration is not needed in most cases.

Figure 2.3: Probability of switching plans due to price changes

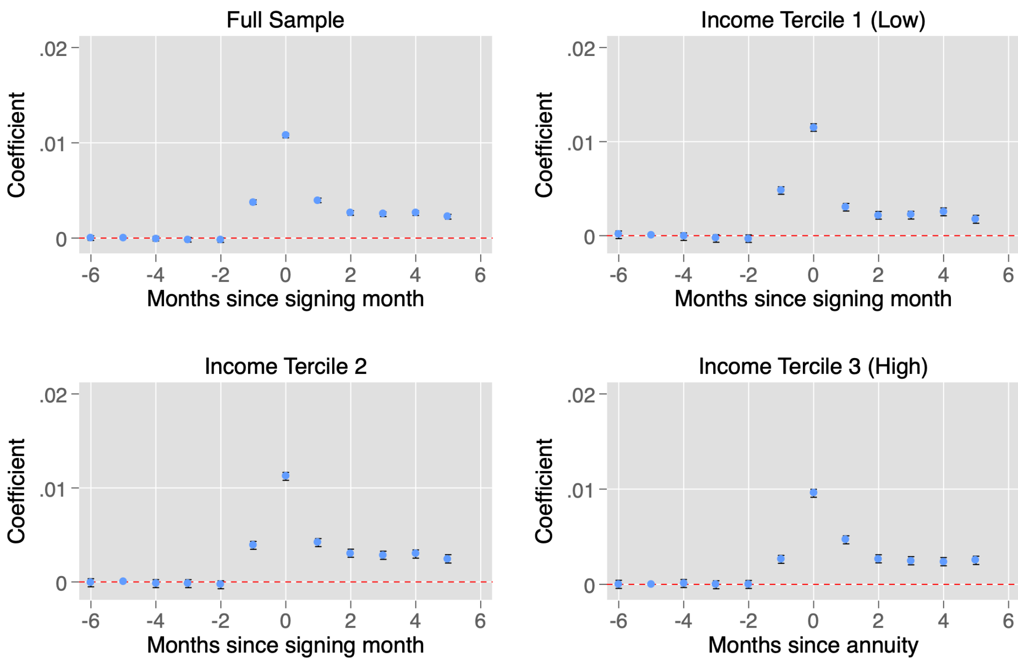

Notes: This figure shows an event study regression where the dependent variable is a dummy equal to one if a consumer switches plans within an insurer. The event is the month in which price changes are applied to health plans. Controls include individual fixed effects and date (month-year) fixed effects. I restrict the estimation sample to policyholders that do not switch insurance companies and that do not leave the private market and re-enter in later dates. Additionally, I drop individuals with zero or missing income at any month. This exercise is done on a 10% random sample of policyholders.

dates. This exercise is done on a 10% random sample of policyholders.

As documented in the figure, I find a large spike in the probability of switching plans in response to premium changes, and the effect appears consistent across income terciles. In terms of magnitude, the lapsing probability goes from 0.3% on average in any month of the year to 1.2% during the signing month. Furthermore, in Figure D.8 in the Appendix, I repeat the same exercise but with the dependent variable now being switching from the private market to the public option. I show that policyholders on the lowest income tercile are prone to leave the private sector in response to premium changes.[66]

What is the story behind these results? As noted in section 2.2, every year a company decides to increase prices, they must send a letter to their enrollees informing them of this change and offering an additional plan that keeps their premiums more or less constant. However, when prices for all plans are increasing at similar rates within a firm, the only way to do this is by offering a lower quality plan. To confirm this, I zoom-in into the enrollees that are switching plans during the signing month of the contract in Figure 2.3. Specifically, I compare the plan scores of their health insurance plans before and after switching. If these switches are being triggered by premium changes, then, on average, plan scores should go down and prices should remain more or less constant. These two trends are confirmed in Figure 2.4 across income terciles.

Under optimal and smooth front-loaded premiums and reclassification risk, pol-

[66] To confirm that these lapses are induced by premium changes, in Figure D.9 in the Appendix I repeat the analysis of Figure 2.3 but only for the 2011-2012 period and the 2013-2014 period. During the first period, insurers increased their prices by 6% on average, but in the latter period they increased premiums by only 1.9%. In line with the hypothesis of premiums being the main reason explaining lapsing, I find that in the 2011-2012 period the effect of the signing month on lapses is much larger than in the 2013/2014 period.

Figure 2.4: Plans before and after switching

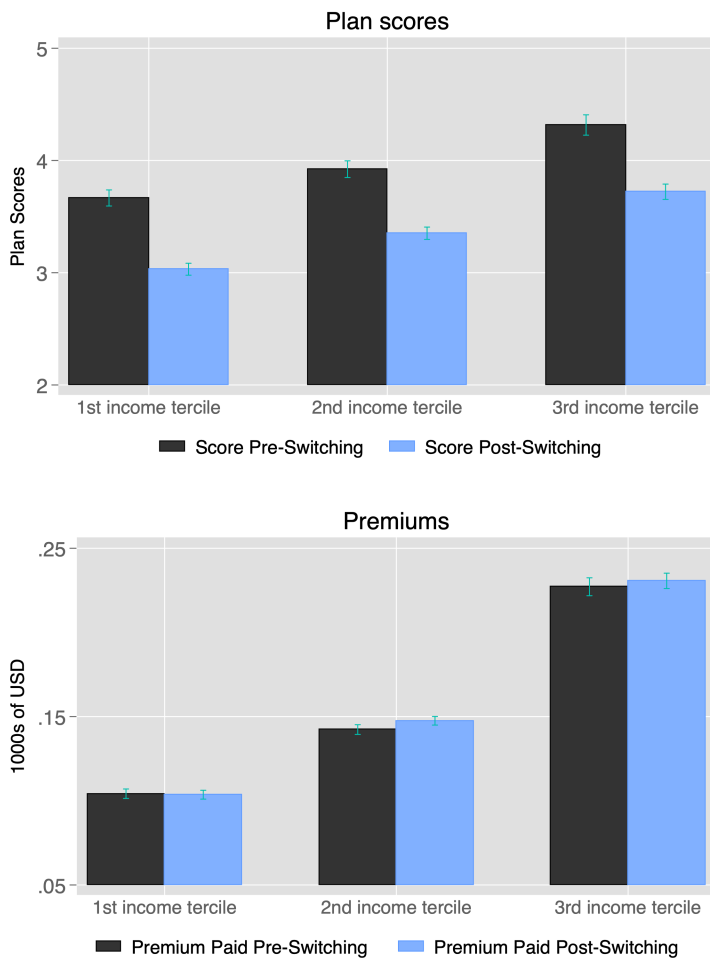

The upper figure shows the plan scores before and after switching during the signing month of the contract. The lower figure shows premiums paid by policyholders before and after switching during the signing month of the contract. Green lines report the 95% interval.

icyholders lapsing their contracts should spend less in health care than stayers. However, in a liquidity constraint environment where premiums are adjusted constantly, it is plausible that those individuals with higher health care costs are the ones leaving their plans in order to avoid paying higher premiums as well (Ericson and Sydnor, 2022). Motivated by this, I compare individuals lapsing during the signing month in a 12-month period (6 months before and 6 months after the signing month) in Figure 2.3 to those individuals that do not lapse during the signing month by estimating the following regression:

$$p(claim_{it} > 0) = \beta_0 + \beta_1 1\{Lapsers\}_{it} + \beta_2 X_{it} \tag{13}$$

where $p(claim_{it} > 0)$ is the probability that policyholder i incurs in a positive number of claims in the 6 months before and the 6 months after the signing month. $1\{Lapsers\}_{it}$ is a dummy equal to one if the policyholder lapses her insurance plan during the signing month in Figure 2.3, and X_{it} is a vector with additional controls, which includes insurer FE, year FE, age FE, gender FE, income tercile FE, region FE and policyholder FE. Table 2.2 displays the results, where each column is a different specification with different controls. Finally, to better match individuals to health care spending, I restrict the sample to single policyholders.

The results from the table show a clear picture. Policyholders lapsing their plans during the signing month are more likely to file at least one claim in the 12-month period surrounding the lapse. In particular, after controlling for insurer FE, those that lapse their plans in the signing month have a 6 to 7 percentage points higher probability of filling at least one claim in a 12-month period. This indicates that those individuals leaving their plans after their premiums rise are more likely to be high-risk, a result that is in stark contrast with reclassification risk theory, but in line with a setting where liquidity constraint individuals face premium fluctuations. To confirm this last hypothesis, in Table D.5 in the Appendix I find that the difference in the probability of filling at least one claim between lapsers and stayers is the largest for individuals in the lowest income tercile.[67]

Table 2.2: Regression - Lapsers and probability of positive spending

	(1)	(2)	(3)	(4)	(5)
$1\{Lapsers\}$	0.038	0.076	0.077	0.066	0.056
	(0.006)	(0.006)	(0.006)	(0.006)	(0.007)
Firm FE	No	Yes	Yes	Yes	Yes
Year FE	No	No	Yes	Yes	Yes
Demographic Characteristics	No	No	No	Yes	No
Policyholder FE	No	No	No	No	Yes
F-stat	36	726	510	464	17
Observations	278,463	278,463	278,463	278,463	278,463

Notes: This table shows the results of a regression where the dependent variable is a dummy equal to one if policyholder i fills a positive number of claims 6 months before and after the signing month, and the main independent variable is a dummy $1\{Lapsers\}$ equal to one if the policyholder lapsed her plan during the signing month in Figure 2.3. Each column is a different specification with different controls. Demographic characteristics include insurer FE, year FE, age FE, gender FE, region of residency FE, income tercile groups FE and policyholder FE. I restrict the estimation sample to single policyholders that do not switch insurance companies and that do not leave the private market and re-enter in later dates. Additionally, I drop individuals with zero or missing income at any month. This exercise is done on a 10% random sample of policyholders.

Finally, to better characterize health care spending patterns by lapsers, in Figure 2.5 I run an event study regression on policyholders that lapse during the signing month in Figure 2.3, where the dependent variable is a 3-month rolling average of health care spending, and the event is the month in which the contract was signed in the first place. The figure highlights an interesting pattern in which individuals start spending more on health care the month previous to lapsing, and then keep

[67] In unreported results, I repeat the analysis in Table 2.2 but running instead a logistic regression,

spending more in the months after leaving their plans. In particular, the average monthly spending in the first five months before lapsing is USD$64, which increases to USD$95 in the last 5 months after switching. It is important to notice that these results do not imply that markups gained by insurers from these lapsers will necessarily decrease after they switch because, as shown in Figure 2.3, most of them are moving to lower quality plans with lower coverage. In section 2.4, I link these lapses to companies' profits over time.

Figure 2.5: Health care spending for lapsers

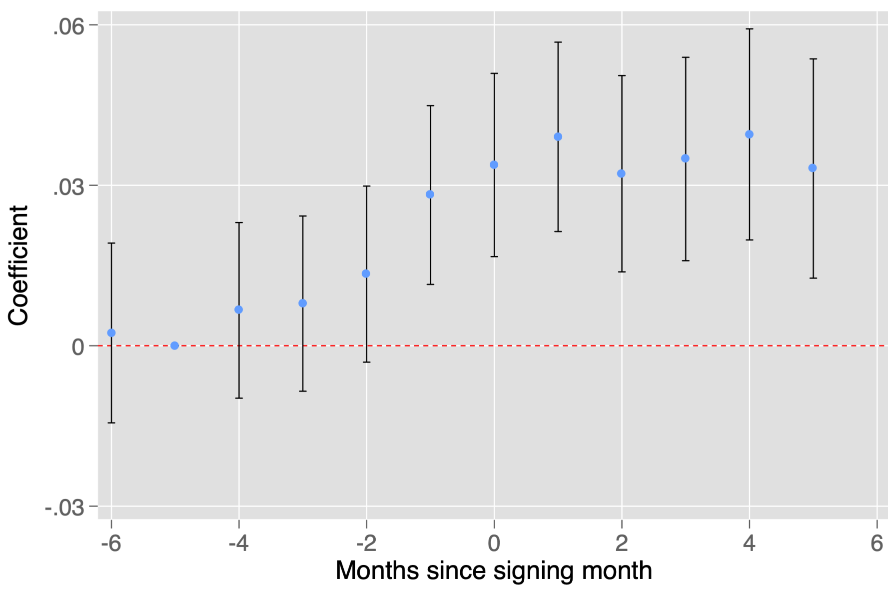

Notes: This figure shows an event study regression where the dependent variable is a 3-month rolling average of health care spending. The event is the month in which price changes are applied to health plans. Controls include individual fixed effects and date (month-year) fixed effects. I restrict the estimation sample to policyholders that switch plans because of price changes in Figure 2.3.

In Chile, prices increasing in policyholders' plans have a direct effect on income as premiums are deducted from policyholders' wages. Therefore, income changes more generally are a natural variable that could potentially lead to lapsing in the market. Recall that the theory of optimal long-term contracts assumes, for simplicity, that income plays no role in consumers lapsing their plans.[68] Thus, empirically studying whether policyholders respond to income changes by lapsing their plans is relevant in order to design effective health insurance markets offering long-term contracts. In particular, if the income of a policyholder goes up, and if she is healthy enough, she might decide that with her new income she would be better off with a higher quality insurance plan, so she could lapse and upgrade her plan. Similarly, if her income goes down, she might decide to downgrade her insurance plan in order to pay lower premiums. That is, whether consumers leave their plans after a change in income depends on how preferences for health insurance plans vary with income.

To empirically test this hypothesis, I run a regression of a dummy equal to one if the policyholder lapses in a particular year (within the company or to another company) and zero otherwise on a dummy equal to one if the policyholder is exposed to a change in income of at least 30% in that year (not accounting for premium changes). I use a sample of policyholders active in the private market from 2013 to 2016. The results are displayed in Table 2.3. Each column displays the results for a different specification with different controls. Standard errors are clustered at the policyholder level.

The results from the table confirm that policyholders do respond to income changes by lapsing their insurance plans. Specifically, from an average probability of lapsing of 9% in the sample, a change in income of at least 30% in a particular

finding similar outcomes.

[68] In theoretical models with long-term contracts, lapsing only occurs by healthy consumers trying to find a lower price in the spot market. This is because these models treat policies as securing a certain level of income (or consumption), which is feasible by assuming *ex-ante* known income paths.

Table 2.3: Regression results - Income changes and lapsing

Lapsing	(1)	(2)	(3)	(4)	(5)
$1\{\lvert\Delta Income\rvert \geq 30\%\}$	0.030	0.020	0.023	0.023	0.008
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
Age and Gender FE	No	Yes	Yes	Yes	No
Other Characteristics	No	No	Yes	Yes	No
Year FE	No	No	No	Yes	Yes
Policyholder FE	No	No	No	No	Yes
F-stat	1,583	556	530	517	35
Observations	1,557,660	1,557,660	1,557,660	1,557,660	1,557,660

Notes: This table shows the results of a regression of a dummy equal to one if the policyholder lapses in a particular year (within the company or to another company) and zero otherwise on a dummy equal to one if the policyholder is exposed to a change in income of at least 30% in that year. Each column is a different specification with different controls. The sample is composed by policyholders that were active in the private market throughout 2013 to 2016. Standard errors are in parenthesis and are clustered at the policyholder level. The mean of the dependent variable in the sample is 0.09.

year increases the probability of a policyholder lapsing her insurance plan by 2 to 3 percentage points, depending on which controls are included.

One caveat to these results is that unobservable characteristics of the policyholders might be correlated with income changes, which would bias my coefficients. For instance, maybe driven people are more likely to receive positive changes in income and also are more likely to search for new plans in the market. To control for this, in column (5) of the table I exploit the panel nature of the sample and I add policyholder fixed effects. That is, now I am identifying the impact of income fluctuations on lapsing by looking at how income changes for the same policyholder affect her probability of leaving her plan. Noteworthy, even with less variation in the independent variable, I still find a sizable effect of income changes on lapsing probability.

To summarize, I find that both premium changes and income changes induce lapsing in the Chilean private health insurance market, and that, in the case of premium induced switching, lapsers are higher risk than stayers. This is in contrast to reclassification risk theory and an optimal smooth premium schedule, but in line with empirical research in other industries offering long-term contracts. For example, Gottlieb and Smetters (2021) find that in the term life insurance market in the U.S., consumers not being able to forecast income shocks explain why most individual policies are terminated by the policyholder before the policies expire or pay a death benefit. Regarding price fluctuations, Aizawa and Ko (2023) document that adjustments to premiums over time is also common in the long-term care insurance industry. In the next section, I document evidence that lapses, particularly those explained by premium changes, are an important component of insurance companies' profits. This is relevant as policymakers designing health insurance systems with long-term contracts might want to consider that, in the absent of regulation, lapse-based insurance, or individuals leaving their plans in response to how premiums are set, could be a potential equilibrium in the market.

2.4 Lapsed-Supported Pricing

The previous section documented substantial lapsing rates in the Chilean private health insurance market, and that premium changes and income changes induce policyholders to lapse their contracts. As shown by Gottlieb and Smetters (2021), an equilibrium with consumers leaving their plans in response to how premiums are set in the market could emerge endogenously if these lapses are an important component of insurers' profits. In Chile, insurance companies are allowed to change prices from year to year as long as they do it at similar rates across all their plans in the market (*i.e.* no reclassification risk). A relevant question is then how profits depend on firms changing prices, and, thus, on policyholders leaving their plans. That is, does the prediction from the theoretical model of Gottlieb and Smetters (2021), which was used to explain lapsing patterns in the life insurance industry in the U.S., also hold in the Chilean private health insurance market?[69]

Anecdotal evidence from local news suggests that insurers increasing their prices over time is a common practice to keep their profits high and that consumers are highly affected by these changes. This has led to many policyholders suing their insurance companies to stop their premiums from going up. As an article from 2013

However, if income paths are uncertain, then insurance against these income shocks is needed as well, which is not provided by real applications of long-term insurance policies (*e.g.* the GR health insurance contracts offered in Chile and Germany).

[69]Firms in the U.S. long-term care insurance (LTCI) market, which offer similar GR contracts, adjust their premiums in a similar fashion. As Aizawa and Ko (2023) argue, they do so to transfer aggregate risks to consumers. In this section, I argue that in the Chilean health insurance market

states:

> *They [insurers] can get rid of expensive, old and sick beneficiaries, whom, faced with non-stop ruthless annual [price] readjustments, are forced to migrate to the public option. These enrollees have "enjoyed" the private system when they did not actually use it and lose it when they desperately need it... Currently, nearly 50,000 members apply for protection annually to avoid the rise in the base price. This number may seem small if one considers the universe of around 1.5 million affiliates, but it is constantly growing.*
> — *Ciper (2013a)*

The information of the article is in line with the results documented in Table 2.2, that is, that price changes affect mainly high-risk enrollees. Furthermore, this has led to a market in which consumers do not feel certain nor secure about their health insurance coverage in the future, which is the main purpose of designing a health insurance market offering long-term contracts. This was the case of a policyholder that was suing her insurance company to stop it from increasing her premium:

> *In her presentation before the judges, she argued that she was not in a position to bear the increase in the value of her plan or to forcibly migrate [to the public option]... because she is 57 years old and she suffers from two serious illnesses, chronic and incurable that can cause her death and that imply high monthly expenses in tests and medications. She also noted that she has a dependent child and that between 1999 and 2013 her income has decreased by half, but the value of her health insurance plan has increased threefold.*
> — *Ciper (2013b)*

To empirically investigate how important are these price changes in firms' profits, I simulate expected actuarial profits for each plan assuming that a 20 years old male

lapses also motivate these price changes because, in expectation, if individuals do not lapse and premiums are stable, policies are not profitable (see Figure 2.6). This is not the case in LTCI as policyholders in this market barely lapse.

individual signed a new contract in 2013 and decided to stay in that plan until age 80, without prices changing over time. This is how these markets should work under an optimally designed premium schedule (Ghili et al., 2022). The upper panel of Figure 2.6 shows the result of this exercise for a representative firm, with profits for each policy and each age of the policyholder calculated year-by-year. The lower panel calculates cumulative profits for each policy and each age of the policyholder. Expected costs are calculated as the average medical spending by age across active male policyholders in 2013 multiplied by the coverage rate offered by each plan.

The figure confirms an enormous reliance on companies increasing their prices over time in order to obtain positive profits. For example, the median plan from this representative firm would gain almost USD$500 in annual profits when this policyholder is 20 years old, but if he remained in the same plan until age 80, the company would lose around USD$2,500 annually in the final years of the policy. Importanty, this median plan would lose almost USD$24,500 of cumulative profits with this policyholder in the span of 60 years, and across all policies the company would lose USD$50 millions. Moreover, Figure D.17 plots year-by-year and cumulative profits for a 20 years old female. In this situation, prices do allow for small cumulative profits thanks to regulation that allows companies to charge higher prices to women, but not enough to compensate the losses from men.[70] This fact, coupled with men accounting for a much larger share of the private system, suggests that companies increasing their prices over time is a key part of their profits.[71]

What is the impact of these price changes in realized profits? In Figure 2.7 I plot

[70]For example, the median plan from a representative firm would gain almost USD$800 of cumulative profits with a female policyholder in the span of 60 years. This is much smaller than the USD$24,500 that the median plan would lose with a male policyholder.

[71]Figure D.15 and Figure D.16 in the Appendix repeat the same exercise for the six insurers, for

Figure 2.6: Expected actuarial profit

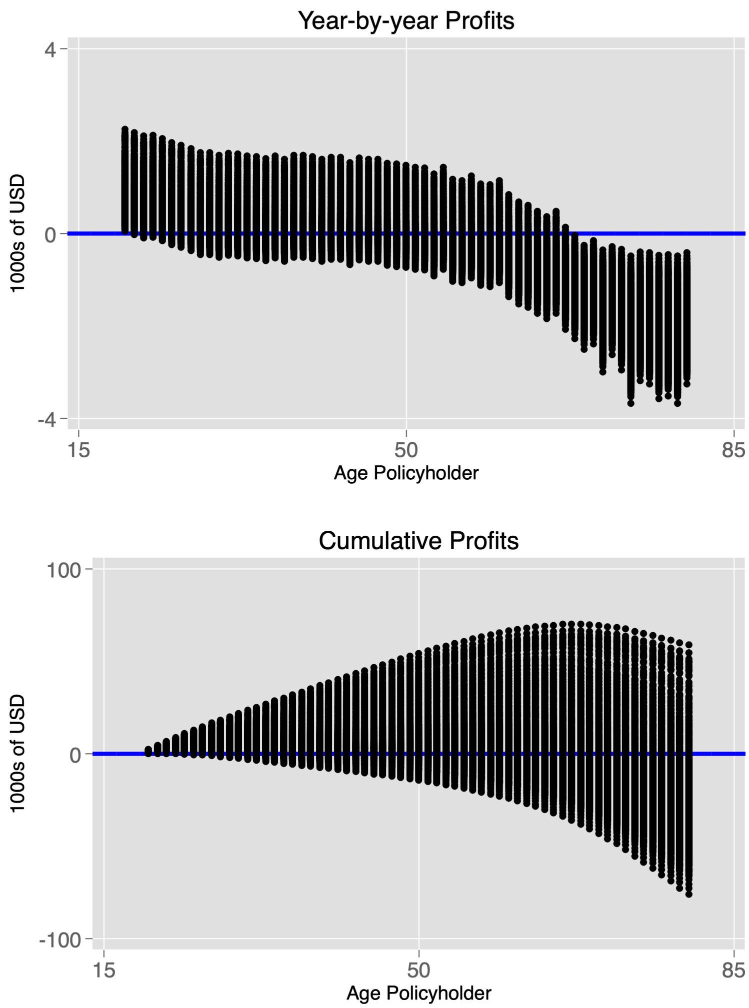

Notes: Each dot in the figure represents the expected actuarial profit for a particular plan of enrolling a male policyholder of a particular age. The upper panel computes annual profits for each age and each plan. The lower panel computes cumulative profits for each age and each plan. Data come from one representative insurance company. The price is calculated for a 20 years old male that signed a new contract in 2013. The cost is the expected medical spending by age across active policyholders in 2013 multiplied by the coverage rate of the plan. Profits are measured in U.S. dollars using the exchange rate on December 2013.

cumulative realized actuarial profits for policyholders that signed a new contract in a representative insurance company in 2013 and I follow them until December of 2016 (or until they lapse). For simplicity, in the figure I only look at single males of age less than 35 when signing the contract, but the results are very similar for the rest of the enrollees (see Figure D.18 in the Appendix). As can be seen in the figure, most policies accrue positive profits for the insurer, especially those that lapse early.[72] Importantly, on average, policyholders that remain in their contracts accrue high profits for the company. As I will show in Figure 2.8, this is due to them paying higher and higher premiums over time, and, on average, not becoming high-risk individuals. Finally, total profits for this company are USD$27 million over this period for these policies, with only 35% of individuals remaining in their initial policies by December of 2016. Thus, because they induce policyholders to lapse their plans or, at least, to pay higher premiums if they remain, price changes are a relevant component of insurers' profits.[73]

In order to analyse this strategy more closely, in Figure 2.8 I extend my data to 2019 for one particular company.[74] In particular, the upper left panel plots monthly average profits across active policies, the upper right panel plots the total number of active policies in each month, and the lower left panel plots total monthly profits across active policies. Finally, the lower right panel plots average medical spending by policyholders in each month. Black dots denote policyholders that remain in the same plan, blue dots denote policyholders that lapse their plans but remain in the

year-by-year profits and cumulative profits respectively, finding similar results.

[72]For this particular firm, 95% of policyholders that lapse their contracts during the first 12 months are profitable. That share decreases to 85% for policyholders that stay around 48 months in their policies.

[73]Figure D.19 repeats the same exercise for the six companies, with similar findings. Additionally, Figure D.20 plots realized profits for policyholders that signed a new contract in a representative insurance company in 2014, again finding similar results.

[74]Many changes happened after 2016 that make it unfeasible to analyse additional insurance

Figure 2.7: Realized actuarial profit

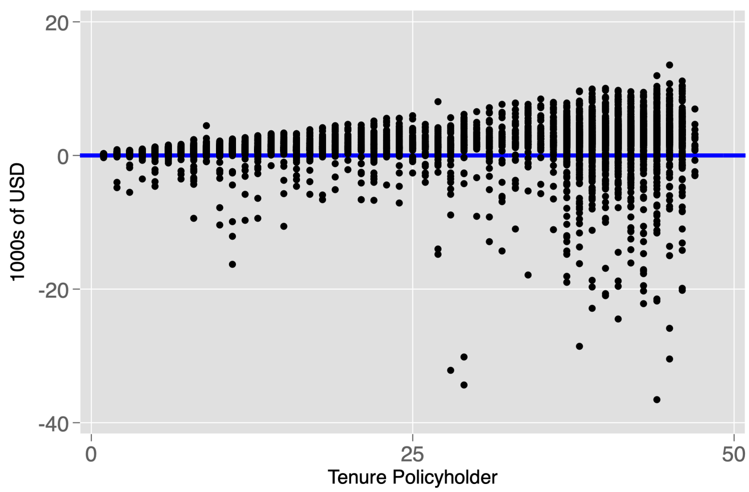

Notes: Each dot in the figure represents the accumulated realized actuarial profit for a particular single male policyholder of age less than 35 that signed a new contract in 2013 until he lapses. Tenure measures how many months the policyholder remained in the contract. Data come from one representative insurance company. Profits are measured in U.S. dollars using the exchange rate on December 2013.

same company and red dots denote policyholders that lapse their plans and leave the company.

The figure highlights four main points; first, average profits, and their variance, from policyholders that do not lapse slightly increase over time. Specifically, average (variance) profits in 2013 were USD$71 (USD$220), and they rise to USD$109 (USD$291) in 2019. This result might be counterintuitive considering that Figure 2.1 shows that premiums in this system are front-loaded. However, front-loading and increasing average realized profits can happen at the same time if the policyholders lapsing their plans early are high-risk (or will become high-risk with a higher probability). Figure 2.8 suggests that this is indeed what is happening, which is in line with the results from section 2.3.2. For example, average profits for lapsers that remain in the company in 2013 were USD$67 (USD$4 less than stayers), but in 2019 those profits were only USD$75 (USD$34 less than stayers). This is a combination of those policyholders paying lower premiums than stayers (they tend to lapse to cheaper lower quality plans) and also spending more in health care over time.

Second, the number of active policies in the market declines substantially over this period. From 2,400 active policies in the beginning of 2013, only 764 of them are active in December of 2019. Similarly, policies of lapsers that remain in the company increase over time. Third, total profits decrease over this period, which is explained by the last two points. Lastly, average medical spending for those that remain in the company but lapse their plans is much higher than for stayers (or even than movers that leave the company). For example, in 2019, lapsers that remain in the company spent on average USD$167 in health care, while stayers spent on average USD$125 (movers that leave the company spent USD$116). This is explained by the fact that if individuals want to switch companies, they have to fill a health declaration before joining a new insurer and they might get denied coverage. This is not required if the policyholder is switching plans within an insurer, especially if the switch is to a lower quality plan.

Figure 2.8: Realized actuarial profit 2013-2019 for one company

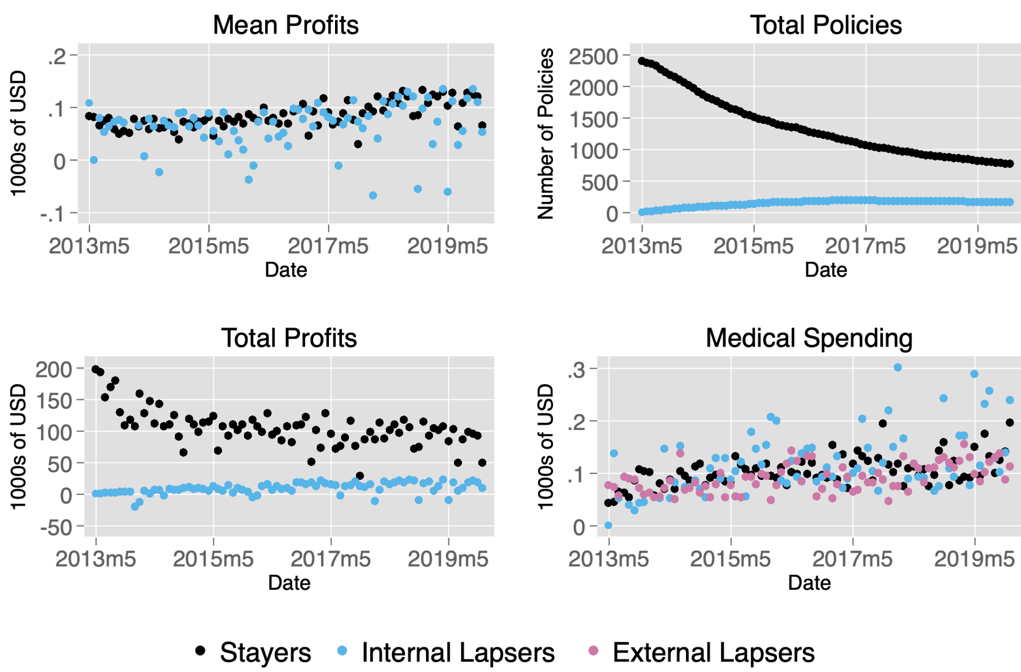

Notes: The figure looks at single policyholders that signed a new contract in the first quarter of 2013. The upper left panel plots monthly average profits across policies. The upper right panel plots the total number of active policies in each month. The lower left panel plots total profits in each month. The lower right panel plots average medical spending in each month. Black dots denote only policyholders that remain in the same plan. Blue dots denote only policyholders that lapse their plans but remain in the same company. Red dots denote only policyholders that lapse their plans and leave the company. Data come from one insurance company. Profits are measured in U.S. dollars using the exchange rate on December 2013.

---

companies. For example, one company filled for bankruptcy and two other companies merged in 2017. Also, data from other firms are unreliable after 2016 or they just do not have a large sample of policies in the market. For those reasons, I focus only on one company with robust and reliable information.

To summarize, policies in the Chilean private health insurance system are not profitable if individuals stay in their plans and premiums do not change. In practice, these policies accrue positive profits for insurance companies, which is explained by real prices increasing over time, and by policyholders lapsing their contracts. Importantly, insurers benefit from these lapses because these movers tend to be higher risk than stayers.

2.5 Policy Implications and Discussion

Section 2.4 shows both anecdotal and empirical evidence that premium changes are a relevant feature of the Chilean private health insurance market. This has led to a system with consumers lapsing their plans often and, hence, feeling uncertain about their prospective health insurance coverage in the future. This section discusses the implications of these market characteristics and possible solutions to address these issues.

One of the consequences of a system in which individuals are being induced to lapse their plans in response to premium changes is that a large proportion of policyholders in Chile are suing their insurance companies to stop prices from increasing in their contracts. For example, in 2016, 143,000 policyholders sued their insurers, with judges most of the time favoring insurees. At the same time, insurers spent more than USD$40 million dollars in these lawsuits, which has led to higher costs in the system and, thus, to companies to increase their premiums further (Libertad y Desarrollo, 2017).[75]

[75]By 2023, the national discussion has shifted to whether the market is viable in the long-term. Because of the higher costs of COVID-19 and, more importantly, because of a new policy that bans gender-based pricing and add restrictions to age-based pricing (see Figueroa, 2023 for details), insurers argued that they needed to increase prices over 20% in recent years. The regulator responded

Why are these contracts, that are effective and welfare improving in the theoretical models of Ghili et al. (2022) and Atal et al. (2020), among many others, not working effectively in Chile? One important difference is that, as noted in section 2.3.2, these models assume that income paths are flat or that consumers have perfect foresight of their income paths. In practice, income changes often and consumers normally cannot predict these changes. For example, premiums from health insurance plans in Chile are deducted from policyholders' wages, such that when an insurer increases prices, consumers' income goes down. Section 2.3 shows that these fluctuations lead to individuals lapsing their plans at high rates, even though, theoretically, these contracts are supposed to be for the long-term. Ghili et al. (2022) argue that one way to solve the problem of income paths being uncertain would be to make long-term contracts insure against income changes as well. That is, instead of long-term health insurance contracts, policymakers would need to design and implement long-term income contracts that provide protection against changes in income more generally. Of course, the implementation of such a contract in practice is not an easy task.[76] A similar, but maybe easier, solution would be for the government to provide subsidies in the scenario of big fluctuations in income such that policyholders do not need to leave their health insurance contracts.[77]

In the case of premium adjustments, and the lapses induced by them, they are

by prohibiting them to do so, arguing that such a price hike was not reasonable given the current situation in the country. As a consequence, in the last few years, companies have reported big losses and they are concerned that, at this pace, in the near future the private system will go bankupt and all their policyholders will be left uninsured or will have to switch to the low quality public option (La Tercera, 2022).

[76]Insurance against income changes is not provided by real applications of long-term health insurance policies (*e.g.* the GR health insurance contracts offered in Chile and Germany).

[77]In Chile this would likely be politically unfeasible as the private system provides coverage to the high-income population. The public option offers coverage to the low- and middle-class. Thus, subsidies to the high-class would not be a popular idea.

relevant because they make the market profitable for insurers, as predicted and tested by Gottlieb and Smetters (2021) in the life insurance industry. Specifically, as described in section 2.4, without prices changing, and consumers leaving their plans, especially those with high health care costs, most of the policies sold in the market would not be profitable. A way to solve this issue would be for the regulator to be more involved in how base premiums are set initially. That is, to help insurers set prices similar to the optimal ones designed by Ghili et al. (2022), such that real prices do not have to change over time for companies to make reasonable profits. Atal et al. (2020), for example, show that, in Germany, pricing regulation works well enough in designing front-loaded premiums that will cover expected costs, even if policyholders remain in the same plan for many years, and still leave room for profits. The Chilean private market then, or any country designing insurance markets with long-term contracts, could learn from the German experience in order to set front-loaded base premiums that would not require real prices changing over time.[78]

Additionally, dynamic pricing regulations (*i.e.* regulations that limit insurers' ability to adjust rates) are also an alternative to stop firms from increasing their premiums over time. For example, in the context of U.S. private long-term care insurance (LTCI), many states adopted new standards in their oversight of the LTCI industry in the early 2000s to deter rate increases for existing consumers. However, as studied by Aizawa and Ko (2023), there is a trade-off in that these regulations might increase consumer welfare by decreasing uncertainty about future rate increases, but they

[78]An important caveat is that in Germany there is not detailed data for the full private market. Hence, it is hard to assess with a high level of confidence whether policyholders in that country do not suffer from the same lapsing problems as policyholders in Chile. This is especially important as premium adjustments take place in Germany based on changes in health care costs (Browne and Hoffmann, 2013). Investigating lapsing in the German private health insurance market is an interesting area for future research.

might also induce insurers to exit from the market or charge a higher markup, which will adversely affect consumer welfare. Studying this trade-off in health insurance markets with long-term contracts is an important avenue for future research.

2.6 Conclusion

This paper studies the Chilean private health insurance system, a market characterized by offering GR contracts, and documents that there is substantial lapsing in this market and lapse-supported pricing. I argue that the rational model with optimally designed long-term contracts developed by Ghili et al. (2022), among others, cannot explain key observable features of this insurance market: the high annual switching rate, even compared to health insurance markets offering short-term contracts in the U.S.; the short tenure of most policyholders in their GR contracts; the lapsing response of consumers to price changes and income changes, with lapsers being higher risk; and, the fact that these price changes, and the corresponding lapses, are a key component of insurers' profits.

These features have led to a system in which consumers feel uncertain about their future health care coverage, which is the main purpose of designing a health insurance market offering long-term contracts. Regulating the way in which base premiums are set initially can fix the need to increase prices over time, while allowing contracts to protect consumers from income changes can prevent lapses in the market.

3 Chapter 3: Earthquakes and Brand Loyalty: A quasi-natural experiment to investigate brand loyalty under severe product unavailability

(Coauthored with Andrés Musalem and Carlos Noton)

3.1 Introduction

Brand loyalty and brand equity have received substantial attention in the marketing literature. Several papers have investigated the role of prices, advertising, and sales promotions as drivers of consumer choices and brand loyalty (e.g., Bronnenberg et al. 2008; Grover and Srinivasan 1992; Yoo et al. 2000). In terms of its value, it has been argued that brand equity may be associated with several possible benefits, one of them being a potentially lower vulnerability to crises (Ahluwalia et al. 2000; Dawar and Pillutla 2000). Accordingly, strong brands might be in a better position to recover from challenging incidents (Keller 2013, p. 522).

In this paper, we empirically examine whether the brand loyalty towards the top players in a market is affected by a severe supply chain disruption. More specifically, an earthquake in 2010 caused a shortage of the two leading beer brands for several weeks. We are interested in assessing whether the brands of these two top players are sufficiently strong to fully recover from this incident or whether this prolonged unavailability induces systematic changes in brand loyalty.

Since the disruption under consideration is related to product (un)availability, we present four essential differences relative to previous research investigating the effects of stockouts on buyer behavior. First, instead of considering stockout effects in the same or subsequent visit to a retail store, we investigate whether prolonged stockouts lead to systematic changes in preferences on the next shopping trips. We are particularly interested in studying whether leading brands can regain initial market share losses once product availability returns to normal levels. We note that other researchers have considered longer impacts but in different settings such as catalog purchases (Anderson et al., 2006) and online transactions (Jing and Lewis, 2011). Both catalog and online purchases are different contexts when compared to physical stores. In both cases, there is no possibility to physically examine substitutes. In addition, for online purchases, store switching costs are negligible, as opposed to physical stores which require additional trips. Hence, there is a very small cost to search for the unavailable product at other online retailers. Both differences could potentially make substitution towards other brands carried by the same physical store more likely.

Second, we consider stockout events that lasted several weeks, as opposed to the shorter events observed in prior work (e.g., Vulcano et al. 2012, Conlon and Mortimer 2013). As opposed to an isolated event, prolonged unavailability might be more likely to induce loyal customers of the top brands to substitute these purchases with other brands. As discussed later in this section, this product trial may lead to learning and future purchase consideration with potential systematic preference changes and consequences for subsequent trips.

Third, we provide evidence of a consumer learning mechanism explaining the observed systematic changes in consumer preferences. In particular, we can identify consumers who tried some of the small brands for the first time during this period of prolonged product unavailability. For this set of first-time consumers, we quantify the extent to which they remained purchasing these small brands even after product availability was back at normal levels.

Fourth, to identify causal effects, we rely on a quasi-natural experiment that led to the absence of the top beer brands from the shelves. The nature of this supply shock implies that the treatment is independent of demand shocks and because it could not be unanticipated by consumers it ensures no planned stockpiling. The unexpected supply shock is an important distinction since most of the literature studying product unavailability relies on stockout episodes primarily driven by the observed demand instead of an exogenous event. Given the nature of the scarcity and that it affected all retailers, the stockout episodes were not informative to consumers about the quality of the affected products or the retailer's assortment. We use loyalty card data to observe consumer-specific exposures to product-specific stockouts at different stores, causing a considerable variation in the severity of the stockout-treatment across products and consumers. This exogenous variation allows us to study whether these prolonged stockouts changed purchase behavior even after leading products were available again. Furthermore, the earthquake that led to this product's unavailability did not imply price changes. Cavallo et al. (2014) studied online prices and found compelling evidence that most prices remained unchanged in Chile after the earthquake in 2010. Besides other reasons, the Chilean law considers illegal a price increase after a significant catastrophe, and most retailers choose not to change prices. Moreover, national statistics indicate that employment and economic activity were only affected for a brief period and then showed a speedy recovery due to the significant fiscal expenditure to rebuild public infrastructure.

Since our study focuses on the ability of top brands to recover from a severe disruption, it is also related to the empirical research on the interplay between the origins of brand loyalty (Bronnenberg et al., 2019; Dubé et al., 2010; Horsky et al., 2006) and consideration sets (Bronnenberg et al., 2016; Nedungadi, 1990; Roberts and Lattin, 1991). Our unique and novel data from a quasi-natural experiment provides an exogenous change in availability that allows us to identify various drivers of preferences while accounting for individual unobserved persistent heterogeneity and state dependence. Our evidence suggests that the prolonged stockouts changed the consideration set for a substantial share of consumers, who became aware or learned about competing products with long-lasting consequences in equilibrium market shares. Thus, our evidence is consistent with Bronnenberg et al. (2021), who emphasize the importance of product availability for brand loyalty relative to other arguments in the US market of craft beers.

Our analysis focuses on 5,668 frequent buyers of the top leading brands and considered 21 weeks before the earthquake, seven weeks of frequent stockouts, and 16 weeks when the top products gradually became available on shelves again. We find that 6 percent of the most frequent buyers of leading brands persistently stopped purchasing them after the shortage. Moreover, many consumers tried the less popular brands for the first time (in our data) when the top brands were unavailable, and a substantial fraction of these consumers did not switch back to the top leading products. Overall, this analysis indicates that the leading brands only partially recovered their pre-stockout market shares even months after this severe shortage.

A priori, these findings could potentially be explained by other factors different from product availability, such as price changes. To evaluate and quantify the relevance of alternative mechanisms, we estimate a discrete choice model that incorporates the effect of the stockout exposure on choices in the post-treatment period.

In particular, we estimate a random coefficients logit model accounting for prices, state dependence, seasonality, availability, and unobserved heterogeneity in preferences (Dubé et al., 2010; Heckman, 1981). Thus, the discrete choice model allows the valuations of leading brands to be permanently affected by the degree of stockout exposure each consumer faced. We also include a state dependence term to distinguish between transitory and more permanent changes in purchasing behavior. We estimate the demand model using Bayesian methods accounting for unobserved heterogeneity in preferences (Rossi and Allenby, 2003).

Our key finding is that, after controlling for differences in prices, state dependence, seasonality, and product availability, the smaller brands systematically increase their valuations (and market shares) at the expense of the top brands among those consumers who experienced more significant exposure to stockouts. Overall, we observe that the stockout treatments negatively affected the leading brand valuations, leading to decreased market shares in the post-treatment period. We use our structural estimates to compute the counterfactual market shares after the stockout treatments and the price discount needed to offset the adverse stockout effects. Specifically, we find a reduction of more than 5% in the market shares of leading brands, with the most affected products losing between 14% and 21% of their shares. On the other hand, the small brands gain more than 5% in market share. We also quantify the market share losses of an additional week of shortage to shed light on the optimal resources to prevent stockouts. Finally, we characterize the first-time purchasers' preference parameters of small brands to shed light on potential mechanisms at play.

We interpret our estimates as evidence that removing top products from the stores made consumers aware of or willing to learn more about competing products that became their top choices for a significant share of these first-timers. The empirical study of consideration sets is remarkably challenging as endogenous consideration sets typically preclude researchers from disentangling whether consumers have a strong taste for the leading brands or have not explored enough for competing products (Roberts and Lattin, 1991). Ideally, the identification of consideration sets would rely on an exogenous change in product availability that is uncorrelated with taste shock and (perceived) product quality. Our setting is consistent with that ideal scenario since our stockouts are exogenous and unanticipated. Another common difficulty in most settings, including ours, is that consideration sets are unobservable as researchers do not observe which products are being inspected by consumers during their shopping trips. Nevertheless, we can show that the weekly average of first-time consumers of non-top products grew substantially during the stockout period, suggesting that the quasi-experimental shortage enlarged their choice set. The excellent match value of the initially unknown products implied that, at least for a subset of consumers, the new choices remained preferred over the leading brands after these prolonged stockouts.

We discuss whether our results could be explained by other brand preference mechanisms, such as gradual or instantaneous customer learning of new products, switching costs, advertising, habit formation, peer influence, or evolving quality beliefs (Bronnenberg et al., 2019). Our results imply that the observed market share changes are primarily driven by the first-time purchasers of small brands who tried those products only after being exposed to the leading brands' unavailability (Ching et al., 2013; Erdem and Keane, 1996; Shin et al., 2012). Since the first purchase reveals most of the uncertainty about beer's match value, we see trying a product for the first time equivalent to one-shot learning. Furthermore, we track these first-timers' purchases and verify that their new preferences persist even after the stockouts are over. In addition to one-shot learning, we see the gradual learning hypothesis as a complementary force. However, incremental learning requires a greater level of product quality variability that seems somewhat limited in the beer industry relative to other sectors like, for instance, restaurants.

Other alternative mechanisms seem less relevant in our setting. First, leading brands have substantial incentives to recapture the market through massive advertisements, so it is unlikely that small brands' marketing campaigns could explain some consumers persistently remaining away from the leading brands. Second, switching costs, which are time-invariant in the supermarket industry, could not support our findings either.

The rest of this paper is organized as follows. Section 3.2 describes the data and the market. Section 3.3 provides a statistical analysis of the effect of the leading brands' unavailability on consumer purchase behavior. Section 3.4 presents our structural econometric model and results. Finally, Section 3.5 summarizes our main findings and briefly discusses opportunities for future research.

3.2 Empirical Setting

The beer market in Chile is highly concentrated, as it is often the case worldwide (Adams, 2006). CCU is the largest supplier accounting for over 70 percent of the beer market and produces the two leading brands in the market: Cristal and Escudo. We describe next the data used to measure customer behavior and characterize the shopping environment.

3.2.1 Description of the transactional data

We use loyalty card data from the second largest big-box supermarket chain in Chile, covering 64 stores in Santiago's metropolitan area. The point-of-sale (POS) individual-level data include quantities and prices paid for each stock keeping unit (SKU) within each transaction involving the beer, water, and soft drink categories. These shopping baskets account for a large number of our relevant consumer trips. We have access to panel data since the retailer's loyalty program identifies transactions where the same loyalty identification number was provided. We note however that most customers belonging to the same household use a single number to accumulate loyalty points at a faster rate. Hence, we consider our panel data to be at the household instead of at the individual consumer level. According to the retailer, purchases of brand loyalty card members account for about 80 percent of its total revenues.

Within this market, *Cristal* and *Escudo* are the top two leading brands. We focus on their most popular formats and group their SKUs into six alternatives: Cristal one-liter bottles (1000cc), Cristal individual cans (350cc), Escudo one-liter bottles (1000cc), Escudo individual cans (350cc), other SKUs from Cristal, and other SKUs from Escudo.[79]

The recorded transactions took place between early October 2009 and late July 2010. This period includes 21 weeks before the earthquake on February 27th (labeled as the pre-treatment period), 7 weeks immediately after the earthquake where fre-

[79]We combine returnable and disposable bottles into the same alternative since their prices and content are identical.

quent stockouts were observed (treatment period), and 16 weeks when the availability of top brands was gradually restored (post-treatment period). Figure 3.1 illustrates the start and end dates and labels of the different periods that we use in our analysis. The full sample contains 28,005 households who purchased any beer products at least ten times during the pre-treatment period. The selected consumers made 586,989 beer transactions in the pre-treatment period and 244,622 transactions during the post-treatment period. The average consumer spent approximately 21 dollars and purchased 10.18 items per visit.[80] A store in our sample generated, on average, approximately 2,300 daily transactions including a product from one of the four categories in our data set. Variation in the total number of transactions and revenue across stores reflects differences in store size and location.

Figure 3.1: Timeline

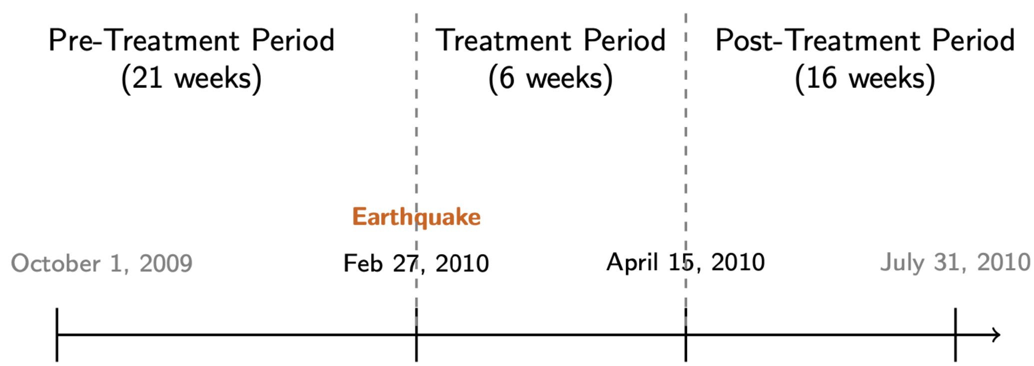

We focus on the sub-sample of consumers most loyal to the leading brands. Hence, we consider 5,674 households having at least ten purchase events of the top brand products in the pre-treatment period. This sub-sample made 169,986 beer transactions in the pre-treatment period and 71,845 transactions during the post-

[80]Amounts in US dollars, using the average exchange rate for that period.

treatment period. This reduction in purchases is consistent with an expected seasonality as the Fall starts in mid-March in the Southern hemisphere.

Table 3.1 presents summary statistics for these frequent buyers of the leading brands. The table shows the average price, the percentage of trips purchasing each product combining the pre- and post-period data. It also presents the market shares before and after the treatment period. Panel A displays the figures for the leading brand in various formats, while Panel B shows the summary statistics for the small brands. The fractions of trips (or incidence rates) are similar to the market shares, indicating that consumers of different brands buy similar quantities. Finally, excluding bottles, there are small price differences among leading and small brands, since most price gaps are caused by one-liter bottles being considerably more expensive than individual cans.

Table 3.2 shows detailed summary statistics for prices (Panel A), value market shares (Panel B) and incidence rates (Panel C) for each product across the three periods. Panel A shows that prices did not suffer significant changes during the three episodes, consistent with the fact that the Chilean law forbids abusive price increases after catastrophes like earthquakes.[81] Instead, Panel B and C show changes in market shares and incidence. We see how product shortage during the treatment period creates a substantial decrease in shares and incidence among Escudo's bottle and can products in that period, which as we will show later exhibited frequent stockouts. Then, the post-treatment market structure resembles their pre-earthquake configuration. However, we observe that leading brands did not quite reach their initial market shares. In relative terms, small brands gained a sizable increase in

[81] Notice that even though prices decline for some products in the post-treatment period compared to the pre-treatment period, consistent with seasonality, these prices changes are small for most

Table 3.1: Summary statistics for frequent buyers of leading brands

Panel A: Leading Brands	Average Price (US Dollars) (1)	Trips (2)	Market Share (Pre-Treatment) (3)	Market Share (Post-Treatment) (4)
Cristal (1L bottle)	1.26	9.5%	10.99%	6.98%
Cristal (350cc can)	0.47	19.1%	18.03%	17.67%
Escudo (1L bottle)	1.25	11.3%	11.82%	8.24%
Escudo (350cc can)	0.47	23.4%	23.64%	25.20%
Other Cristal	0.60	4.6%	5.50%	5.23%
Other Escudo	0.63	5.0%	7.04%	8.20%
All Escudo and Cristal			77.02%	71.52%

Panel B: Small Brands	Average Price (US Dollars) (1)	Trips (2)	Market Share (Pre-Treatment) (3)	Market Share (Post-Treatment) (4)
Baltica (350cc can)	0.39	3.0%	1.61%	2.81%
Becker (350cc can)	0.41	2.7%	1.83%	3.10%
Stella Artois (354cc can)	0.65	0.8%	0.83%	1.39%
Heineken (350cc can)	0.70	2.9%	3.66%	3.96%
Royal Guard (350cc can)	0.62	1.3%	1.62%	2.12%
Other Beers	0.87	16.3%	13.46%	15.09%
All Small Brands			22.98%	28.48%

No. of households	5,674	Av Trips per household	33.9
No. of Stores	64	Av Top Brand per household	25.9

Notes: Column (1) shows the average price for each product, Column (2) shows the percentage of purchases for each product, conditional on a beer purchase combining data from the pre- and post-treatment periods. Column (3) and (4) are the sales market shares before the Treatment period and after the Treatment period respectively. We only consider the 5,674 households that have at least ten beer transactions of the Leading Brand beers within the initial 21 weeks of data.

market shares.

Based on the summary statistics, we observe a 5.5 percentage point decline in the combined market share for the Cristal and Escudo brands after the stockout treatment period (see Table 3.1, Panel A). This noticeable decrease is mostly driven by households who are frequent buyers of the leading brands (Cristal and Escudo). In contrast, we find a smaller reduction of their market shares in the full sample.[82]

3.2.2 Identification of Stockouts

As mentioned above, CCU is the dominant beer producer in Chile and owns two major bottling plants. The closest plant to Santiago suffered severe damages after the February 27, 2010 earthquake. Because of this disruption, there was a substantial shortage of CCU's leading beer brands in the forthcoming weeks.

Table 3.3 presents the number of stores that faced high and low frequencies of stockouts per week. We consider a product to be out-of-stock on a given day and store if no sales are observed. The suggested stockout measure could be misleading for products infrequently sold (slow-moving products). However, for leading brands in fast-moving product categories, as it is the case in our setting, our measure should provide a good approximation of product availability.

We observe substantial variance of stockouts across products and periods, with the treatment period (i.e., between February 26th, 2010, and April 15th, 2010) exhibiting stockout episodes more frequently. Escudo (1L bottle) was the most heavily affected by the factory disruption, while Cristal (350cc can) remained unaffected. Also, the data show that the production shortage impacted more severely the bot-

products (around 1% or 2%) and, importantly, we control for prices in the structural model.

[82] Table D.6 and D.7 presents the summary statistics for the entire sample of 28,005 households,

Table 3.2: Summary statistics of prices and market shares

Panel A: Prices	Pre Treatment			Treatment			Post-Treatment		
	Mean	p5	p95	Mean	p5	p95	Mean	p5	p95
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Cristal (1L bottle)	1.29	1.15	1.62	1.23	1.16	1.55	1.29	1.15	1.53
Cristal (350cc can)	0.50	0.44	0.59	0.48	0.44	0.56	0.47	0.43	0.55
Escudo (1L bottle)	1.31	1.14	1.71	1.31	1.17	1.62	1.25	1.13	1.59
Escudo (350cc can)	0.50	0.44	0.59	0.50	0.44	0.59	0.49	0.43	0.58
Other Cristal	0.65	0.51	0.77	0.66	0.53	0.73	0.67	0.52	0.75
Other Escudo	0.75	0.53	1.00	0.71	0.53	0.96	0.74	0.52	0.95
Baltica (350cc can)	0.43	0.37	0.50	0.41	0.37	0.48	0.41	0.35	0.47
Becker (350cc can)	0.46	0.40	0.56	0.47	0.41	0.54	0.42	0.37	0.53
Stella Artois (340cc can)	0.75	0.63	0.94	0.76	0.68	0.90	0.65	0.54	0.88
Heineken (350cc can)	0.72	0.67	0.79	0.74	0.68	0.81	0.72	0.66	0.80
Royal Guard (350cc can)	0.68	0.57	0.82	0.66	0.60	0.78	0.66	0.58	0.77
Other Brands/Formats	1.02	0.42	2.06	1.05	0.41	2.17	0.99	0.40	1.94

Panel B: Market Shares	Pre Treatment			Treatment			Post-Treatment		
	Mean	p5	p95	Mean	p5	p95	Mean	p5	p95
Cristal (1L bottle)	12.66	6.84	20.47	11.33	5.07	18.17	8.77	2.70	17.99
Cristal (350cc can)	18.84	13.48	25.46	27.69	18.49	38.90	18.87	11.40	31.77
Escudo (1L bottle)	13.98	5.59	27.12	3.82	1.44	7.49	9.78	3.28	16.70
Escudo (350cc can)	25.23	13.99	32.94	16.70	10.22	26.14	26.08	17.04	37.82
Other Cristal	6.61	1.69	11.98	4.71	1.53	11.47	6.48	1.48	11.49
Other Escudo	7.73	1.38	17.76	7.67	3.50	15.56	9.22	2.92	18.82
Baltica (350cc can)	2.35	0.68	6.78	2.82	1.26	4.56	3.95	1.16	11.56
Becker (350cc can)	2.31	0.68	4.83	3.70	1.02	7.40	4.01	1.39	13.00
Stella Artois (354cc can)	1.12	0.33	2.06	2.10	0.53	3.68	1.79	0.42	3.44
Heineken (350cc can)	4.17	1.68	8.42	4.46	1.87	9.24	4.48	1.32	9.60
Royal Guard (350cc can)	1.83	0.44	4.08	4.09	1.14	7.26	2.43	0.45	4.93
Other Brands/Formats	14.40	8.82	22.74	18.72	9.54	26.66	16.03	7.81	28.31

Panel C: Incidence	Pre Treatment			Treatment			Post-Treatment		
	Mean	p5	p95	Mean	p5	p95	Mean	p5	p95
Cristal (1L bottle)	12.42	5.83	24.90	11.69	5.34	19.87	8.52	3.87	15.82
Cristal (350cc can)	20.07	12.96	25.56	29.19	19.07	39.39	19.71	13.86	31.20
Escudo (1L bottle)	14.15	7.53	24.60	4.08	2.10	7.08	10.07	5.47	17.56
Escudo (350cc can)	24.61	13.74	33.57	16.28	11.54	24.52	25.23	17.47	36.32
Other Cristal	5.99	1.78	10.00	4.52	1.68	8.52	6.02	1.35	12.75
Other Escudo	5.72	1.45	11.09	6.00	3.26	8.91	7.65	2.48	13.17
Baltica (350cc can)	3.56	0.95	9.53	4.05	1.56	6.34	5.39	1.87	13.17
Becker (350cc can)	2.80	0.77	5.08	4.43	1.25	8.10	4.44	1.82	8.53
Stella Artois (354cc can)	0.92	0.21	1.50	1.85	0.52	3.11	1.47	0.32	2.47
Heineken (350cc can)	3.31	1.40	5.48	3.60	1.87	8.63	3.65	1.14	7.94
Royal Guard (350cc can)	1.45	0.52	2.51	3.17	1.17	5.50	1.94	0.51	4.49
Other Brands/Formats	16.77	11.02	24.36	19.29	12.35	26.79	18.29	9.96	29.25

Notes: The table shows the mean prices across transactions (top Panel A), the average value market shares calculated across stores (middle Panel B), and the incidence rate calculated as the average presence in consumer's trip across stores (bottom Panel C). For each period described in Figure 3.1, we report the mean and the percentiles 5 and 95 of the corresponding distribution. The statistics consider the sample of frequent beer purchasers that comprises 5,674 households.

Table 3.3: Number of stores under different level of stockouts

	Pre-Treatment (1)	Treatment (2)	Post-Treatment (3)
Panel A: Cristal (1L bottle)			
Less than 2 Stockouts per week	58	38	41
More than 2 Stockouts per week	6	26	23
Total	64	64	64
Panel B: Cristal (350cc can)			
Less than 2 Stockouts per week	63	63	63
More than 2 Stockouts per week	1	1	1
Total	64	64	64
Panel C: Escudo (1L bottle)			
Less than 2 Stockouts per week	60	0	45
More than 2 Stockouts per week	4	64	19
Total	64	64	64
Panel D: Escudo (350cc can)			
Less than 2 Stockouts per week	63	11	63
More than 2 Stockouts per week	1	53	1
Total	64	64	64
Panel E: Others Cristal			
Less than 2 Stockouts per week	61	17	49
More than 2 Stockouts per week	3	47	15
Total	64	64	64
Panel F: Others Escudo			
Less than 2 Stockouts per week	54	6	52
More than 2 Stockouts per week	10	58	12
Total	64	64	64

Notes: The table shows the number of stores under different levels of stockouts. We compute level of stockouts as the weekly average number of days with out-of-stock episodes for each product, across all 64 stores. Column (1), (2) and (3) reports those statistics for the pre-treatment, Treatment and Post-Treatment period, respectively.

tle format products. In addition, different products were more affected than others across different stores.[83]

It may be argued that the retailer could have strategically and selectively managed the frequency of stockouts at different stores. For example, it may have prioritized certain stores with greater demand for the affected products. To investigate this possibility, we run a regression of the stockout indicator on store and time fixed effects. The idea is that the explaining power of the store fixed effects should capture the ability of the retailer to offset stockouts. Table 3.4 shows the marginal contribution of the store fixed effects to the total R-squared on the primary regression. We find that the store fixed effects only explain less than four percent of the variation in three of our main products. Cristal 1L bottle is the exception, with the marginal contribution of the store fixed effects being close to ten percent. Hence, we conclude that the retailers displayed limited efforts to selectively avoid stockouts at certain stores. This finding supports our approach of considering these stockout episodes as a quasi-natural experiment.

We also note that even if retailers had in fact strategically avoided more stockout events at certain stores, our identification strategy will still be valid. This is because, we will also rely on within store variation, where customers visiting a particular store on different dates were exposed to different stockout frequencies. Our data allow us to complement this stockout exposure variation across stores with variation across consumers within a store. We construct an individual measure of stockout exposure for each of the six leading brand products considered. The specific product-consumer

similar to Tables 3.1 and 3.2.

[83]In Table 3.3 we look at stores with more or less than 2 stockouts per week. Similar variation across periods and across products is observed if we use instead 1 stockout per week or 3 stockouts per week.

Table 3.4: Stockout regressions

	(1)	(2)	(3)
Panel A: Cristal (1L bottle)			
Adjusted R-squared	0.0935	0.2097	0.3033
Number of observations	3,136	3,136	3,136
Panel B: Cristal (350cc can)			
Adjusted R-squared	0.0927	0.2443	0.2875
Number of observations	1,760	1,760	1,760
Panel C: Escudo (1L bottle)			
Adjusted R-squared	0.4140	0.5141	0.5515
Number of observations	2,752	2,752	2,752
Panel D: Escudo (350cc can)			
Adjusted R-squared	0.2102	0.4361	0.4686
Number of observations	3,008	3,008	3,008
Week FE	Y	N	N
Date FE	N	Y	Y
Store FE	N	N	Y

Notes: The table shows the results of OLS regressions of the stockout indicator on store and time fixed effects. Each panel is one of the four main products in the paper, and each column is a different specification. Column (1) adds only week FE. Column (2) adds date FE. Finally, column (3) adds both date and store FE. For each product the table reports the number of observations and the adjusted R-squared.

measurement is the number of store visits where the consumer faced a leading brand being unavailable. Figure 3.2 shows the distribution of stockout treatment across individuals and products. From the figure, we can see that the variation across consumers is substantial. Table 3.5 summarizes the considerable heterogeneity of stockout exposure we observe in the data. Therefore, this quasi-natural experiment provided us with a significant exogenous variation in product availability, which will allow us to identify the causal effects of stockouts on future purchase behavior of the more affected individuals.

Table 3.5: Summary statistics of stockouts across consumers

	Mean (1)	p5 (2)	p50 (3)	p95 (4)
Cristal (1L bottle)	3.47	0	2	11
Cristal (350cc can)	0.59	0	0	3
Escudo (1L bottle)	12.76	1	10	31
Escudo (350cc can)	7.33	0	6	19
Other Cristal	6.15	0	4	18
Other Escudo	7.75	0	6	20
Total Stockout Exposure	38.03	4	30	97

Notes: The table shows the statistics of the stockout episodes for each given product across consumers during the treatment period. Column (1) shows the mean, Column (2)-(4) presents the 5th, 50th, and 95th percentile of the distribution, respectively. Total stockout exposure is the sum of stockout episodes across products for a given consumer.

Our proposed metric of stockout treatment has significant advantages over previous papers on stockouts. First, our unanticipated supply shock implies that the treatment variable is independent of demand shocks, ensuring a necessary exogeneity of stockouts to identify their causal effect.

Second, we obtain considerable variation in the severity of the stockout-treatment across products and consumers, ideal for econometric identification. Hence, this con-

Figure 3.2: Histogram of stockout treatment across products and individuals

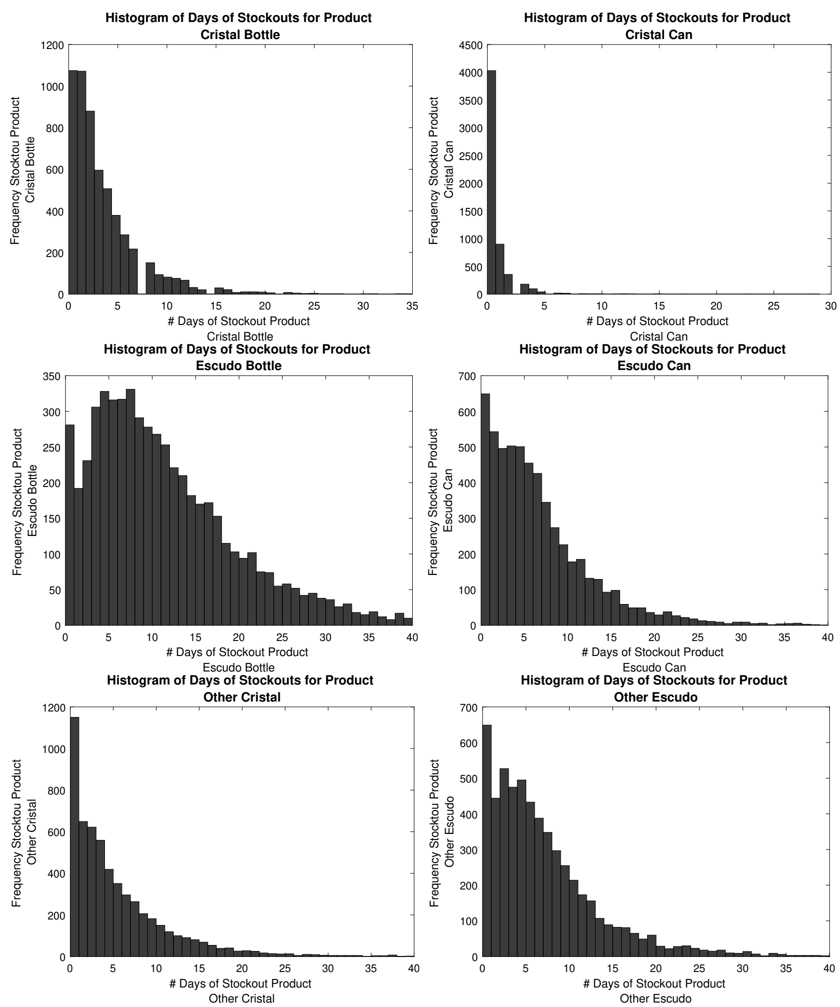

Notes: Each histogram shows the distribution of the days with stockouts that each consumer faced during the 7 weeks of the Treatment period. An stockout is defined as a day with no sales of a given product in a given store. A consumer visiting more than one store in the same day may face more than one stockout episode per day.

tinuous treatment allows us to compare the behavior of consumers exposed to different levels of product unavailability. Furthermore, since we have panel data prior to these prolonged stockouts, we will able to control for differences among consumers in their prior predisposition to purchase each of the brands in the market. Third, given the nature of the shortage, the stockout episodes were not informative to consumers about the quality of the products nor the quality of the retailer's assortment. Arguably, one should not expect that these massive stockouts lead to consumer migration between supermarket chains or between stores within a chain.

3.3 The Treatment Effect of Stockouts on Consumers

This section analyzes the impact of the leading brands' unavailability on consumer purchase behavior. Thus, we consider the exposure to out-of-stock products as a (continuous) treatment on consumers and seek to estimate the average treatment effect (see details in Imbens and Rubin, 2015).

Our analysis examines whether the increase in market shares of the small brands after facing prolonged stockouts is driven by consumers who have no records of purchasing those products before, labeled as *first-timers*. Moreover, we will consider the extent by which the probability of buying one of the small brands for the first time correlates with the severity of the stockouts faced by each consumer.

Table 3.6 describes the statistics of first-time consumers of small brand products over time. Each panel reports the number of new and total buyers for a specific small brand product, the ratio of these two quantities, the (potential) number of consumers who could become first-time buyers in each period and the weekly average of first-time buyers for each product. Column (1) shows the pre-treatment period, which is our baseline. Column (2) in Table 3.6 shows that for four out of these six small brands, the new buyers are about 30 percent of their consumers (see fraction of first timers in this table). This impressive growth in a few weeks is not replicated in Column (3), ruling out a potential market trend in the post-treatment period. Furthermore, since the treatment period is only seven weeks long, the weekly average of new consumers is remarkably higher during the weeks after the shortage.[84]

So far we have shown evidence that the market share increase observed in small brands occurs after the earthquake. However, the connection between this market share changes and the treatment can be strengthened in our analysis. Thus, we exploit the variation across consumers with different stockout treatments to shed light on this issue.

We focus on the sub-sample of 1,225 potential first-time consumers, i.e., those customers who have not purchased any of the small brands during the initial 21 weeks in our data. We estimate probability models where the dependent variable is whether the consumer becomes a first-timer of any of the small brands in a given week. The explanatory variable of interest is the stockout treatment (number of visits with unavailable leading brands). We also include store fixed effects and the number of pre-treatment visits as controls, although different specifications yield the same conclusions.[85]

Table 3.7 shows the estimated average treatment effect on the probability of being a first-time purchaser of small brands. Panel A and B considers a linear probability model and logit model, respectively. Columns (1)-(2) in both panels of Table 3.7 show

[84]As shown in section 3.2.2, the smaller formats were less affected by stockouts, which can justify the switching from large bottles to the can format. However, the format cannot explain the brand switching taking place away from Cristal and Escudo towards smaller brands.

[85]In this analysis, we normalize the maximum treatment per product to be one (the average

Table 3.6: Summary statistics of first-time consumers of small brands products

	Pre-Treatment (21 weeks) (1)	Treatment (7 weeks) (2)	Post-Treatment (16 weeks) (3)
Panel A: Baltica (350cc can)			
# First timers	-	325	139
# Total buyers	697	1,022	1,161
Fraction of first timers	-	0.32	0.12
# Potential first timers	5,674	4,977	4,652
# First timers per week	-	46.43	8.69
Panel B: Becker (350cc can)			
# First timers	-	354	237
# Total buyers	983	1,337	1,574
Fraction of first timers	-	0.27	0.15
# Potential first timers	5,674	4,691	4,337
# First timers per week	-	50.57	14.81
Panel C: Stella Artois (354cc can)			
# First timers	-	273	191
# Total buyers	559	832	1,023
Fraction of first timers	-	0.33	0.19
# Potential first timers	5,674	5,115	4,842
# First timers per week	-	39.00	11.94
Panel D: Heineken (350cc can)			
# First timers	-	284	266
# Total buyers	1,545	1,829	2,095
Fraction of first timers	-	0.16	0.13
# Potential first timers	5,674	4,129	3,845
# First timers per week	-	40.57	16.63
Panel E: Royal Guard (350cc can)			
# First timers	-	374	181
# Total buyers	845	1,219	1,400
Fraction of first timers	-	0.31	0.13
# Potential first timers	5,674	4,829	4,455
# First timers per Week	-	53.43	11.31
Panel F: Other Brands/Formats			
# First timers	-	516	249
# Total buyers	3,859	4,375	4,624
Fraction of first timers	-	0.12	0.05
# Potential first timers	5,674	1,299	783
# First timers per week	-	73.71	15.56

Notes: The table describes the number of new buyers of small brand products (panels) in each period (columns). We define new buyers (first timers) as those who have not purchased the corresponding SKU in our data. Thus, the pre-treatment period is our baseline, with the potential of new buyers being all of the 5,674 households. Column (1) shows the records for the 21 weeks of Pre-treatment period; Column (2) for the 7 weeks of the treatment period; and Column (3) for the 16 weeks of the Post-treatment period. The figures highlight the remarkable peak of new consumers of the small brands during the treatment period as compared to the post-treatment period.

that consumers facing more stockouts of leading brands' during the treatment period are more likely to be a first-timer of small brands, even when controlling for store-fixed effects and the number of store visits during the pre-treatment period. Regarding the size of the stockout effect, we find that the first-purchase probability is about 7 or 8 percent higher for the average stockout exposure relative to the full availability baseline of 29 and 35 percent in Columns (1) and (2) of Panel A, respectively. Hence, the observed purchasing behavior is consistent with stockouts changing the set of products that consumers typically consider to purchase making households more likely to try new products during the treatment period. We also observe in Columns (3)-(4) of Table 3.7 that the impact of stockouts during the treatment period on being a first-timer in the post-treatment period is not significant in both panels. Hence, during the post-treatment period, new buyers are not significantly driven by the stockout treatment.[86]

To account for heterogeneity across stores, we estimate the logit model above considering store-specific stockout effects (i.e., including the interaction of stockouts and store fixed effect). Figure 3.3 shows the histogram of the estimated stockout coefficients during both periods. The histograms confirm that, despite some heterogeneity across locations, the estimates are positive for most stores during the treatment period. In contrast, the estimated effects in the post-treatment data are noisy, with most estimates around zero.

We now turn to investigate whether consumers who became first-timers of small brands during the treatment period stopped buying them in the subsequent weeks

normalized treatment is 0.122 after dividing by the maximum exposure observed in the data).

[86] We validate the quality of our stockout treatment by replicating the estimates using a stockout measure using pre-treatment stockouts. Table D.10 in the Appendix shows that measures of pre-treatment stockouts do not explain consumer behaviour of first-timers.

Table 3.7: Effects of stockouts on the first-time purchase probability of small brands

	Treatment Period (7 weeks)		Post-Treatment Period (16 weeks)	
Panel A: OLS Linear Probability Model	(1)	(2)	(3)	(4)
Stockouts	0.606***	0.678***	0.066	0.063
	(0.160)	(0.174)	(0.127)	(0.125)
Constant	0.253***	0.314***	0.151***	0.139***
	(0.027)	(0.023)	(0.022)	(0.017)
R-squared	0.03	0.08	0.00	0.07
Panel B: Logit	(1)	(2)	(3)	(4)
Stockouts	2.594***	3.065***	0.539	0.606
	(0.709)	(0.809)	(0.997)	(1.079)
Constant	-1.049***	-0.819***	-1.702***	-1.803***
	(0.133)	(0.122)	(0.212)	(0.174)
Log-Likelihood	-793.36	-758.18	-494.22	-447.00
Store FE	N	Y	N	Y
Number of Observations	1,225	1,221	1,225	1,091

Notes: The table shows the average treatment effect of stockouts on the probability of a first-time purchase in small brand products (any brand different from Cristal and Escudo). Panel A uses a linear probability model and Panel B uses a logit model. The stockout variable is the consumer-specific sum of visits with unavailable leading brands during the six weeks of the treatment period, as described in section 3.2.2. As a normalization, we divide the stockout variable by the maximum value. All specifications include the number of pre-treatment visits. Columns (2) and (4) add store fixed effects. Two stores have no variation in stockouts, and we must drop four observations when including store fixed effects (for the same reason we have less observations in Column (4)). Cluster-robust standard errors (at the store level) in parenthesis. P-values notation: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

(post-treatment period). Consider the case where first-timers during the treatment period correspond to consumers loyal to the leading brands being forced by heavy stockouts to try new products. In that case, we should expect no (or at least very limited) repurchase of those small brands by these consumers in the post-treatment period. Figure 3.4 shows the purchase behavior of first-timers for each specific small brand product. The fraction of first-timers repurchasing the particular product they tried for the first time during the treatment period ranges between 16.5 percent (Heineken) and 27.7 percent (Becker). Therefore, at least a fraction of households kept buying the new product, thus their purchase behavior is not reversed in the post-treatment period, suggesting a persistent effect. Nevertheless, we also note that a majority of first-timers stopped buying that specific product in the post-treatment period. They split between buying leading brands only or mixing leading brands with other small brands or not buying any beer products (probably due to seasonality reasons).

Figure 3.3: Store-specific stockout effects on first-time purchase of small brands

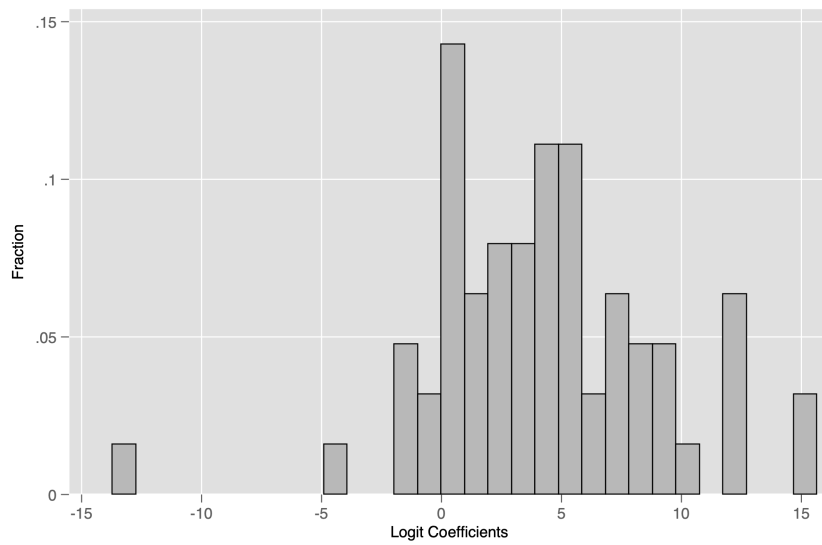

(a) Treatment Period

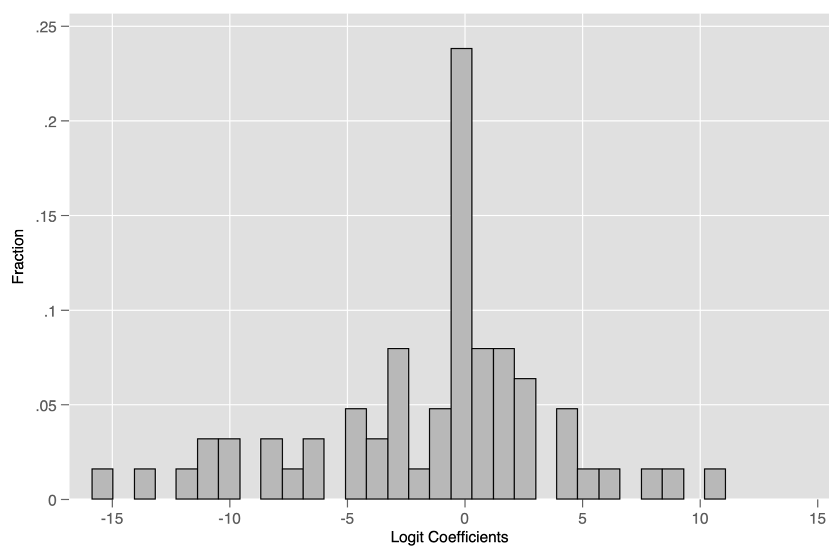

(b) Post-Treatment Period

Notes: The histogram shows store-specific estimates of the effect of stockouts on the probability of a first-time purchase of small brand, conditional on having no records of previous small brand purchases.

Figure 3.4: Purchase behavior of treatment first-timers of small brands in the post-treatment period

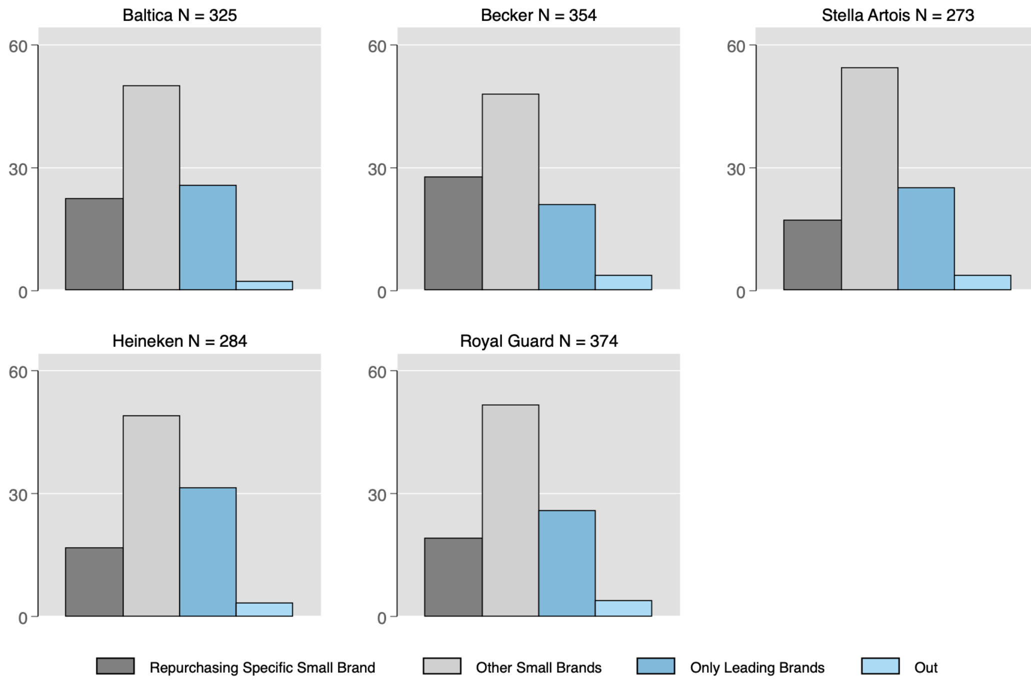

Notes: The figure shows the purchase behavior in the post-treatment period of the subset of consumers who purchased small brands for the first time during the treatment period. The figure shows that about one quarter of these first timers purchased the same small brand during the post-treatment period. Y-axis is in percentage terms relative to total purchase in post-treatment period.

We now replicate this analysis focusing on consumers who purchased one of the small brands during the pre-treatment period. Figure 3.5 shows their distribution of choices in the pre- and post-treatment periods. We observe that the fraction of consumers repurchasing each specific small brand remains virtually unchanged between treatment and post-treatment periods. Hence, we conclude that the leading brands' unavailability did not boost small brand purchases among consumers who have already tried them in the past. This finding supports the idea that the effect of stockouts on choices is driven by consumers who learned about new products during the treatment period. We formalize this argument in the next section and in the theoretical model in section G of the Appendix.

Figure 3.5: Purchase behavior of pre-treatment first-timers of small brands

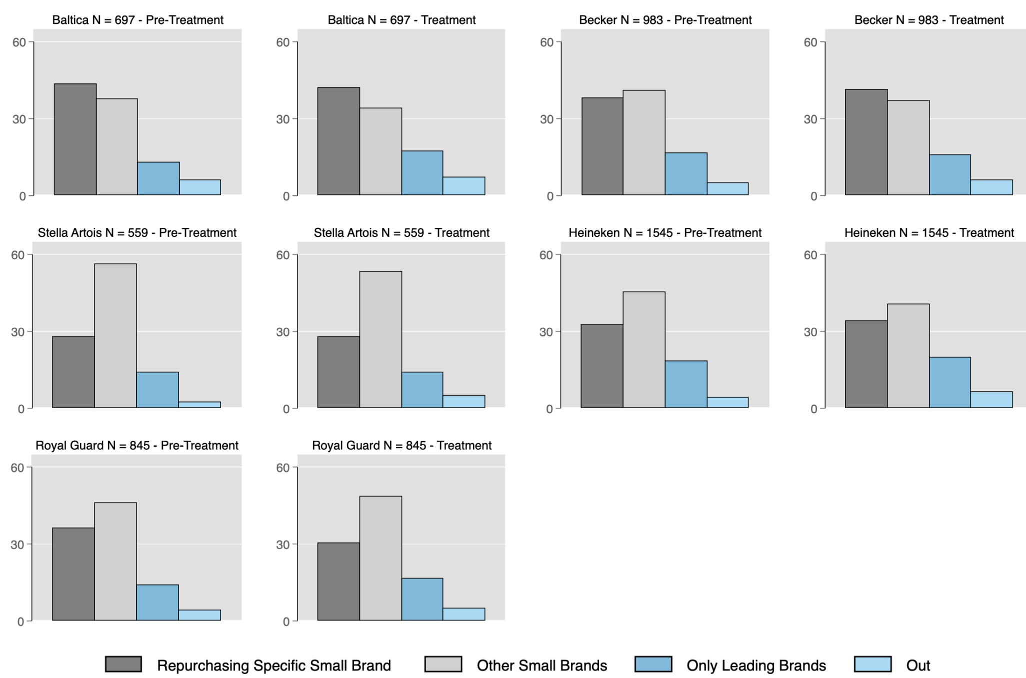

Notes: The figure shows the purchase behavior over time of the subset of consumers who purchased small brands for the first time during the pre-treatment period. The figure is consistent with stockouts (during the treatment period) not altering the distribution of choices for this segment during the post-treatment period. Y-axis is in percentage terms relative to total purchase in the corresponding period.

In summary, we find suggestive evidence that the frequent and prolonged stockouts change purchase behavior for a sizable fraction of consumers. However, a thorough analysis needs to weigh alternative explanations like a change in relative prices, potential state-dependence in consumer choices, and product availability (as some top products were slowly becoming available in the post-treatment period).

3.4 Structural Demand Estimates

We use a structural demand model to quantify the costs of stockouts beyond the same trip purchase decisions but accounting for relative price changes, potential state-dependence in consumer choices, and product availability. Thus, we estimate a discrete choice model focused on the two leading brands, allowing the other alternatives and the outside good to become more attractive in the post-treatment period. These changes in brand valuations are modeled as a function of a consumer's exposure to the leading brands' stockouts.[87]

3.4.1 Econometric Model and Results

We assume that each household h makes discrete choices among the J available products and the outside option (0) in each visit to the supermarket. The close relationship between purchase incidence and market shares shown in section 3.2 (see Table 3.2) suggests that modeling purchase incidence should yield similar insights compared to the analysis of quantity choices. We capture inertia (or variety-seeking behavior) by including the previous product choice in current utilities (Guadagni and Little, 1983). Thus, the utility of alternative j for consumer h in week t of the pre-treatment period is given by:

$$u_{jt}^h = \alpha_j^h + \eta^h \ln(p_{jt}) + \gamma^h \, \mathbb{I}\{s_t^h = j\} + \delta^h X_t + \varepsilon_{jt}^h \tag{14}$$

[87] Appendix section G presents one theoretical model of consideration sets that is consistent with the permanent effects of stockouts in purchase behavior despite other alternative models with the same prediction, such as models of updated beliefs.

where p_{jt} is product j's price in period t, $\mathbb{I}\{s_t^h = j\}$ equals one if product j is the last product that was purchased by the consumer, where $s_t^h \in \{1, ..., J\}$ is the index of the previous alternative purchased by the consumer; X_t is a control variable to parsimoniously account for seasonality calculated as the mean temperature registered in Santiago for each week in our data set; and ε_{jt}^h is a random utility shock i.i.d. according to a Type I extreme value distribution.[88] Consequently, the parameter η^h is the price sensitivity coefficient, while γ^h is the state dependence coefficient for household h.[89]

The product-specific intercepts α_j^h represent the household's persistent brand valuation for product j relative to the outside option. In our estimation, we consider $J = 12$ alternatives in addition to the no purchase option. The first six products correspond to the top leading brands: Cristal 1 liter bottle, Cristal can, Escudo 1 liter bottle, Escudo can, Other Cristal products, Other Escudo products. Alternatives 7-11 correspond to products from smaller brands: Baltica, Becker, Stella Artois, Heineken, Royal Guard, while the 12th alternative considers all other beer products.

If frequent stockouts of the leading brands enlarged consumers' awareness set with better products than the initial inside goods, then we should observe a reduction in the relative valuations of leading brands. Furthermore, we expect that the more stockouts the consumer faced, the greater the product valuation reduction should be for that specific unavailable product.

Therefore, we incorporate the potential effects of stockout treatments in the

[88] A typical concern when estimating demand is the potential endogeneity of prices. In our setting, prices are identical across consumers as the retailer follows a national pricing policy eliminating a possible correlation with the individual demand shocks. See a more comprehensive discussion in Chintagunta et al. (2005).

[89] The model allows for inertia in brand choices if $\gamma^h > 0$. Conversely, $\gamma^h < 0$ predicts variety-seeking behavior.

utility function for the post-treatment periods. Thus, the utility for small brand products (7-12) follows Equation (14), whereas the utility of the top leading brand products (1-6), is modelled as follows:

$$u_{jt}^h = \alpha_j^h + \rho_j ST_j^h + \eta^h \ln(p_{jt}) + \gamma^h \mathbb{I}\{s_t^h = j\} + \delta^h X_t + \varepsilon_{jt}^h, \quad j = 1, .., 6 \quad (15)$$

where the stockout treatment, ST_j^h, is the number of stockout episodes of product j that consumer h was exposed to during the treatment period. If the stockouts for product j led to lower preference for this product, then the changes in product j valuation should be captured by a negative parameter ρ_j. There is experimental evidence of consumers' response to stockouts being positive in some cases (Fitzsimons, 2000; Moore and Fitzsimons, 2014). Hence, we do not restrict the valence of these changes in the relative valuations of the brands due to the stockout treatment and let the data inform us about the sign of this effect.

Note that Equations (14) and (15) include consumer-specific coefficients. We allow for unobserved heterogeneity among consumers with a random coefficients specification. We use Bayesian estimation via Markov chain Monte Carlo simulation. Letting $\theta^h \equiv (\alpha_1^h, .., \alpha_{12}^h, \eta^h, \gamma^h, \delta^h)'$, we specify the following prior distribution: $\theta^h \sim N(\bar{\theta}, \Lambda)$. We also specify the following weak prior and hyper-prior distributions: $\rho \sim N(0, 100^2)$, $\bar{\theta} \sim N(0, 100^2)$, $\Lambda \sim \text{InverseWishart}(17, 17I_{15})$, where I_{15} denotes an identity matrix with 15 rows and columns.

Finally, we account for changes in product availability in our estimation since stockouts were observed not only during the treatment period, as shown in Table 3.3. Thus, we adjust the choice set appropriately for each transaction during the pre- and post-treatment periods. Note that the stockout treatments are not affected by this inclusion, as we do not use the seven weeks of the treatment period in our estimation.[90]

We estimate several specifications and perform model selection given the marginal log-likelihood of each model. Estimation results of our preferred specification are presented in Tables 3.8 and also in Table D.8 in the Appendix.[91]

Consistent with the evidence in section 3.3, we confirm that stockouts explain purchase behavior in the post-treatment period. In effect, the best model in terms of marginal likelihood includes the interaction between the product-level stockout treatment with their correspondent product-dummies for all the six alternatives manufactured by CCU. The log Bayes Factor for the comparison of this model against the same specification but without treatment variables is 99.9, suggesting very strong evidence in favor of the inclusion of the stockout treatment variables.

The effects of stockouts are significant and negative for four of the leading brand alternatives (Cristal bottle, Cristal can, Escudo bottle and Escudo can) and negative and marginally significant for Cristal Other. Thus, our estimates imply that stockouts have long-lasting effects on consumer preferences, decreasing brand-specific valuations. These estimates are consistent with more frequent stockouts making consumers more likely to try different products that may eventually yield higher match values. The model captures this effect by reducing the brand valuation in the post-treatment period for those consumers who faced more stockouts during the treatment

[90] To estimate the discrete choice model, we use the subsample of 5,674 households corresponding to the most frequent buyers of the two leading brands. We discard transactions with more than one beer product to ensure mutually exclusive options consistent with the discrete choice model, dropping six customers that only have multiple beer transactions. Thus, the final estimation sample contains 5,668 households. We provide further details about the data used in the structural estimation in section H in the Appendix.

[91] For computational convenience and ease of interpretation, the treatments were normalized by the overall average stockout exposure across consumers (6.34).

Table 3.8: Empirical results: estimated posterior mean, standard deviation, 2.5% and 97.5% quantiles for $\overline{\theta}$, the square root of the diagonal elements of Λ and the treatment effects (ρ).

		Mean (1)	Std Dev (2)	pc 2.5% (3)	pc 97.5% (4)
Mean Preferences	Cristal Bottle	1.76	0.27	1.62	1.90
$\overline{\theta}_j$	Cristal Can	2.17	0.25	2.04	2.28
	Escudo Bottle	2.05	0.28	1.90	2.20
	Escudo Can	2.44	0.24	2.33	2.55
	Cristal Other	0.30	0.30	0.12	0.47
	Escudo Other	0.54	0.25	0.42	0.66
	Baltica	-2.34	0.28	-2.49	-2.19
	Becker	-0.95	0.31	-1.11	-0.75
	Stella Artois	-1.04	0.37	-1.28	-0.77
	Heineken	0.42	0.25	0.30	0.53
	Royal Guard	-0.81	0.27	-0.94	-0.64
	Other Brands	2.80	0.24	2.69	2.91
	Temperature	1.23	0.16	1.18	1.28
	State Dependence	0.56	0.11	0.53	0.58
	ln(Price)	-1.05	0.10	-1.07	-1.03
Preference	Cristal Bottle	4.57	0.05	4.47	4.68
Heterogeneity	Cristal Can	3.70	0.04	3.62	3.78
(diagonal elements)	Escudo Bottle	3.96	0.05	3.88	4.05
$\sqrt{\Lambda_{jj}}$	Escudo Can	3.55	0.04	3.47	3.63
	Cristal Other	3.56	0.05	3.47	3.65
	Escudo Other	3.45	0.04	3.37	3.54
	Baltica	4.29	0.06	4.17	4.40
	Becker	4.07	0.06	3.96	4.18
	Stella Artois	4.40	0.07	4.26	4.54
	Heineken	3.58	0.05	3.49	3.67
	Royal Guard	3.40	0.05	3.30	3.51
	Other Brands	3.58	0.04	3.51	3.66
	Temperature	1.24	0.02	1.19	1.28
	State Dependence	0.52	0.01	0.49	0.54
	ln(Price)	0.65	0.01	0.64	0.66
Treatment Effects	Cristal Bottle	-0.44	0.03	-0.50	-0.37
ρ_j	Cristal Can	-0.48	0.08	-0.65	-0.32
	Escudo Bottle	-0.19	0.01	-0.21	-0.17
	Escudo Can	-0.06	0.01	-0.09	-0.04
	Cristal Other	-0.04	0.02	-0.08	0.00
	Escudo Other	0.04	0.02	0.00	0.08

period. We also find positive and significant effects for Escudo Other formats, which might be consistent with the experimental findings of Moore and Fitzsimons (2014), which suggest that certain individuals increase their valuation of out-of-stock products after their availability is restored. Furthermore, there is some evidence that the purchases of Escudo Other resemble those of small brand products regarding the existence of first-time purchasers during the treatment period. In particular, consumers not finding the 1 liter bottle and the 350cc can products from this brand may have switched to other formats during the treatment period (see section I in the Appendix).

Note that the model does allow for a degree of transitory inertia in addition to the lasting stockout effects. We estimate a positive state dependence coefficient for the average consumer implying that consumers are on average prone to repurchase their previous choice. However, unlike the decrease in product valuations, this inertia can be reversed by adverse changes in relative prices or shocks. Thus, the model allows for stockouts to alter the last purchase and induce the consumer to repeat that purchase away from the leading brand product temporarily. However, the fact that our treatments are significant after controlling for state dependence in our specification are consistent with a lasting instead of a transitory impact of stockouts on preferences.

The other parameters are in the expected range. Weekly temperature significantly captured seasonality in the beer demand as the post-treatment period covers a typically colder (autumn) low-season for the beer market. As expected, the price estimates are negative for almost all consumers.

We now use the estimates of our structural demand model to assess the impact of product unavailability on consumer behavior.

3.4.2 Quantifying the Effects of Stockouts on Market Shares

We quantify the effects of stockouts on purchase behavior under different counterfactual scenarios. Unlike our descriptive analysis, the structural assessment accounts for observed prices, availability, state dependence, seasonality, and the average stockout-treatment.

First, we compare the steady-state market shares under full availability relative to the scenario with the average stockout level we observe. Hence, as our baseline, we evaluate our estimated demand function at average levels of price, temperature and state dependence and assuming no consumers had been exposed to stockouts during the treatment period. Next, we compute market shares, using the same average prices and state-dependence, but under the stockout exposure observed in the treatment period.

Table 3.9 shows the posterior mean of the market shares and their changes due to the observed stockout exposure. Column (1) shows the baseline market shares assuming no exposure to stockouts during the treatment period, while Column (2) shows the same calculations but under the stockout exposure observed for each consumer in the treatment period. Finally, Column (3) shows the relative change in the market shares due to the stockouts observed in the treatment period.

We see that, in general, the combined market share of the leading brands exhibits a significant reduction. This is driven by the market shares of Cristal bottle, Escudo bottle, and Escudo can which have substantial decreases. Cristal bottle market share decreases from 6.2 to 5.4 percent, which is a 14 percent reduction. Both the bottle and can formats of Escudo are impacted by the stockouts: bottles decrease its preference from 7.0 to 5.5 percent (-21%), while cans decrease from 15.6 to 15.2 (-2%). No

Table 3.9: Structural model - Market share estimates

	Baseline	Post Treatment	% Change
	(1)	(2)	(3)
Cristal Bottle	6.23	5.35	-14.07
Cristal Can	13.19	13.11	-0.56
Escudo Bottle	7.01	5.52	-21.31
Escudo Can	15.54	15.19	-2.24
Cristal Other	2.30	2.32	0.77
Escudo Other	2.54	2.80	10.54
Leading Brands	46.81	44.30	-5.36
Baltica	1.30	1.35	4.58
Becker	1.56	1.63	4.14
Stella Artois	0.52	0.54	4.13
Heineken	1.46	1.52	3.82
Royal Guard	0.73	0.76	4.31
Other Brands	6.64	7.05	6.22
Small Brands	12.21	12.85	5.29
No purchase	40.99	42.85	4.55

Notes: Market shares are calculated using demand estimates and average explanatory variables observed in the data. Column (1) shows the posterior mean of the market shares evaluating the estimated demand function at the average price, state dependence and temperature observed in the Pre-treatment period, imposing full availability. Column (2) shows the same calculation of Column (1) but assuming that all consumers faced the average stockout exposure observed during the Post-Treatment period. Column (3) shows the relative change in the market shares caused by the presence of stockouts.

significant changes are observed for Cristal cans and Cristal Other, which is consistent with these products being less affected by stockouts, as shown in Figure 3.2. At the same time, the market share of "Escudo Other" increases. We conjecture that this increase may be associated with a segment of consumers that try this option for the first time (in our data) during the treatment period.[92]

Regarding the impact on small brands, we observe an increase in total market shares, which is somewhat small when expressed in percentage points (column 3), but non-negligible when considered in relative terms (column 6). The magnitudes of the market share increases are significant and in the order of 0.1 percent for all small brands. However, their relative increases are substantial for all small brands and range between 3.8 percent for Heineken and 6.2 percent for "Other brands". We stress that these sizable increases did not disappear, at least, several months after the stockout episodes.

Our results in Table 3.9 also imply that the outside good increased substantially after the stockout period. The corresponding share increased from 41 percent to 42.9 percent. Therefore, we see that the top brands' unavailability also led to a shrinking in category sales. Notice that seasonality is controlled for in this exercise (as measured by the average temperature).[93] Hence, we interpret this finding as being consistent with substitution of beer consumption through purchases in other categories, probably wine and soft drinks.

Second, to assess the economic magnitude of this effect, we compute the price discount that would offset the negative impact of stockouts on consumer utility for the average consumer. The required price discount for product j, denoted by d_j^*,

[92] See section I in the Appendix for further details.

[93] To test whether our measure of seasonality is driving the results, we perform a robustness check

should satisfy the following condition: $\overline{\eta}\ln((1-d_j^*)\bar{p}_j) = \overline{\eta}\ln(\bar{p}_j) - \rho_j\overline{ST}_j$, where $\bar{p}_j$ is the average price for product j over time; $\overline{\eta}$ and $\overline{ST}_j$ are the average price coefficient and the average stockout treatment across consumers, respectively. The estimated discounts are 20 and 30 percent for the bottle format products for Cristal and Escudo, respectively (see detailed results in Table D.9 in the Appendix). These sizable discounts result from large values of both ρ_j and $\overline{ST}_j$ (especially for Escudo bottle). The remaining products require a single-digit discount to offset the observed stockout effect.

Finally, we also compute the marginal impact of additional stockout episodes during the treatment period for every consumer. These estimates provide a useful benchmark regarding the financial consequences of stockouts and give insights on the resources that may be allocated towards avoiding them. Table 3.10 shows the post-treatment market shares for each alternative when adding one week of stockouts to the average treatment for different products. The first column shows the baseline market share. The next columns contain the counterfactual market shares under one additional week of stockout exposure to all consumers in each specific top brand product. The numbers in bold (diagonal of upper sub-matrix) show the effect of the own-product unavailability. The largest market share loss is -3.9 percent for Cristal can (i.e., 13.11-9.24%) and -1.5 percent for Cristal bottle (i.e., 5.35-3.82%), which are the products with the largest estimated stockout coefficients. The lower panel and the last row shows the corresponding market shares for the small brands and the outside good, respectively, summarizing the estimated substitution from leading brands towards specific small brand products and the no purchase option.

In summary, Table 3.10 shows winners and losers from an additional episode

Table 3.10: Marginal changes in market shares due to an extra week of stockouts

		Additional Stockout for					
	Baseline	Cristal Bottle	Cristal Can	Escudo Bottle	Escudo Can	Cristal Other	Escudo Other
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Cristal Bottle	5.35	**3.82**	5.60	5.39	5.36	5.36	5.35
Cristal Can	13.11	13.33	**9.24**	13.14	13.16	13.13	13.11
Escudo Bottle	5.52	5.60	5.57	**4.77**	5.55	5.52	5.51
Escudo Can	15.19	15.24	15.51	15.31	**14.58**	15.20	15.17
Cristal Other	2.32	2.39	2.58	2.33	2.33	**2.24**	2.32
Escudo Other	2.80	2.82	2.86	2.84	2.85	2.81	**2.89**
Baltica	1.36	1.37	1.40	1.37	1.37	1.36	1.35
Becker	1.63	1.65	1.73	1.64	1.64	1.63	1.63
Stella Artois	0.54	0.55	0.57	0.55	0.55	0.54	0.54
Heineken	1.52	1.53	1.61	1.52	1.53	1.52	1.51
Royal Guard	0.76	0.77	0.81	0.76	0.77	0.76	0.76
Other Small Brands	7.05	7.19	7.34	7.14	7.11	7.06	7.05
Outside Good	42.85	43.77	45.18	43.26	43.20	42.88	42.82

Notes: The matrix shows the market shares for each product resulting in the demand function when using baseline parameters plus the post-treatment estimates at the average observed price and state dependence, but adding an additional week of stockout to the observed average stockout treatment. The difference between the baseline market share (in the first row) is the marginal effect in market shares of an extra week of stockout. The expected effect is a reduction in the same product market share and a weekly increasing in competitor and outside good, that includes not buying beer.

of stockouts for each product of the leading brands. These exercises provide useful information to decision-makers regarding the resources that could be economically justified to prevent stockouts and their resulting losses or gains in market shares.

3.4.3 Characterizing First-time consumers of Small brands

As we obtain estimates at the individual level, we can characterize the segment of consumers that tried the small brands during the treatment period. This exercise can help managers identify which consumer segment is more sensitive to stockouts and hence consider measures to gain them back.

We re-estimate the model using *only* pre-treatment data, and compare the structural parameter estimates of those who tried a small brand product for the first time during the stockout period and those who did not. Using only the pre-treatment data allows us to obtain utility coefficients for each consumer that do not rely on the post-treatment behavior. These coefficients are then used to compare first-timers to all other consumers.

Table 3.11 presents the posterior mean of the 15 estimated coefficients for both sub-samples. Column (1) shows estimates for the subsample of consumers who tried at least one small brand for the first time during the treatment period, while Column (2) considers all other consumers. The third column provides the significance of the difference of mean coefficients across the two samples.

Based on Table 3.11, we find that the first-timers are less price-sensitive than non first-timers, although the mean difference is relatively small. There are no significant differences between the two sub-samples in terms of the coefficients associated with state dependence or seasonality.

Table 3.11: Demand estimates of first-timers and non first timers

Pre-Treatment Average Coefficients	First Timers Mean (1)	 Std Dev (2)	Non First-Timers Mean (3)	 Std Dev (4)	Comparison significance (5)
Price	-0.81	0.01	-0.84	0.01	<0.01***
State Dependence	0.46	0.02	0.46	0.01	0.46
Temperature	0.60	0.08	0.60	0.06	0.50
Cristal Bottle	0.94	0.07	1.65	0.04	<0.01***
Cristal Can	1.71	0.05	2.27	0.03	<0.01***
Escudo Bottle	1.81	0.05	1.63	0.05	<0.01***
Escudo Can	2.65	0.04	2.18	0.03	<0.01***
Cristal Other	-0.33	0.07	0.19	0.05	<0.01***
Escudo Other	0.57	0.06	0.11	0.07	<0.01***
Baltica	-3.04	0.13	-3.13	0.13	0.17
Becker	-1.37	0.10	-1.26	0.09	0.03**
Stella Artois	-1.51	0.19	-1.44	0.18	0.17
Heineken	0.02	0.08	0.09	0.08	0.07*
Royal Guard	-1.17	0.13	-1.13	0.12	0.26
Other Small Brands	2.33	0.04	2.44	0.03	<0.01***
Sample size	1,736		3,932		

Notes: Estimates of the discrete choice demand model using pre-treatment data only (first 21 weeks). The specification follows Equation (14) for all 13 products. Columns (1) and (2) consider the sub-sample of 1,736 consumers who tried at least one small brand product for the first-time during the treatment period. For this sample, we display the estimated posterior mean and standard deviation of their average utility coefficients. Columns (3) and (4) report the same quantities for the remaining 3,932 households. Column (5) shows the significance of the difference between the average coefficients of first timers and the remaining consumers. This test is performed by determining the fraction of MCMC iterations in which the average of each k^{th} coefficient for the sample of first timers is greater than the corresponding average for the remaining customers and then computing the minimum between this fraction f_k and its complement $1-f_k$. We use the following notation: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$ and denotes the probability that the conclusions based on the mean estimates are reversed.

Regarding brand valuations, we observe that first-timers have greater intrinsic preference for Escudo products than non first-timers and are also relatively less prone to prefer Cristal products. In addition, there are some statistical differences in their preferences towards some of the small brands, particularly for Becker and the last alternative which combines all other small brands.

Therefore, from the pre-treatment behavior, we observe that first-timers are more likely to be frequent purchasers of Escudo. When facing Escudo stockout episodes, they are less prone to switch to Cristal and hence are more likely to explore small brands despite being slightly more expensive. This characterization of first timers may be useful to identify and potentially target different groups that might be at risk to switch to other brands when facing prolonged stockouts.

3.4.4 Robustness Check

We perform and discuss several robustness checks in this section. The corresponding tables for these analyses are provided in the Appendix. We first focus on those consumers who faced a minimum exposure to stockouts during the treatment period. These consumers are suitable for a placebo test, and we can check whether their purchase behavior was affected after the earthquake. We define the low-treatment sample as those facing a sum of stockouts at the bottom five percentile of the aggregate exposure to stockouts. Thus, the placebo exercise is conducted using consumers who experienced at most three episodes of leading brand unavailability across all products and visits during the treatment period.

Table D.11 in the Appendix focuses on the placebo sample and shows that the fractions of first-timers within this sample is much smaller relative to those in Table

3.6. Therefore, this evidence suggests that consumers with minimum exposure to stockouts did not change their purchasing behavior by trying small brand products.

Second, another potential concern is that the treatment period coincides with last weeks of the summer, thus, our findings could potentially be explained by seasonality not entirely captured by our average temperature measure, as the beer sales are higher during the summer than the rest of the year. To address this concern, we re-estimated our counterfactual analysis by conditioning on buying beer. The corresponding counterfactual demand estimates show the redistribution of market shares within buyers, regardless of size of the market in the post-treatment period.

Table D.12 in the Appendix replicates the counterfactual calculations for the market shares, conditional on buying.[94] Our conclusions are qualitatively very similar: the stockouts decrease the leading brand market shares from 79.3 to 77.5 (a relative decrease of -2.2 percent). Small brand products gain a significant market share from 20.7 to 22.5 (a relative increase of 8.8 percent) after the treatment period. Notice that some of the less treated products, namely, can products, increase their market shares when conditioning on buying, consistent with substitution to the closest product. Nevertheless, this substitution does not eliminate our main finding that small brands increase their market shares (by a sizable 8.7 percent after conditioning on buying). Table D.13 in the Appendix then shows the marginal stockout effect, conditional on buying. Naturally, the estimated changes are larger, but most findings are qualitatively similar to those in Table 3.10.

In sum, our results seem robust to capturing seasonality by including average temperature or conditioning on buying beer. However, we acknowledge that if dif-

which conditions on buying beer in section 3.4.4.

[94]Notice that market shares are greater than those in Table 3.1 since these calculations do not

ferent beers are especially appealing in certain seasons, we may have a confounding factor not included in our analysis. To the best of our knowledge, there is no evidence that this is the case in the Chilean beer market.

3.4.5 Discussion on Mechanisms and Brand Loyalty

Given our findings, we further discuss the connection between product unavailability and alternative sources of brand loyalty. We follow closely Bronnenberg et al. (2019), who point out that the capital stock of a brand can be explained by evolving quality beliefs through learning, switching costs, advertising, habit formation, and peer influence.

Our most plausible explanation is given by stockouts causing brand loyal consumers to try small brands for their first-time. Erdem and Keane (1996), Ching et al. (2013) and Shin et al. (2012) argue that quality beliefs about products evolve through gradual learning. Our results seem consistent with full learning after the first purchase, equivalent to one-shot learning, where the initial consumption removes most uncertainty about the product match value. Thus, we see the potential gradual learning hypothesis as a complementary force to our preferred explanation. We think that this explanation requires a certain level of volatility on product quality that seems limited in the beer industry relative to other sectors like, for instance, restaurants or airlines.

Other alternative mechanisms seem less relevant in our setting. First, although we have no data on advertising expenditure, we believe that the leading brands had all the incentives to recapture their original market shares through massive advertising. However, our evidence seems that for a subset of consumers, any marketing campaigns were ineffective, so we do not believe that a specific marketing campaign from leading or small brands could explain the persistence of altered market shares. Second, we argue that switching costs in the supermarket industry remained constant across periods and cannot explain the diverse consumer behavior we observe in the post-treatment period. Third, given the shortage's random nature, we believe consumers did not update their beliefs about the leading brands or the retailer's quality.

3.5 Conclusion

A quasi-natural experiment changes the availability of the leading brands in the Chilean beer market and allows us to study whether prolonged stockouts have persistent consequences in equilibrium market shares. After controlling for prices, heterogeneous preferences, state dependence, seasonality, and product availability, we find that the small brands increase their valuations (and market shares) at the expense of the leading brands among the consumers who were exposed to more extensive stockouts.

We find evidence that suggests that the prolonged unavailability changed the consideration set for a substantial share of consumers, who might have become aware of competing products with long-lasting purchase behavior. Our evidence stresses the importance of product availability for brand loyalty relative to other arguments as in Bronnenberg et al. (2021). Despite the advantages of our identification strategy, we acknowledge that our data comes from a specific retailer and category, limiting the generalization of our findings. However, we provide an empirical approach that could be replicated in the future when similar supply-side shocks generate exogenous stockouts.

Future research should follow the pool of first-time consumers several years after the incident to study more permanent consequences. Also, we believe that the purchase behavior of the first-time consumers could be useful to test competing theories of learning (Shin et al., 2012) and the origins of brand loyalty (Bronnenberg and Dubé, 2017). Finally, future research could also investigate whether shelf space allocation changes in response to stockouts. Such information was not available for our sample, but recent innovations in retailing (e.g., Internet of Things) are now making this information more likely to be obtained.

We hope that our findings might be useful to researchers interested in understanding brand loyalty and it is affected by product availability. We also believe that our findings should be helpful for scholars and practitioners concerned with improving assortment and distribution decision making.

consider the outside option.

References

Abaluck, J. and Adams, A. (2021). What do consumers consider before they choose? identification from asymmetric demand responses. *Quarterly Journal of Economics, Forthcoming*, 136:1611–1663.

Abaluck, J. and Gruber, J. (2011). Choice inconsistencies among the elderly: Evidence from plan choices in the medicare part d program. *American Economic Review*, 4:1180–1210.

Abaluck, J. and Gruber, J. (2016). Evolving choice inconsistencies in choice of prescription drug insurance. *American Economic Review*, 8:2145–2184.

Abaluck, J., Gruber, J., and Swanson, A. (2018). Prescription drug use under medicare part d: A linear model of nonlinear budget sets. *Journal of Public Economics*, 164:106–138.

Adams, W. J. (2006). Markets: beer in Germany and the United States. *Journal of Economic Perspectives*, 20(1):189–205.

Ahluwalia, R., Burnkrant, R. E., and Unnava, H. R. (2000). Consumer response to negative publicity: The moderating role of commitment. *Journal of Marketing Research*, 37(2):203–214.

Aizawa, N. and Ko, A. (2023). Dynamic pricing regulation and welfare in insurance markets.

Anderson, E. T., Fitzsimons, G. J., and Simester, D. (2006). Measuring and mitigating the costs of stockouts. *Management Science*, 52(11):1751–1763.

Atal, J., Fang, H., Karlsson, M., and Ziebarth, N. (2020). Long-term health insurance: Theory meets evidence.

Atal, J. P. (2019). Lock-in in dynamic health insurance contracts: evidence from chile.

Atkinson, A. (1970). On the measurement of inequality. *Journal of Economic Theory*, 2:244–263.

Azevedo, E. and Gottlieb, D. (2017). Perfect competition in markets with adverse selection. *Econometrica*, 85:67–105.

Bronnenberg, B. J. and Dubé, J.-P. (2017). The formation of consumer brand preferences. *Annual Review of Economics*, 9:353–382.

Bronnenberg, B. J., Dubé, J.-P., and Moorthy, S. (2019). The economics of brands and branding. In *Handbook of the Economics of Marketing*, volume 1, pages 291–358. Elsevier.

Bronnenberg, B. J., Dubé, J. P., Mela, C. F., Albuquerque, P., Erdem, T., Gordon, B., Hanssens, D., Hitsch, G., Hong, H., and Sun, B. (2008). Measuring long-run marketing effects and their implications for long-run marketing decisions. *Marketing Letters*, 19(3/4):367–382.

Bronnenberg, B. J., Dubé, J.-P. H., and Joo, J. (2021). Millennials and the take-off of craft brands: Preference formation in the U.S. beer industry. NBER Working Paper 28618, National Bureau of Economic Research.

Bronnenberg, B. J., Kim, J. B., and Mela, C. F. (2016). Zooming in on choice: How do consumers search for cameras online? *Marketing Science*, 35(5):693–712.

Browne, M. and Hoffmann, A. (2013). One-sided commitment in dynamic insurance contracts: Evidence from private health insurance in germany. *Journal of Risk and Uncertainty*, 1:81–112.

Cabral, M. (2017). Claim timing and ex post adverse selection. *The Review of Economic Studies*, 84:1–44.

Capps, C., Dranove, D., and Satterthwaite, M. (2003). Competition and market power in option demand markets. *The RAND Journal of Economics*, 34:737–763.

Cavallo, A., Cavallo, E., and Rigobon, R. (2014). Prices and supply disruptions during natural disasters. *Review of Income and Wealth*, 2(60):S449–S471.

Chile, C. (2016). Dimensionamiento del sector de salud privado en chile.

Ching, A. T., Erdem, T., and Keane, M. P. (2013). Learning models: An assessment of progress, challenges, and new developments. *Marketing Science*, 32(6):913–938.

Chintagunta, P., Dubé, J.-P., and Goh, K. Y. (2005). Beyond the endogeneity bias: The effect of unmeasured brand characteristics on household-level brand choice models. *Management Science*, 51(5):832–849.

Ciper, C. (2013a). Reajuste de precios de isapres: una historia de abuso, imposición, desigualdad y lucro. *https://www.ciperchile.cl/2013/04/02/reajuste-de-precios-de-isapres-una-historia-de-abuso-imposicion-desigualdad-y-lucro/*.

Ciper, C. (2013b). Tribunal constitucional falla a favor de isapre por alza en el valor de planes de salud. *https://www.ciperchile.cl/radar/tribunal-constitucional-falla-a-favor-de-isapre-por-alza-en-el-valor-de-planes-de-salud/*.

Cochrane, J. (2017). Here's what healthcare looks like in a perfect world. *The Hill, https://thehill.com/blogs/pundits-blog/healthcare/318906-heres-what-health-insurance-looks-like-in-a-perfect-world/.*

Conlon, C. T. and Mortimer, J. H. (2013). Demand estimation under incomplete product availability. *American Economic Journal: Microeconomics*, 5(4):1–30.

Cuesta, J. I., Noton, C., and Vatter, B. (2019). Vertical integration between hospitals and insurers.

Dalton, C., Gowrisankaran, G., and Town, R. (2020). Salience, myopia, and complex dynamic incentives: Evidence from medicare part d. *The Review of Economic Studies*, 87:822–869.

Dawar, N. and Pillutla, M. M. (2000). Impact of product-harm crises on brand equity: The moderating role of consumer expectations. *Journal of Marketing Research*, 37(2):215–226.

Diamond, R., Dickstein, M., McQuade, T., and Persson, P. (2021). Insurance without commitment: Evidence from the aca marketplaces.

Dickstein, M., Ho, K., and Mark, N. (2021). Market segmentation and competition in health insurance.

Duarte, F. (2012). Price elasticity of expenditure across health care services. *Journal of Health Economics*, 31:824–841.

Dubé, J.-P., Hitsch, G. J., and Rossi, P. E. (2010). State dependence and alternative explanations for consumer inertia. *The RAND Journal of Economics*, 41(3):417–445.

Duffy, E., Dworsky, M., and Whaley, C. (2017). Can a continuous coverage requirement produce a healthy insurance market? *RAND Blog*.

Einav, L., Finkelstein, A., and Cullen, M. (2010). Estimating welfare in insurance markets using variation in prices. *The Quarterly Journal of Economics*, 125:877–921.

Einav, L., Finkelstein, A., Ryan, S., Schrimpf, P., and Cullen, M. (2013). Selection on moral hazard in health insurance. *American Economic Review*, 1:178–219.

Erdem, T. and Keane, M. P. (1996). Decision-making under uncertainty: Capturing dynamic brand choice processes in turbulent consumer goods markets. *Marketing Science*, 15(1):1–20.

Ericson, K. and Sydnor, J. (2022). Liquidity constraints and the value of insurance.

Ericson, K. M. (2014). Consumer inertia and firm pricing in the medicare part d prescription drug insurance exchange. *American Economic Journal: Economic Policy*, 38:38–64.

Fang, H. and Kung, E. (2021). Why do life insurance policyholders lapse? the roles of income, health, and bequest motive shocks. *The Journal of Risk and Insurance*, 88:937–970.

Figueroa, C. (2023). Adverse selection and equity in insurance markets with guaranteed renewable contracts: Evidence from chile.

Fitzsimons, G. J. (2000). Consumer response to stockouts. *Journal of Consumer Research*, 27(2):249–266.

Fleitas, S., Gowrisankaran, G., and Lo Sasso, A. (2020). Reclassification risk in the small group health insurance market.

FONASA (2019). Cuenta pública participativa.

Galetovic, A. and Sanhueza, R. (2013). Un analisis económico de la integración vertical entre isapres y prestadores.

Gensch, D. H. (1987). A two-stage disaggregate attribute choice model. *Marketing Science*, 6(3):223–239.

Geruso, M., Layton, T., McCormack, G., and Shepard, M. (2021). The two margin problem in insurance markets. *Review of Economics and Statistics, Forthcoming.*

Ghili, S., Handel, B., Hendel, I., and Whinston, M. (2022). Optimal long-term health insurance contracts: Characterization, computation, and welfare effects. *Review of Economic Studies, Forthcoming.*

Gottlieb, D. and Smetters, K. (2021). Lapse-based insurance. *American Economic Review*, 111:2377–2416.

Grover, R. and Srinivasan, V. (1992). Evaluating the multiple effects of retail promotions on brand loyal and brand switching segments. *Journal of Marketing Research*, 29(1):76–89.

Guadagni, P. M. and Little, J. D. (1983). A logit model of brand choice calibrated on scanner data. *Marketing Science*, 2(3):203–238.

Handel, B. (2013). Adverse selection and inertia in health insurance markets: When nudging hurts. *American Economic Review*, 7:2643–2682.

Handel, B., Hendel, I., and Whinston, M. (2015). Equilibria in health exchanges: Adverse selection versus reclassification risk. *Econometrica*, 83:1261–1313.

Handel, B., Kolstad, J., Minten, T., and Spinnewijn, J. (2021). The social determinants of choice quality: Evidence from health insurance in the netherlands.

Harris, M. and Holmstrom, B. (1982). A theory of wage dynamics. *The Review of Economic Studies*, 49:315–333.

Heckman, J. J. (1981). Heterogeneity and state dependence. In *Studies in Labor Markets*, pages 91–140. University of Chicago Press.

Heiss, F., McFadden, D., Winter, J., Wuppermann, A., and Zhou, B. (2021). Inattention and switching costs as sources of inertia in medicare part d. *American Economic Review*, 111:2737–81.

Hendel, I. and Lizzeri, A. (2003). The role of commitment in dynamic contracts: Evidence from life insurance. *The Quarterly Journal of Economics*, 1:299–328.

Ho, K. (2006). The welfare effects of restricted hospital choice in the u.s. medical care market. *Journal of Applied Econometrics*, 7:1039–1079.

Horsky, D., Misra, S., and Nelson, P. (2006). Observed and unobserved preference heterogeneity in brand-choice models. *Marketing Science*, 25(4):322–335.

Huang, S. and Salm, M. (2019). The effect of a ban on gender-based pricing on risk selection in the german health insurance market. *Health Economics*, 29:3–17.

Imbens, G. W. and Rubin, D. B. (2015). *Causal inference in Statistics, Social, and Biomedical Sciences*. Cambridge University Press.

Jing, X. and Lewis, M. (2011). Stockouts in online retailing. *Journal of Marketing Research*, 48(2):342–354.

Keller, K. L. (1993). Conceptualizing, measuring, and managing customer-based brand equity. *Journal of Marketing*, 57(1):1–22.

Keller, K. L. (2013). Strategic brand management. 4-th ed.

KKF (2022). Medicare beneficiaries rarely change their coverage during open enrollment.

La Tercera, C. (2022). La molestia de las isapres por el reajuste de precios en planes que definió la super de salud: alegan que subestima sus costos. *https://www.latercera.com/pulso-pm/noticia/asociacion-de-isapres-asegura-que-el-reajuste-de-precios-definido-por-la-super-de-salud-subestima-los-costos-que-registro-el-sector/DMBVSNPKGVH2TNDMAICMHX7UI4/*.

Libertad y Desarrollo, C. (2017). Judicialización en isapres: antecedentes y alternativas de solución. *https://lyd.org/wp-content/uploads/2017/07/SISO-166-Judicializacion-en-Isapres-antecedentes-y-alternativas-de-solucion-Junio2017.pdf*.

Moore, S. G. and Fitzsimons, G. J. (2014). Yes, we have no bananas: Consumer responses to restoration of freedom. *Journal of Consumer Psychology*, 24(4):541–548.

Nedungadi, P. (1990). Recall and consumer consideration sets: Influencing choice without altering brand evaluations. *Journal of Consumer Research*, 17(3):263–276.

Pardo, C. (2019). Health care reform, adverse selection and health insurance choice. *Journal of Health Economics*, 67.

Pardo, C. and Schott, W. (2013). Health insurance selection: a cross-sectional and panel analysis. *Health Policy and Planning*, 29:302–312.

Pauly, M. and Herring, B. (2006). Incentive-compatible guaranteed renewable health insurance premiums. *Journal of Health Economics*, 25:395–417.

Pope, C. (2020). Continuous renewable coverage: Rx for america's dysfunctional health-insurance system. *Manhattan Institute*.

Roberts, J. (1989). A grounded model of consideration set size and composition. *ACR North American Advances*.

Roberts, J. H. and Lattin, J. M. (1991). Development and testing of a model of consideration set composition. *Journal of Marketing Research*, 28(4):429–440.

Robles-Garcia, C. (2022). Competition and incentives in mortgage markets: The role of brokers.

Rossi, P. E. and Allenby, G. M. (2003). Bayesian statistics and marketing. *Marketing Science*, 22(3):304–328.

Shepard, M. (2022). Hospital network competition and adverse selection: Evidence from the massachusetts health insurance exchange. *American Economic Review*, 2:578–615.

Shin, S., Misra, S., and Horsky, D. (2012). Disentangling preferences and learning in brand choice models. *Marketing Science*, 31(1):115–137.

Shugan, S. M. (1980). The cost of thinking. *Journal of Consumer Research*, 7(2):99–111.

Small, K. and Rosen, H. (1981). Applied welfare economics with discrete choice models. *Econometrica*, 49:105–130.

Superintendencia (2006). Circular if 15.

Vulcano, G., Van Ryzin, G., and Ratliff, R. (2012). Estimating primary demand for substitutable products from sales transaction data. *Operations Research*, 60(2):313–334.

Yoo, B., Donthu, N., and Lee, S. (2000). An examination of selected marketing mix elements and brand equity. *Journal of the Academy of Marketing Science*, 28(2):195–211.

A Stylized Model

In this section, I develop a simple stylized model of insurance contract choice and insurer pricing that highlights the effects of a ban of gender-based pricing in health insurance markets, and how the effects of the ban change under GR contracts and non-GR contracts.

The general setting of the model follows the one in Einav et al. (2010) and Geruso et al. (2021). First, I assume perfect competition in the private market (zero profit condition) and that each company has zero administrative costs. This means that, in a competitive equilibrium, if allowed, firms will set plan prices equal to their average costs.[95] Second, I will consider two fixed contracts offered by each firm, $j = \{H, L\}$, where H is a high-coverage policy and L is a low-coverage policy. Finally, there is a safety net public option U, which has lower coverage than any plan in the private market. A price vector is defined as $P = \{P_H, P_L, P_U\}$.

Demand and costs

The model's primitives are consumers' willingness-to-pay (WTP) for each plan and expected insurer costs for consumers of specific types in each plan. Regarding demand, let $W_{i,H}$ be WTP of consumer i for plan H, and $W_{i,L}$ be WTP for L, both defined as WTP relative to U ($W_{i,U} \equiv 0$). I make the following assumptions:

Assumption 1. *Vertical Ranking*: $W_{i,H} > W_{i,L} > 0\ \forall i$

The vertical ranking assumption implies that the products are vertically ranked for all consumers. Notice that this does not allow for plans being horizontally differ-

[95]In Einav et al. (2010) the setting is a market with short-term contracts in which firms set plan prices equal to their annual average costs. Here, for simplicity, I am assuming that the same will hold even though the plans in this market are long-term contracts. A way to think about this assumption is that firms will set plan prices equal to their long-term average costs, thus, when defining expected costs, I actually mean long-term expected costs, and profits actually mean long-term expected profits.

entiated, which is one of the reasons why an empirical model is needed.

Assumption 2. *Single dimension of WTP heterogeneity and only 4 types*: there are only four types of consumers in the market; "high-income males" ($i = HM$), "low-income males" ($i = LM$), "high-income females" ($i = HF$) and "low-income females" ($i = LF$). These types can be ranked based on declining WTP, with $W_{HF,L} > W_{HM,L} > W_{LF,L} > W_{LM,L}$ and $(W_{HF,H} - W_{HF,L}) > (W_{HM,H} - W_{HM,L}) > (W_{LF,H} - W_{LF,L}) > (W_{LM,H} - W_{LM,L})$.

This assumption implies that consumers' WTP for H and L—which in general could vary arbitrarily over more dimensions—are assumed to collapse to a single-dimensional type s with $s = \{HM, LM, HF, LF\}$. Thus, if $W_{s,H} - W_{s,L} > P_H - P_L$, then consumer type s selects plan H, and if $W_{s,L} > P_L - P_U$ and $W_{s,H} - W_{s,L} < P_H - P_L$, the consumer chooses plan L. In any other case, the consumer would choose U.

In terms of expected insurer costs for consumers, I define "type-specific costs" for each plan j as $C_j(s) = E[C_{ij}|s_i = s]$ with $C_H(s) > C_L(s)$ and $C_L(s) > C_U(s)$ $\forall s$. Supported by empirical evidence in Figure 1.6, I assume that females have higher health care costs than males, thus $C_j(HF) = C_j(LF) > C_j(HM) = C_j(LM)$ $\forall j$. Notice that this implies adverse selection as consumers with higher willingness to pay also have higher health care costs. Relatedly, plan-specific average costs $AC_j(P)$ are defined as the average of $C_j(s)$ for all types who buy plan j at a given set of prices.

Initial condition and equilibrium under gender-based pricing

Assumption 3. *Initial condition and dynamic price adjustments:* there is a period $t = 0$, at the onset of the market, in which insurers set prices equal to average costs across all potential enrollees. After that, premiums are adjusted dynamically until an equilibrium is reached.

The assumption implies that, under a ban of gender-based pricing, at $t = 0$, $P_H = \frac{C_H(HF)+C_H(HM)+C_H(LF)+C_H(LM)}{4}$ and $P_L = \frac{C_L(HF)+C_L(HM)+C_L(LF)+C_L(LM)}{4}$. This allows for dynamic price adjustments, after consumers start selecting into plans, until an equilibrium is reached. Under GR contracts, the assumption forces cross-subdization between plans. In the case of non-GR contracts, the assumption is inconsequential.

The main problem with the assumption is that it prevents forward-looking behavior from insurers, which might not be realistic in this simple setting with only two vertically differentiated policies and four consumer types. However, in practice, health insurance markets have multiple plans differentiated both vertically and horizontally, and multiple consumer types. Moreover, the way companies adjust prices over time after creating new plans in the ACA Marketplaces and in the Chilean private market (before and after the ban) provides strong support for the assumption. Nonetheless, as the ban in Chile was implemented after the market was created, for the empirical application of this paper, the assumption is not required. The assumption is only needed in order to generalize the results to markets that start with new plans under a ban and GR contracts.

If gender-based pricing is allowed in the market, plans can charge females X times more than males for the same plan. In Chile, before the ban, $X \approx 3$. I take the equilibrium under gender-rated prices as given and study how the equilibrium changes under a ban and GR contracts, or under a ban and non-GR contracts. Particularly, for this case, based on empirical evidence in Figure 1.1, I assume that in equilibrium high-income males select H, low-income males and high-income females select L, and low-income females select U. Prices for males and X are set such that profits in each plan, and for each gender, are equal to zero.

Banning gender-based pricing

Compared to the last scenario, in a market with a ban of gender-based pricing and $t = 0$, both $P_H - P_L$ and $P_L - P_U$ will be lower for females. In that case, $P_H = \frac{C_H(HF)+C_H(HM)+C_H(LF)+C_H(LM)}{4}$ and $P_L = \frac{C_L(HF)+C_L(HM)+C_L(LF)+C_L(LM)}{4}$. Given consumers' ordering of WTP, high-income females select H and low-income females choose to enter the private market and select L.[96] Notice that for males the opposite is true, that is, $P_H - P_L$ and $P_L - P_U$ are higher. For expositional reasons, I will assume that, at those premiums, $W_{HM,L} > P_L - P_U$ and $W_{HM,H} - W_{HM,L} < P_H - P_L$, and $W_{LM,L} > P_L - P_U$ and $W_{LM,H} - W_{LM,L} < P_H - P_L$ Thus, both high-income males and low-income males select L. Even though this is a strong assumption, it is an approximation to the actual response of consumers to premium changes observed in section 1.6.

As females have higher health care costs than males, after these choices, firms will face pressure to increase prices due to the zero profit condition. The way they do that depends on whether the market offers GR contracts or non-GR contracts. I will start with the common case of non-GR contracts, such as the short-term contracts offered in health insurance markets in the U.S., and then move to the setting with GR contracts.

Non-GR contracts: no reclassification risk protection

In this scenario, firms will set plan prices equal to their average costs (Handel et al., 2015 and Azevedo and Gottlieb, 2017). In that case, $P_H = AC_H(P) = C_H(HF)$ and $P_L = AC_L(P) = \frac{C_L(LF)+C_L(HM)+C_L(LM)}{3}$. Now, either this is the final equilibrium or

[96] Here I assume that, at these prices, low-income females will not choose plan H, that is, $W_{LF,H} - W_{LF,L} < P_H - P_L$, $W_{LF,L} > P_L - P_U$.

high-income females switch to L after the increase in $P_H - P_L$ (*i.e.* plan H unravels). For the purposes of this model, this is inconsequential.

GR contracts: reclassification risk protection

In this scenario, firms are not allowed to change prices based on the enrollees that select into each plan (*i.e.* protection against reclassification risk). Instead, the percentage price increase of both plans must be the same. Therefore, $P_H = P_{H,t-1} + \Delta^{\%} P_{H,t-1}$ and $P_L = P_{L,t-1} + \Delta^{\%} P_{L,t-1}$, where $\Delta^{\%} = \frac{P_H - P_{H,t-1}}{P_{H,t-1}} = \frac{P_L - P_{L,t-1}}{P_{L,t-1}}$ and $t-1$ stands for prices before the current premium adjustment. The main difference between this scenario and the previous scenario with non-GR contracts is that now firms must cross-subsidize selection in plan H with plan L. Premiums will keep adjusting in this manner until an equilibrium is reached. The main implication of the model is:

Proposition 1. *In a health insurance market with a ban of gender-based pricing:* P_L(*GR contracts*) $>$ P_L(*non-GR contracts*) *and* P_H(*GR contracts*) $<$ P_H(*non-GR contracts*)

Proof. Both of these inequalities follow directly from the zero profit condition and cross-subsidization. In the case of L, P_L(non-GR contracts) $= AC_L(P) = \frac{C_L(HM)+C_L(LF)+C_L(LM)}{3} < \frac{C_L(LF)+C_L(LM)}{2} + \Delta^{\%}(\frac{C_L(LF)+C_L(LM)}{2}) = P_L$(GR contracts). This can be easily shown after some algebra by finding $\Delta^{\%}$ that solves P_L(GR contracts)$+P_H$(GR contracts) $= C_H(HF) + \frac{C_L(HM)+C_L(LF)+C_L(LM)}{3} = AC_H(P) + AC_L(P)$ and replacing it in the inequality. Once the first inequality has been proven, the second one for H is trivial because if P_L(GR contracts) $> AC_L(P)$, then profits for L are positive. Thus, profits for H must be negative such that the zero profit condition holds, meaning that P_H(GR contracts) $< AC_H(P) = P_H$(non-GR contracts). □

This proposition can be further generalized, but the most important takeaway is that prices of low-coverage plans should be higher under GR contracts and prices of high-coverage plans should be lower in that scenario. This happens because low-coverage plans are forced to cross-subsidize adverse selection in high-coverage plans. Notice that this implies that the number of price sensitive consumers choosing the private market should be lower as some of them are not able to afford the higher premiums in low-coverage plans.

Corollary 1. *Share of consumers in the private market will be lower under GR contracts than under contracts non-GR contracts.*

The empirical outcome of banning gender-based pricing in the Chilean private health insurance market is much more complicated than what is outlined in this simple model. First, the private sector in Chile features multiple plans per firm, with multiple coverage levels, that differ not only on price and coverage, but also on characteristics such as the hospital network offered or the firm's brand (*i.e.* horizontal differentiation). Second, by law, policyholders must spend 7% of their income in health insurance plans, meaning that high-income consumers are likely to remain in the private market purchasing high-coverage plans even if prices go up. Finally, enrollees in the private market, and in the public option, differ not only by gender, but also by age, family composition, and income, and each of these "consumer-types" has different WTP for insurance and expected costs. Therefore, quantifying the actual difference in outcomes of implementing the ban under GR contracts or under non-GR contracts is an empirical question. However, as long as enough healthy consumers remain in low-coverage plans in the private market, none of these features should affect the main predictions of the model: higher prices of high-coverage plans under non-GR contracts, lower prices of low-coverage plans, and a higher share of enrollees in the private market. The reason is that these healthy enrollees will keep the costs, and, hence, the prices of low-coverage plans, at low levels.[97]

B Forward-Looking Behavior

Even though it is hard to conclusively prove that consumers do not exhibit forward-looking behavior in this market, it is possible to show that this is not a first order issue, meaning that the estimates from the demand model in section 1.5 will not be strongly biased by omitting this factor. To do that, I exploit a policy implemented in 2011 that completely changed dynamic incentives, particularly among new enrollees, and show that new consumers did not change their purchasing behavior in response to this regulation.

Before March of 2011, insurers were able to update (increase) the premium of enrollees according to their age in a predetermined way. Specifically, as policyholders aged, firms would increase their premiums by moving down the rows in the left table of Figure D.2. Thus, for example, if a male signed a new contract at age 34, he would get a factor of 1. But once he turned 35, his factor, and, hence, his premium, would have increased by 1.05. This caused public outraged, especially because insurers would only update the premium of consumers when this meant a higher price, but not when it meant a lower price. The regulator then decided to intervene, and by the end of 2010 enacted a policy that prohibited this kind of updating. Instead, as long as consumers stay in their plans, companies must keep the age factor fixed.

[97] If only high-cost enrollees remain in the private market under non-GR contracts, then it might happen that the price of low-coverage plans is actually higher in this scenario than under GR contracts, contradicting then the prediction of the model. This last scenario seems implausible given the multitude of plans and consumer types in the Chilean setting. Nevertheless, generalizing this simple theoretical model in order to provide clear predictions in a more complicated environment, such as the Chilean market, is an interesting area for future research. This is out of the scope of the current paper.

This policy drastically changed dynamic incentives when purchasing a new plan in the market. Specifically, policyholders know that they are likely to remain in their plans for longer because their age factors will not increase as they age. If consumers, then, are forward-looking, I should observe a structural change of how they choose plans in the private market after the regulation was implemented. In Figure A.1 below I plot the premium paid by new enrollees as a percentage of their income, probably one of the most relevant variables to examine. The data is smoothed using a 3-month moving average.

As can seen from the figure, policyholders do not seem to change their behavior when choosing new plans. Besides some normal fluctuations, there is no clear structural change in how much they are spending on health insurance plans before and after the policy was implemented. In addition to this variable, I also look at other measures such as raw premium or type of plan chosen, among others, finding similar results. That is, this is evidence that, at least on aggregate, consumers in this market do not exhibit strong forward-looking behavior. This is in line with findings from the Medicare part D literature, where, for example, Abaluck et al. (2018) and Dalton et al. (2020) have found strong levels of myopia.

Certainly, this is not meant to be a rigorous test of forward-looking behavior among consumers in this market. But the lack of evidence showing a clear change of behavior after the policy was implemented suggests that this is not a first order issue in the Chilean private system, and, hence, my demand estimates will not be significantly affected by omitting this feature in the estimation of the structural model.[98]

[98]An interesting area of study for future research would be to use this policy for a more sophisticated test of forward-looking behavior. However, this is out of the scope of the current paper.

Figure A.1: Premium paid as a percentage of income - New enrollees 2008-2014

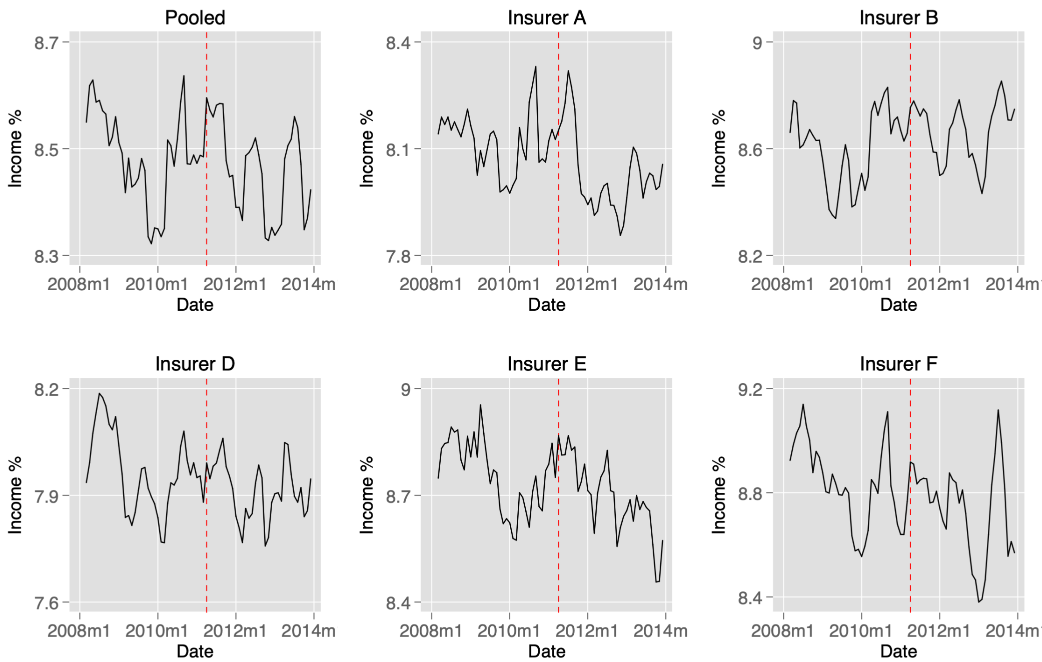

Notes: This figure shows premiums paid by new enrollees as a percentage of their income between 2008 and 2014. The vertical red dashed line represents the introduction of a new policy that fixed the age function to the age at which the policyholder initially enrolled in her health insurance plan. The data is smoothed using a 3-month moving average. Data for insurer C before 2013 is questionable. Therefore, I do not include that firm in the analysis.

C Hospital Choice Model

Construction of coverage rates for hospitals

Prices faced by consumers when choosing a hospital have two main components: the actual price charged by the hospital and the coverage rage offered by the insurance plan. This subsection describes how I construct these coverage rates for each plan in each hospital.

A big portion of this task was completed manually. Specifically, the platform *Queplan.cl* gave me access to their database of contracts for each plan in the market, allowing me to retrieve each contract file and gather information about the coverage rate offered in each hospital. This manual procedure was done for all plans with tiered networks. For hospitals outside these preferred network, and for plans with unrestricted open networks, coverage rates are more complicated to retrieve manually as they involve visit payments cap by insurers. Therefore, for those cases, I impute the coverage rates using the observed empirical coverage rates from the enrollees in those plans in the claims data.[99] Finally, for each plan, I create only four coverage rates: one for the top 5 most expensive hospitals that are part of a preferred network, one for the top 5 most expensive hospitals that are not part of a preferred network, one for the other hospitals that are part of a preferred network, and one for the other hospitals that are not part of a preferred network. I do this by taking the average coverage rates across hospitals.

The final outcome of this procedure is four coverage rates for each health insur-

[99]For a few plans that I could not find their respective contract in the platform or that did not have enough claim information, I extrapole their coverage rates using similar plans. Specifically, I group plans by company, hospital network offered, and extra plan characteristics. Then, for plans with missing coverage rates, I look at plans from the same group but without missing information and I impute the missing coverage rate using the base price difference between these plans.

ance plan in the private market. These are the rates c_{jh} I use to estimate the hospital choice model. Notice that the main advantage of using the *Queplan.cl* information is that I do not have to rely extensively on empirical coverage rates, which are likely to be biased, especially for the most expensive hospitals.

Construction of hospital prices

In order to construct hospital prices I follow Shepard (2022) closely. Therefore, for the interested reader, I recommend looking at his article for more details. The main difference with the current paper is that I only focus on inpatient care because in Chile there are too many outpatient units, making it almost impossible to identify each one of them in the data.

For inpatient care, the most natural service unit is the diagnosis-related group (DRG), which is the standard measure used in hospital price analyses. Nonetheless, because not all admissions are DRG-paid and because even DRG payment allows exceptions due to outlier adjustments, I estimate a pricing model that allows quantity to vary within a DRG or diagnosis based on other patient severity observables. Essentially, this method defines the quantity associated with each hospital admission in a continuous way based on a projection of spending onto DRG/diagnosis categories and other patient observables.

Consider a particular admission a – for enrollee i in plan j in year t for DRG (or diagnosis) d at hospital h. I regress log total payments ($log(Paid_{a,i,j,t,d,h})$) on insurer-hospital dummies ($\alpha_{h,j}$), year dummies (β_t), DRG/diagnosis fixed effects (γ_d), and patient severity factors ($Z_{a,i,t}$) comprised of gender x age groups (in 5-year bins), income groups and diagnoses groups:

$$log(Paid_{a,i,j,t,d,h}) = \alpha_{h,j} + \beta_t + \gamma_d + Z_{a,i,t}\delta + u_{a,i,j,t,d,h} \tag{16}$$

Using estimates from this regression, I define the quantity unit as the component of payment arising from DRG/diagnosis and severity:

$$\tilde{Q}_{a,i,t} \equiv exp(\hat{\gamma}_d + Z_{a,i,t}\hat{\delta}) \tag{17}$$

The remainder of the regression is defined as price:

$$\tilde{P}_{a,i,j,t,h} \equiv exp(\hat{\alpha}_{h,j} + \hat{\beta}_t) \tag{18}$$

The price I use then in the hospital choice model is $p_{ijhdt} \equiv \tilde{P}_{a,i,j,t,h} \times \tilde{Q}_{a,i,t}$. That is, with this procedure I get rid of the residual portion of the total payment to the hospital. This makes sense because the variable I want to create is the *expected price* for a particular diagnosis in a hospital in a year for a particular group (*e.g.* low-income young males).[100]

C.1 Empirical Model

In this subsection I describe the hospital choice model I estimate in order to construct the expected utility for a consumer from the hospital network offered by a specific plan, that is, EU_{ijkt}. These models are now standard in the health insurance literature and, hence, I will brief in each part of the model. For more detailed descriptions, see

[100] See Table D.1 for descriptive statistics showing the results of this decomposition for the 14 main private hospitals in Santiago.

both Cuesta et al. (2019) and Shepard (2022), which I follow closely.

In the specification used in this paper, the choice of a hospital is conditional on the diagnosis and the health insurance plan of the individual. Specifically, an enrollee with a certain health condition will choose a hospital among those available in her choice set, and this choice will depend on how much she needs to pay to treat that condition in each hospital, on the distance between her house and each hospital, and on the quality of each hospital. That is, the utility of consumer i in period t with plan j for choosing hospital h to treat diagnosis d takes the following form:

$$u^H_{ijhdt} = \alpha^H_i c_{jh} p_{ijhdt} + \beta^H_i Dist_{ih} + \delta_h + \epsilon_{ijhdt} \tag{19}$$

where $c_{jh}p_{ijhdt}$ is how much consumer i must pay to treat condition d at hospital h. This amount will depend on both the coverage offered by her insurance plan in that hospital, c_{jh}, and the price charged by the hospital to treat that condition to consumer i, p_{ijhdt}. $Dist_{ih}$ is the distance between consumer i's house and hospital h. Finally, δ_h are hospital fixed effects.[101] As in the plan choice model of section 1.5, coefficients in the utility function vary by demographic and socioeconomic groups.

The outside option in this case means going to a public hospital or to a smaller private hospital.[102] Specifically, the utility in case of choosing the outside option takes the following form:

$$u^H_{ij0dt} = \alpha^H_i c_{j0} p_{ij0dt} + \upsilon_{l(i)} + \epsilon^H_{i0hdt} \tag{20}$$

where $\upsilon_{l(i)}$ are county fixed effects in order to control for the fact that the quality of

[101] In the preferred specification I use diagnoses interacted with hospital fixed effects.

[102] The choice set consists of the 14 main private hospitals in the capital Santiago. Any other

the outside option might vary between counties. Similarly to the model in section 1.5, this hospital choice model is estimated by maximum likelihood.

Results

Table A.1 below describes the main results from the hospital choice model. Specifically, the first panel of the table displays the premium coeffients and the second panel shows the distance coefficients. Each row represents a different estimate for a different consumer group. Different columns show estimates considering a different set of fixed effects. Column (1) does not allow for any type of fixed effects. Column (2) considers hospital fixed effects, and column (3) includes hospital fixed effects interacted with diagnoses, which is also my preferred specification.

The main takeaways from the table are the following: (i) women, older and richer people are less price sensitive when it comes to choosing a hospital, and (ii) older and richer people dislike more to travel to hospitals. These results are now standard in the health insurance literature, which is reassuring that my specification is correct. Moreover, in Figure A.1 I plot the hospital fixed effects from column (2) of Table A.1 against the final accreditation grade assigned by the regulator to each hospital as a proxy for quality.[103] The high correlation between these two variables supports the fact that hospital fixed effects are capturing the quality of each hospital.

Once I have the estimates from the hospital choice model, I can construct the expected utility obtained from each health insurance plan j for each consumer i as

private hospitals falls into the outside option. The same goes for any public hospital.

[103]The final accreditation grade is a score assigned by the regulator after they do a thorough inspection on each hospital. This is mandatory in order to be able to operate in the country.

Table A.1: Parameters estimates - Hospital choice model

Variable	(1) Coeff	(1) S.E.	(2) Coeff	(2) S.E.	(3) Coeff	(3) S.E.
α^H - Hospital Price						
age ≤ 35	-2.183	(0.020)	-3.192	(0.020)	-3.266	(0.020)
age $\in (35, 45]$	-1.879	(0.020)	-2.831	(0.020)	-2.869	(0.020)
age $\in (45, 55]$	-1.843	(0.020)	-2.697	(0.020)	-2.668	(0.020)
age > 55	-1.736	(0.021)	-2.544	(0.020)	-2.488	(0.020)
Single female	0.347	(0.010)	0.371	(0.011)	0.336	(0.011)
Family	0.021	(0.012)	-0.030	(0.012)	-0.049	(0.013)
Income 2nd Tercile	0.403	(0.021)	0.345	(0.020)	0.343	(0.020)
Income 3rd Tercile	1.398	(0.018)	1.393	(0.017)	1.352	(0.017)
β^H - Distance to hospital						
age ≤ 35	-0.153	(0.002)	-0.180	(0.003)	-0.177	(0.003)
age $\in (35, 45]$	-0.160	(0.002)	-0.183	(0.003)	-0.179	(0.003)
age $\in (45, 55]$	-0.171	(0.002)	-0.196	(0.003)	-0.181	(0.003)
age > 55	-0.189	(0.003)	-0.209	(0.003)	-0.189	(0.003)
Single female	0.007	(0.001)	0.026	(0.002)	0.016	(0.002)
Dependents	0.016	(0.001)	0.031	(0.002)	0.012	(0.002)
Income 2nd Tercile	-0.009	(0.002)	-0.016	(0.002)	-0.010	(0.002)
Income 3rd Tercile	-0.031	(0.002)	-0.037	(0.002)	-0.028	(0.002)
Observations	2,726,739		2,726,739		2,726,739	
Hospital FE	N		Y		N	
Hospital-diagnosis FE	N		N		Y	

Notes: This table shows the logit estimates of the hospital choice model. The first panel displays the premium coeffients and the second panel shows the distance coefficients. Estimates vary across age groups, household composition, and income. Different columns show estimates considering a different set of fixed effects. Column (1) does not allow for any type of fixed effects. Column (2) considers hospital fixed effects. Column (3) considers hospital fixed effects interacted with diagnoses. Standard errors are in parenthesis.

Figure A.1: Hospital fixed effects and accreditation grade

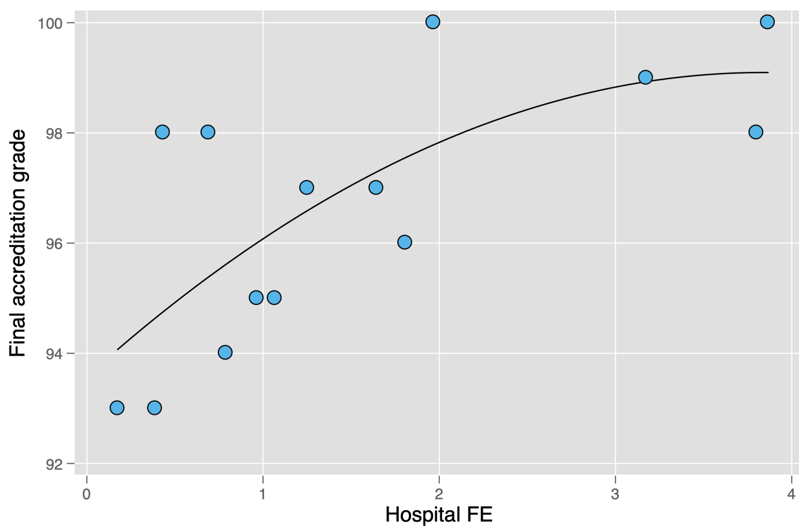

This figure plots the hospital fixed effects from column (2) of Table A.1 against the final accreditation grade of each hospital. The latter is a score assigned by the regulator after they do a thorough inspection on each hospital. This is mandatory in order to be able to operate in the country. Each dot is an observation and the line is a quadratic fit.

follows:

$$EU_{ijkt} = \sum_{d \in D} \gamma_{di} \log \sum_{h \in H} \exp(\alpha_i^H (c_{jh} p_{ijhdt} - c_{j0} p_{ij0dt}) + \beta_i^H Dist_{ih} + \delta_h - \upsilon_{l(i)}) \quad (21)$$

where γ_{di} is the probability of consumer i of being diagnosed with condition d. In practice, this is calculated using a frequency estimator of admission probabilities for each diagnosis for multiple demographic groups. EU_{ijkt} is the variable that is then used in the plan choice demand model of section 1.5.

Finally, I also use this model, and the price and quantity variables constructed above, to measure expected costs for each enrollee in each insurance plan. In particular, let s^H_{ijhdt} be the probability of consumer i in period t with plan j of choosing hospital h to treat diagnosis d. Then, $\rho_{ijhdt} = \tilde{P}_{a,i,j,t,h} \times s^H_{ijhdt}$ is the weighted price that this consumer will pay to treat her diagnosis. In order to get the total payment, I have to sum over each hospital and then multiply this by my quantity measure:

$$C_{i,j,d,t} = \tilde{Q}_{a,i,t} \sum_{h \in H} \rho_{ijhdt} \quad (22)$$

I weight this by γ_{di} in order to get the final measure of expected costs for consumer i in plan j in year t:

$$\tilde{c}_{i,j,t} = \sum_{d \in D} \gamma_{di} C_{i,j,d,t} \quad (23)$$

This is the variable I use in the simulation to construct profits for each insurer. Notice that, as opposed to Shepard (2022), there is no problem in creating this variable because enrollees in Chile are able to access every hospital, no matter the plan they have, thus, within each insurer, I have enough observations to know the average costs for each group of consumers in each plan.[104]

D Sample Construction

Several adjustment are made to the dataset in order to get the sample of enrollees that are actually used for estimation. This section describes the steps that are taken to do this.

1. I keep only individuals that are enrolled in individual plans, with contracts under *open* insurers (*i.e.* enrollment is not limited to specific industries) and whose plans are subject to the standard pricing regime.[105]To have a clean panel, I also drop from the sample those enrollees that at some point between 2007 and 2016 move to any of those plans or to a *closed* insurer.

2. I drop policyholders younger than 20 years old and older than 80 years old (less than 2% of policyholders), and non-employees, such as independent workers or retirees (less than 17% of policyholders). The reason behind the last restriction is that these individuals do not have reliable income data. For the same reason, I also drop observations with invalid or missing wages.

3. I focus only on enrollees in Santiago, the capital of Chile. This is due to the fact that it is hard to know how the choice set looks for someone living in a

[104]The problem in Shepard (2022) is that the hospital network *within* a plan is changing, and, so, it is less realistic to create an expected cost variable that can measure the change in costs after that change. This is because the author does not know how consumers will change their behavior once they are allowed to go to a different hospital (*e.g.* moral hazard).

[105]Firms also offer other type of plans such as partnership plans or employer-based plans. These other plans normally follow different pricing rules, among other differences, so I drop them from the analysis. These plans account for less than 23% of all policyholders in the market.

different city. They might want access to their local hospitals, but also access to a high quality hospital in the capital in case of a major surgery. Therefore, to avoid choice set misspecification, I drop them from the analysis. Importantly, however, most policyholders live in Santiago (over 60% of them).

4. Finally, for computational reasons, the estimation of the demand model, described in section 1.5, uses a 10% random sample. For the simulation of the policy in section 1.6 I use a 5% random sample.

E Health Insurance Plans and Cream-Skimming

In this section, I explore the creation of new plans as a channel used by firms to respond to the regulation banning gender-based pricing. This is important because if companies use this as a way to practice "cream-skimming" in the private market, then the results of the simulation in section 1.6 will be biased (*e.g.* companies create expensive but low quality plans only for females). Luckily, data on plans are publicly available from the regulator even after 2020. However, as noted in section 1.4.1, these data do not have detailed information regarding real coverage rates for each plan. Instead, as a measure of quality, I use "plan scores" provided by the platform *QuePlan.cl*.

According to the regulator, the main reason to create new plans in the market is to respond to the fact that companies must have plans available with prices close enough to the 7% of their potential enrollees' income. Moreover, plan quality is by far the strongest predictor of base premiums. If firms use the creation of new plans as a channel to discriminate after implementing the ban, I should observe a change in the relationship between plan quality and base premiums. The intuition is that they would create new plans solely for women with low quality and higher base premiums, thus disrupting the positive, and strong, relationship between these two variables.

Figure A.1 tests this hypothesis descriptively by plotting base premiums of health plans against their plan scores, separately for plans created before and after 2020. A clear positive and strong relationship emerges from the figure, that is, better plans tend to have higher quality. Importantly, the slope of this relationship does not appear to have changed after the policy change, which goes against the hypothesis of firms using the creation of new plans as a way to practice "cream-skimming".

Figure A.1: Base premiums and plan quality

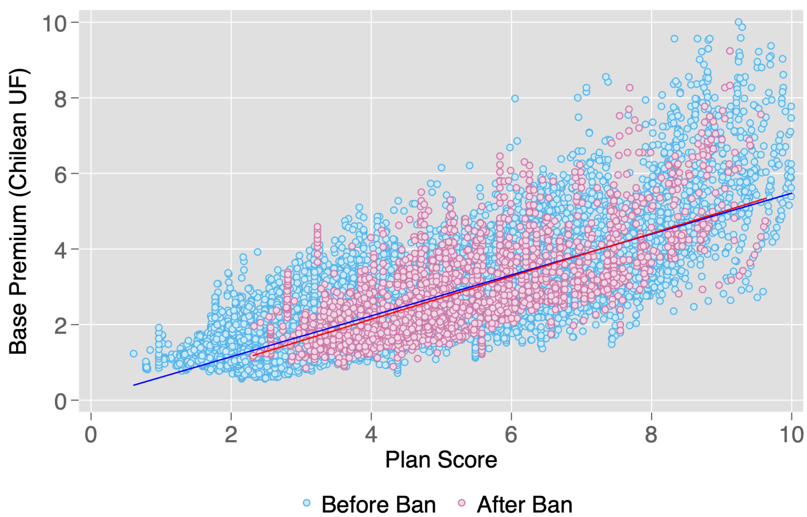

This figure plots the base premiums of plans against their plan scores. Blue dots refer to plans created before 2020 and red dots refer to plans created after 2020. A linear trend is added to each set of plans.

Formally, in Table A.1 I run a regression model with base premiums as the dependent variable and plan scores as the dependent variable. The main variable of interest is an interaction of plan scores with a dummy equal to one for plans created after 2020. Additionally, I control for insurer FE, plan characteristics and year FE. The main results of the table are the following. First, as mentioned above, higher plan scores are the main predictor of base premiums. This can be seen by comparing the R-squared among regressions. Second, plans created after the ban do not appear to have different base premiums than plans created before. Third, and most importantly, there is no evidence of a change in the relationship between base premiums and plan scores. The coefficient in the interaction is negligible and statistically non-significant, which is evidence that companies do not use the creation of plans as a way to practice "cream-skimming".[106]

Table A.1: Regression results - Base premiums and plan quality

	(1)	(2)	(3)	(4)	(5)	(6)
Plan score	0.542	0.536	0.564	0.574	0.574	0.597
	(0.019)	(0.015)	(0.015)	(0.017)	(0.021)	(0.021)
1(plan after ban)				-0.240	-0.228	-1.310
				(0.274)	(0.160)	(0.073)
Plan score x 1(plan after ban)					-0.002	-0.032
					(0.074)	(0.073)
Firm FE	No	Yes	Yes	Yes	Yes	Yes
Plan Characteristics	No	No	Yes	Yes	Yes	Yes
Year FE	No	No	No	No	No	Yes
R^2	0.63	0.66	0.67	0.67	0.67	0.71
Observations	18,866	18,866	18,866	18,866	18,866	18,866

Notes: This table shows the results of a regression of plan scores on a dummy equal to one if the plan was created after the banning of gender-based pricing, controlling for base prices, plan characteristics, company fixed effects, and date fixed effects. Each column is a different specification with different controls. Standard errors are in parenthesis and are clustered at the firm level.

Anecdotal Evidence Post-Regulation

[106] I find similar results if I use logged variables instead.

Some additional anecdotal evidence appears to suggest that changes in prices have been the main mechanism used by firms in response to, among other things, the ban of gender-based pricing. First, firms are not currently allowed anymore to create plans that do not cover maternity expenses, thus forbidding this type of discrimination. Second, in March 2022, as a reaction to pressure from companies to increase prices due to the ban and COVID-19 related expenses, the regulator announced a cap of 8% on how much firms could raise premiums. The firms complained that the cap was too low given the higher costs in the system, as they wanted to increase prices three times more than the cap. This has started an intense debate in the country where insurance companies are demanding higher prices in order to match higher health care costs or otherwise, according to them, the whole private market will become unsustainable (see La Tercera, 2022).

In conclusion, anecdotal evidence suggest that, as a result of banning gender-based pricing, (i) health care costs in the private market have skyrocketed, and (ii) companies have responded to that pressure by trying to increase prices and not by lowering plans' quality. Thus, this evidence supports the analysis and results of the simulation in this paper. Nevertheless, endogenizing the creation of new plans is an interesting topic for future research.

F Sensitivity Analysis

In this section, I briefly describe how the results change when I modify some of the steps taken in section 1.6.2 in order to implement the simulation of the ban. Particularly, I focus on the level of profits that firms have to keep constant, the percentage of people from the public option that are active in each iteration of the simulation, and the percentage of female and old enrollees from the private market that are active right after implementing the ban.

Figure A.1 below shows the percentage change in prices and the change in consumer surplus, relative to the baseline without the ban, for different scenarios. Each row represents what percentage of the original profits used in Table 1.5 firms have to keep constant throughout the simulation. Therefore, the first row is the status quo and each row below that is a lower level of profits. In the case of the columns, they represent the percentage of people from the public option that are active in each iteration of the simulation. Hence, the first column is the status quo and the second column represents a case in which (a random) 50% of the people from the public option are active. Additionally, in this case, I only allow 50% of female and old enrollees from the private market to be active right after implementing the ban. This means that the top left block in each plot is the same result as the one documented in Table 1.5.

In the case of lower profits the results are straightforward. Lower profits for firms are just a transfer to consumers. Hence, it is not a surprise that consumer surplus increases as profits go down. Importantly, 90% of the original profits (the third row) is around the lowest level the profits have ever been for these companies during the years of my data, so any level lower than that would mean that firms would obtain less profits that they have ever done before. As expected as well, prices go down as I reduce profits to the point where the overall price change would be negative if profits are too low.

The results are more interesting when I only allow half of the population in the public option to be active in each iteration, and only 50% of female and old enrollees to be active right after implementing the ban. On the one hand, when profits are close to the original ones, surplus is lower and prices are higher as less people are active. On the other hand, when profits are low, surplus is higher and prices are lower as less people are active. In the first scenario, the intuition is that some young females from the public option are not able to enter the private market because they just never become active but a large portion of young males are still leaving the private market. Therefore, the share of old people in the private system is higher than otherwise, and these groups have higher health care costs than young females, explaining then the higher prices and lower surplus. In the second scenario, still some young females remain in the public option, but now young males, which are less costly than young females, are able to stay at higher rates in the private market, and this force dominates the results.

In sum, the main takeaway of the sensitivity analysis is that, as profits are lower, the benefits from banning gender-based pricing will be higher. This is not surprising as it is just a transfer from firms to consumers. Additionally, if less people from the public option are active in each iteration, the benefits from banning gender-based pricing will be lower if profits are close to the original ones and the opposite is true if profits are lower.

G Model of consumer choice under product unavailability

Many models can explain changes in purchase behavior due to stockouts, although most of them would predict transitory, not permanent, changes. One model that can rationalize the persistent effects of stockouts on purchase behavior is based on

Figure A.1: Sensitivity to lower profits and % of active people

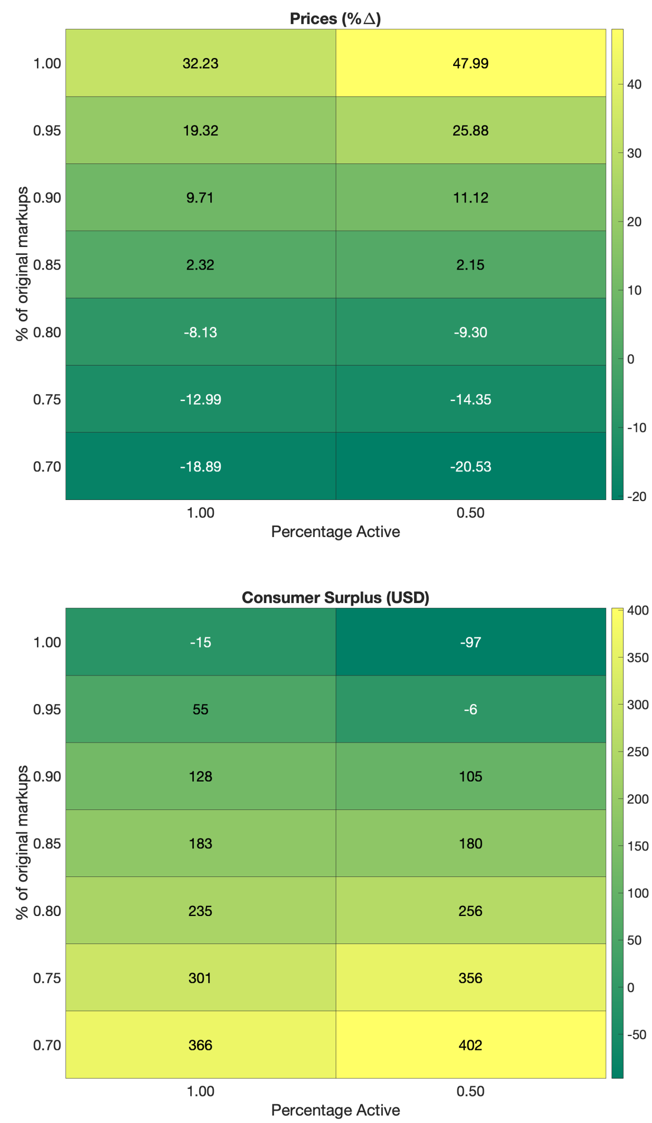

This figure shows prices (top plot) and consumer plus (bottom plot) for multiple scenarios regarding lower profits and lower percentage of people active from the public option. Specifically, each row represents what percentage of the originals profits used in table 1.5 are being considered. Therefore, the first row is the status quo and each lower row is a lower level of profits. In the case of the columns, they represent the percentage of people from the public option that are active in each iteration of the simulation (and also the percentage of female and old enrollees from the private market that are active right after implementing the policy). Hence, the first column is the status quo and the second column represents a case in which (a random) 50% of the people from the public option are active. Numbers inside each block are the percentage change in prices and the change in consumer surplus for each scenario, relative to the baseline without the ban.

long-lasting changes in the consumer's consideration set.

As in Gensch (1987), we model consumer choice as a two-stage process in which brands are first screened and then evaluated for the actual purchase. In the first stage, consumers reduce the relevant information by eliminating alternatives (among those they are aware of) until consumers can deal comprehensively with a smaller set of options. In the second stage, consumers thoroughly compare the subset of alternatives for selection (Shugan, 1980).

The literature on two-stage choice models distinguishes between *brand awareness* (i.e., recalling a brand during a purchase or consumption occasion) and *brand consideration*, which is related to the consumer's endogenous deliberation process before making a brand choice (Keller, 1993). Consistent with this approach, the consumer is only aware of a subset of all products and their expected match valuations. Importantly, the consumer has no information on the products outside her awareness set. In terms of the available options in the awareness set, denoted by S^h_{at}, the consumer will decide ex-ante in the first stage, how much product information to acquire. Thus, the first stage optimization determines the subset of options to be inspected in the second stage (Bronnenberg et al., 2019; Roberts, 1989; Roberts and Lattin, 1991). The optimal consideration set, denoted by $\mathbb{S}^h_{ct}$, then solves the following expected utility maximization problem:

$$\mathbb{S}^h_{ct} \quad = \quad \arg\max_{s^h_{ct} \subseteq S^h_{at}} \left\{ \mathbb{E}(\max_{j \in s^h_{ct}}(u^h_{jt})) - \mathcal{C}(s^h_{ct}) \right\} \tag{24}$$

where $\mathcal{C}(s^h_{ct})$ is the cost of product evaluation associated with assessing the consideration set s^h_{ct}; and u^h_{jt} is the standard utility of alternative j in period t for household

h.

We argue that the extended unavailability of leading brands may change consumers' awareness set, S^h_{at}. In effect, when facing the empty shelves of the most popular products, the competing products previously ignored are now among the only available alternatives. Thus, some previously unaware consumers may then be induced to start learning about the attributes of less popular brands.

Thus, we conjecture that the extended stockouts may have substantially changed consumer awareness of beer brands. The inclusion of new products in S^h_{at} might have temporary effects if the entrant products were perceived as worse than the unavailable top brands. If so, once the shortage episode is over, the leading brands could recapture their pre-stockout market shares. If, however, the new products in the awareness set compare favorably relative to the initially unavailable leading brand products, then the change in purchase behavior can be long-lasting, and in the limit, permanent.

H Data construction for estimation

This section lays out the main steps taken in order to create the sample used in the estimation of the structural model described in section 5. Please note that these steps (particularly steps 1 and 2) apply *only* to the estimation of the structural demand model. Specifically, we implement the following steps:

1. We drop 10.51% of the store visits, which correspond to trips with multiple SKU purchases. This step also leads to the removal of 6 households (0.1%) from our dataset because their visits only included multi-SKU purchases. Accordingly, after this step, we are left with 5,668 of the 5,674 consumers.

2. We only model one choice per week for every consumer. In the case of consumers buying the same product in multiple store visits within a week, we only consider the earliest visit in that week (affecting 5.22% of household-week-product combinations). In the case of consumers buying multiple products within a week, we randomly select one choice (affecting 10.29% of household-week combinations).

3. We determine prices using transaction data to create a panel data for each alternative-week-store combination.

4. We build the availability variable that indicates whether each specific leading brand product was purchased during a particular date in a given store.

I Analysis of Escudo Other formats

In this appendix section, we explore whether the less popular formats within the leading brand Escudo can display the same change in consumer behavior documented for the small brands products. In effect, from the summary statistics in Table 3.1, we see that the market share actually increased after the frequent stockout period from 7 to 8.2 percent.

Figure D.1 shows that 421 consumers purchased the Escudo Other format for the first time in our data during the stockout period. This pattern is similar to those shown in Figure 3.4. Also, similar to the purchase behavior of first timers of small brands depicted in Figure 3.5, the consumers who have tried "Escudo Other" during the pre-treatment period, were not majorly affected by the stockouts of that product, as shown in Figure D.2.

Whether the product awareness mechanism we introduced in the theoretical section is taking place at the brand or brand-format level is not theoretically clear. We think, this interesting feature can help us to rationalize the estimated positive effect for the Escudo Other product. Arguably, the frequent stockouts of Escudo bottle and Escudo might explain the behavior of the 421 first-time purchasers of Escudo Other formats, leading to the same persistent phenomenon we documented for the small brands.

Figure D.1: Purchase behavior of treatment first-timers of Escudo Others over time

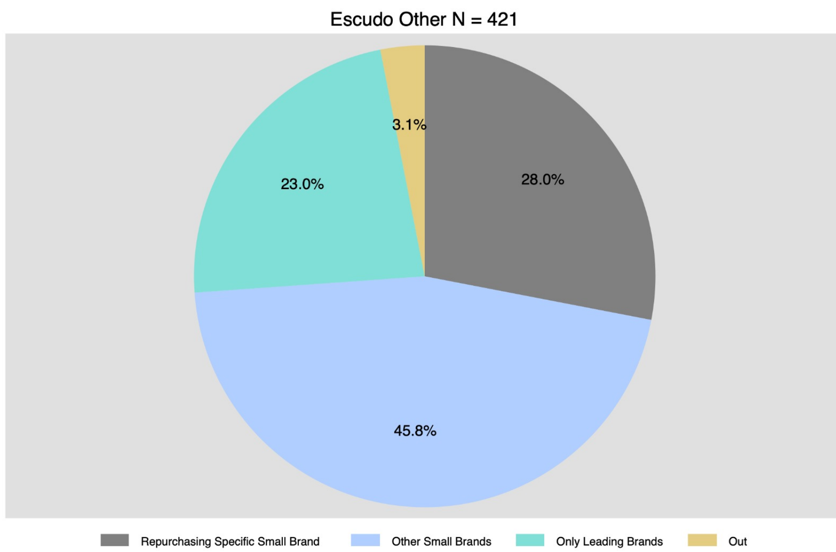

Notes: The figure shows the purchase behavior over time of the subset of consumers who purchased Other Escudo for the first time during the treatment period. The figure shows that 72 percent of these first-timers do not repurchase Escudo Other during the post-treatment period, whereas 28 percent kept buying Escudo Other during the post-treatment period.

Figure D.2: Purchase behavior of pre-treatment first-timers of Escudo Other formats

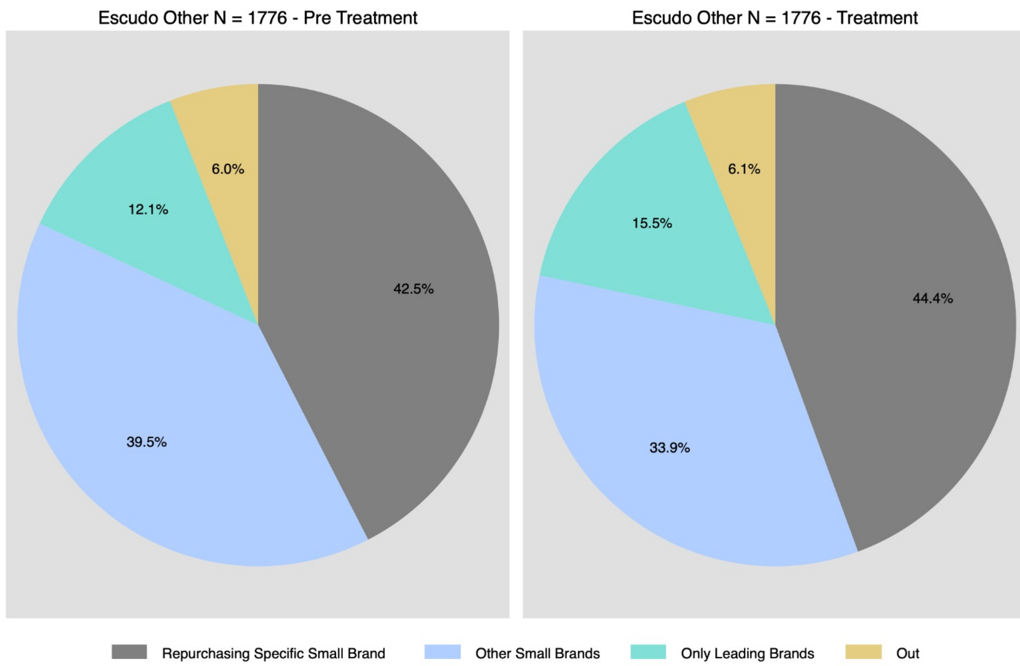

Notes: The figure shows the purchase behavior over time of the subset of consumers who purchased Escudo in Other format for the first time during the pre-treatment period. The figure is consistent with the stockouts (during the treatment period) not increasing the numbers of post-treatment consumers.

J Additional tables and figures

Figure D.1: Risk-rating factors

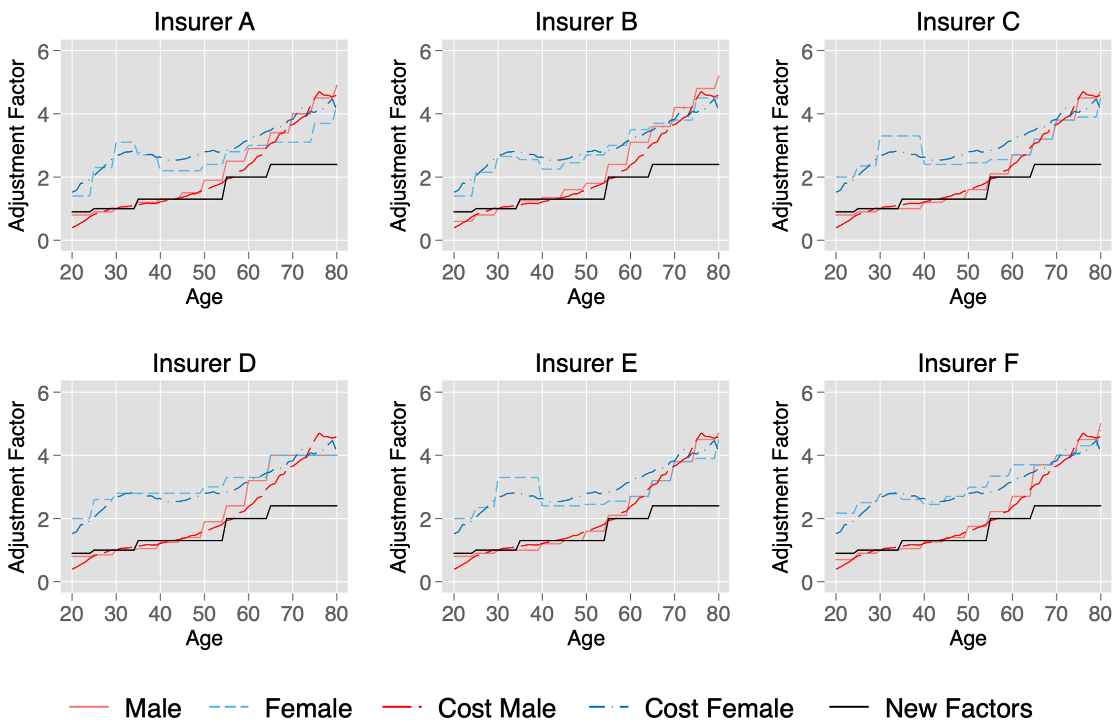

Notes: This figure shows the old risk-rating factors by gender for the six companies, health care costs (relative to a 30-years old male) across enrollees in the private market and the new risk-rating factors implemented in 2020.

Table D.1: Descriptive statistics for hospitals

	Raw Data	Hospital Model		Market
	Avg. Total Payment	Relative Price	Relative Severity	Shares
Hospital 1	7.31	1.89	0.90	5.61%
Hospital 2	6.15	1.39	1.07	9.61%
Hospital 3	5.06	1.27	0.94	0.60%
Hospital 4	4.47	1.16	0.92	13.81%
Hospital 5	3.95	1.14	0.84	11.06%
Hospital 6	4.49	1.14	0.81	1.47%
Hospital 7	3.86	1.08	1.04	2.78%
Hospital 8	4.30	0.94	1.17	5.53%
Hospital 9	3.16	0.88	0.93	4.23%
Hospital 10	2.99	0.77	1.02	15.85%
Hospital 11	2.52	0.71	0.84	3.53%
Hospital 12	2.67	0.69	1.05	4.13%
Hospital 13	2.30	0.69	0.90	4.08%
Hospital 14	2.74	0.66	1.10	7.27%

Notes: This table shows descriptive statistics for the main 14 private hospitals in Santiago using the hospital admissions dataset. The first column includes the hospitals, which are anonymized in the data. The second column documents average total payment per hospitalization. The third and fourth column show the decomposition of these total payments in prices and quantity (severity), which are standarized such that the mean across hospitals is 1. See section C for more details about this decomposition. Finally, the last column focuses on the market shares for each hospital. All prices are measured in thousands of U.S. dollars for December, 2016.

Figure D.2: Risk-Rating factors before regulation (for one firm) versus after regulation

Age groups	Policyholders	
	Male	Female
[0-2)	1.84	1.86
[2-5)	0.87	0.81
[5-10)	0.70	0.66
[10-15)	0.60	0.65
[15-20)	0.66	0.95
[20-25)	0.73	1.87
[25-30)	0.87	2.29
[30-35)	1.00	3.01
[35-40)	1.05	2.91
[40-45)	1.26	2.43
[45-50)	1.38	2.47
[50-55)	1.72	2.65
[55-60)	2.22	2.92
[60-65)	2.79	3.12
[65-70)	3.63	3.42
[70-75)	4.03	3.70
[75-80)	4.66	4.08
[80-	4.97	4.55

Age groups	Policyholders
[0-20)	0.60
[20-25)	0.90
[25-35)	1.00
[35-45)	1.30
[45-55)	1.30
[55-65)	2.00
[65-	2.40

Notes: This figure shows the risk-rating factors before the regulation for one firm (on the left) and the new risk-rating factors (on the right).

Table D.2: Market shares by gender and income

Females	Low-income	Medium-income	High-income	Total
Insurer A	17%	18%	22%	19%
Insurer B	26%	25%	24%	25%
Insurer C	3%	2%	5%	3%
Insurer D	18%	23%	21%	21%
Insurer E	20%	17%	18%	18%
Insurer F	17%	15%	10%	14%
Males	**Low-income**	**Medium-income**	**High-income**	**Total**
Insurer A	10%	12%	16%	13%
Insurer B	22%	22%	20%	21%
Insurer C	2%	2%	6%	3%
Insurer D	13%	17%	18%	16%
Insurer E	21%	23%	24%	23%
Insurer F	32%	24%	16%	24%

Notes: This table shows market shares for the six open insurers in December of 2016 by gender and income group.

Figure D.3: Historical market shares across segments

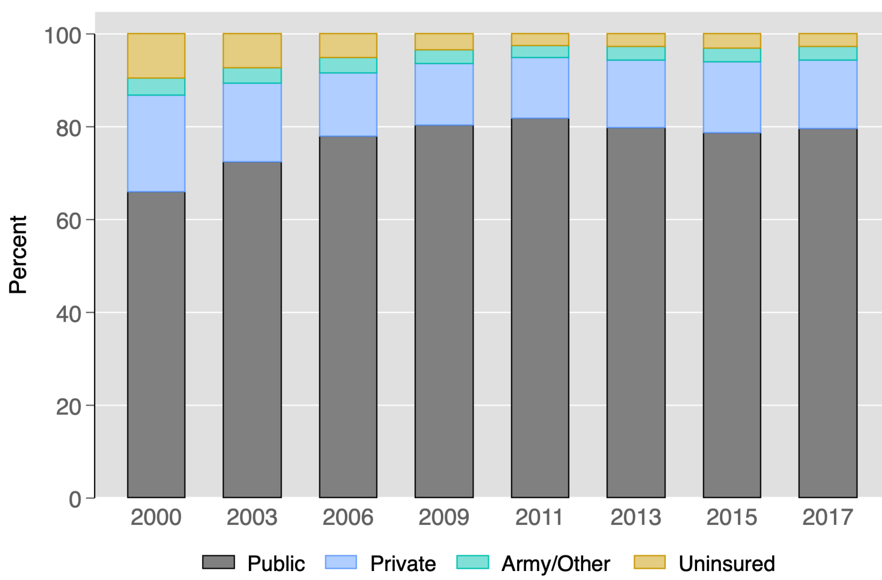

Notes: This figure shows the historical market shares across segments in the health insurance market in Chile. The data come from the Chile National Socioeconomic Characterization Survey (CASEN).

Figure D.4: Base price changes 2013/2014

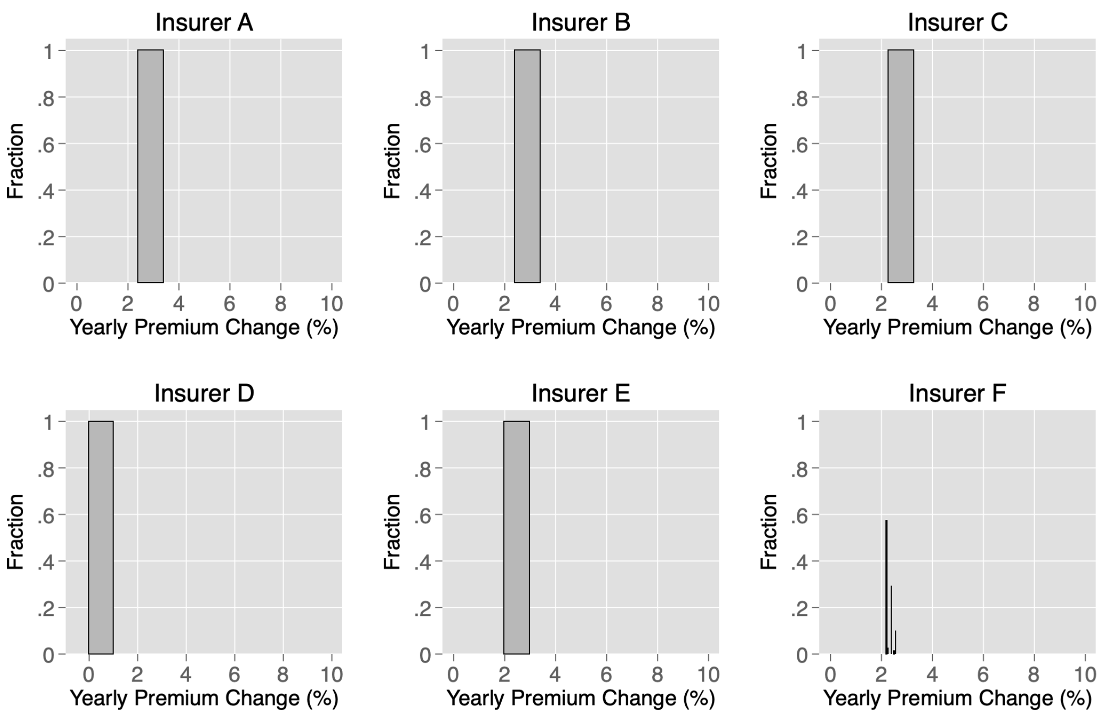

Notes: This figure shows an histogram of the annual base price changes for the six insurance companies in Chile for the period 2013/2014. Plans with less than 50 policyholders by January 2013 are excluded from the figure.

Figure D.5: Base price evolution

Notes: This figure shows the evolution of base prices for five of the six insurance companies in Chile for the period 2013-2018. The company missing filed for bankruptcy in 2017. Plans are divided into terciles according to their plan scores. Chilean UF is the unit, indexed to inflation, in which base prices are measured in Chile.

Figure D.6: Tenure in GR contracts

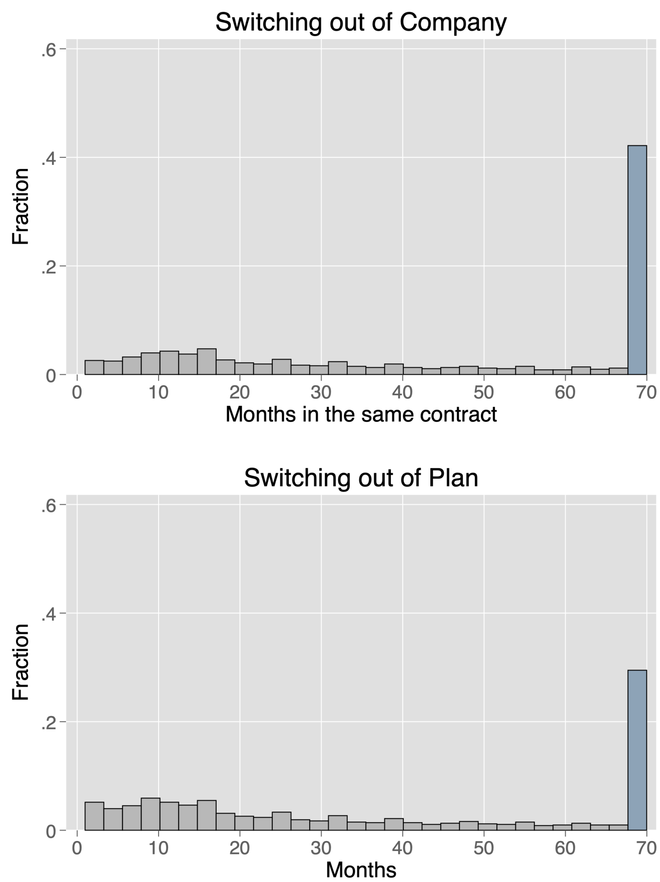

Notes: This figure shows histograms of how many months policyholders stay in their contracts. The top figure shows this for the case where switching plans within the company is not considered. The bottom figure shows the case in which any switching is considered. This exercise is done for policyholders that signed a new contract in May 2011.

Figure D.7: Plans before and after switching

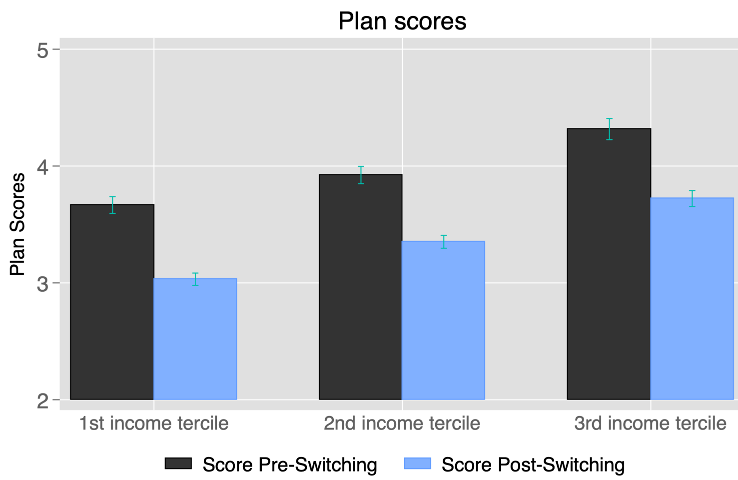

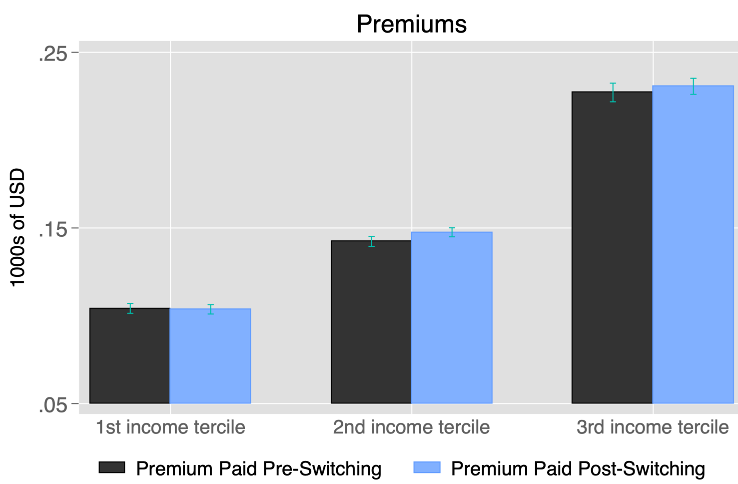

The upper figure shows the plan scores before and after switching during the signing month of the contract. The lower figure shows premiums paid by policyholders before and after switching during the signing month of the contract. Green lines report the 95% interval.

Figure D.8: Probability of leaving the private market due to price changes

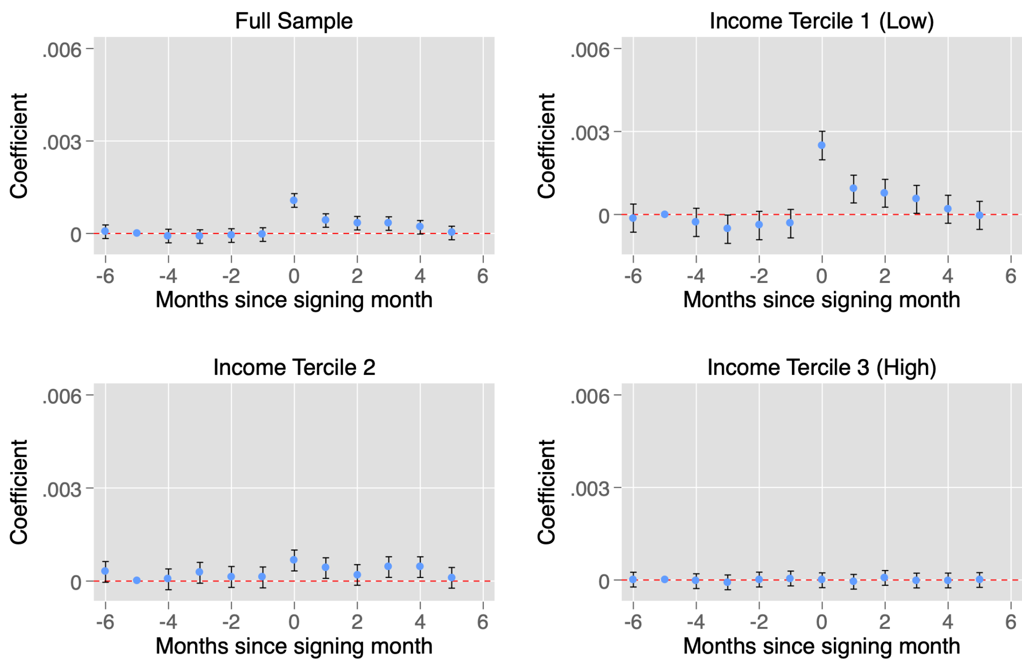

Notes: This figure shows an event study regression where the dependent variable is a dummy equal to one if a consumer leaves the private market and do not re-enter in the future. The event is the month in which price changes are applied to health plans. Controls include individual fixed effects and date (month-year) fixed effects. I restrict the estimation sample to policyholders that are active in the data for at least 12 months and that do not leave the private market and re-enter in later dates. Additionally, I drop individuals with zero or missing income at any month. This exercise is done on a 10% random sample of policyholders.

Figure D.9: Probability of switching plans due to price changes

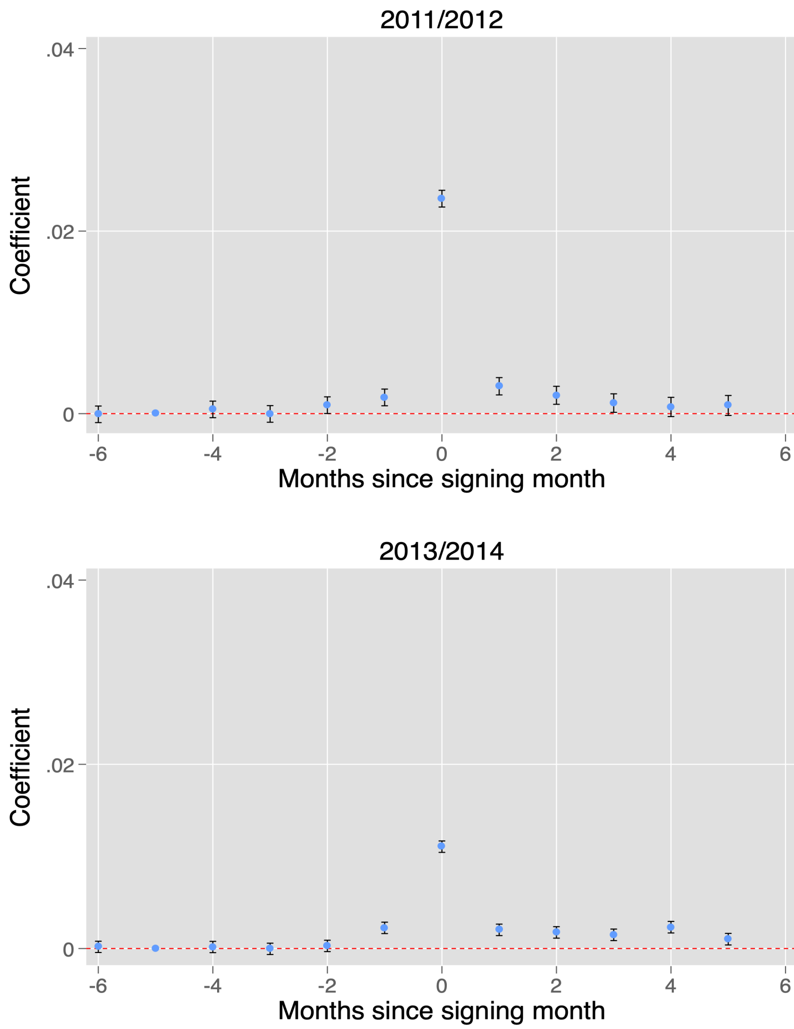

Notes: This figure shows an event study regression where the dependent variable is a dummy equal to one if a consumer switches plans within an insurer. The event is the month in which price changes are applied to health plans. Controls include individual fixed effects and date (month-year) fixed effects. I restrict the estimation sample to policyholders that do not switch insurance companies and that do not leave the private market and re-enter in later dates. Additionally, I drop individuals with zero or missing income at any month. This exercise is done on a 10% random sample of policyholders.

Figure D.10: Probability of switching plans due to price changes - Demographics

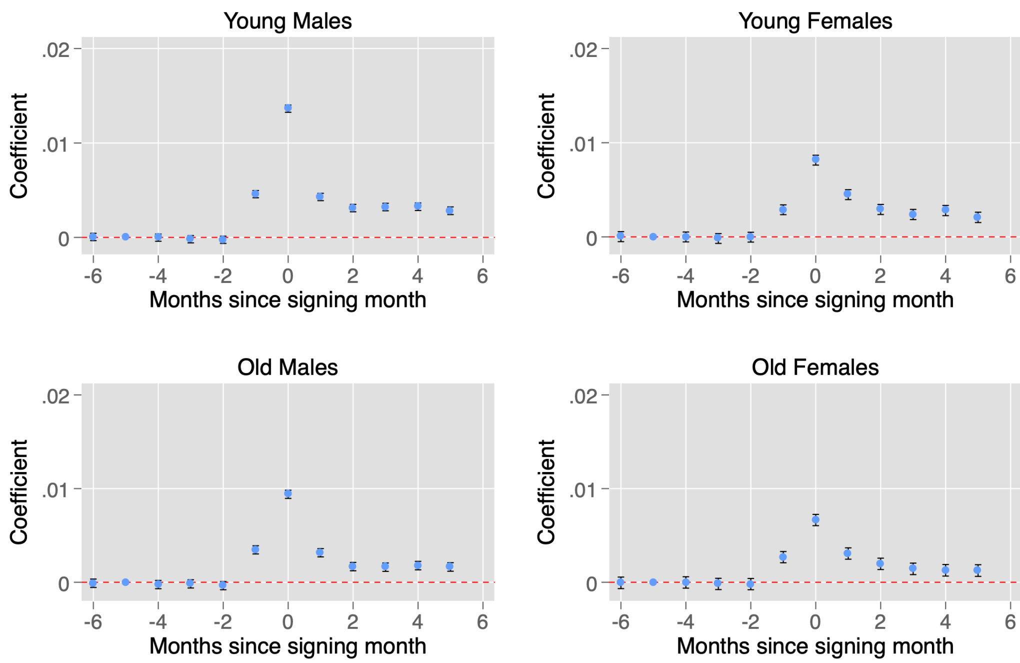

Notes: This figure shows an event study regression where the dependent variable is a dummy equal to one if a consumer switches plans within an insurer. The event is the month in which price changes are applied to health plans. Controls include individual fixed effects and date (month-year) fixed effects. I restrict the estimation sample to policyholders that do not switch insurance companies and that do not leave the private market and re-enter in later dates. Additionally, I drop individuals with zero or missing income at any month. This exercise is done on a 10% random sample of policyholders. Enrollees are split according to their age and gender.

Table D.3: Regression results - Premiums paid

log(premiums)	(1)	(2)	(3)	(4)	(5)
Plan score	0.080	0.101	0.133	0.070	0.147
	(0.010)	(0.022)	(0.008)	(0.010)	(0.015)
Gender					0.441
					(0.035)
Firm FE	No	Yes	Yes	Yes	Yes
Plan Characteristics	No	No	Yes	Yes	Yes
Income Deciles FE	No	No	No	Yes	Yes
Age FE	No	No	No	No	Yes
R^2	0.13	0.28	0.39	0.56	0.83
Observations	55,521	55,521	55,521	55,521	55,521

Notes: This table shows the results of a regression of log of premiums paid on plan scores, firm fixed effects, plan characteristics, income deciles fixed effects, gender and age fixed effects. Each column is a different specification with different controls. The sample is composed by policyholders signing a new contract in the private market in 2016. Standard errors are in parenthesis and are clustered at the firm level.

Table D.4: Model fit

	Model			
	(1)	(2)	(3)	(4)
Log likelihood	135,930	133,720	132,350	130,430
LR test against model (4)	11,000	6,580	3,840	-
AIC	271,910	267,670	264,758	261,098
BIC	272,238	269,179	265,139	262,660
Parameters	25	115	29	119
Observations	3,706,628	3,706,628	3,706,628	3,706,628
First Stage	N	N	Y	Y
Plan Characteristics	Y	Y	Y	Y
Insurer FE	Y	N	Y	N
Insurer-Demographics FE	N	Y	N	Y

Notes: This table shows multiple tests to assess which model provides a better fit to the data. Different columns indicate a different specification. Columns (1) and (3) consider insurer fixed effects, and columns (2) and (4) consider insurer fixed effects interacted with household groups. Additionally, columns (3) and (4) include the first stage in the maximum likelihood estimation.

Figure D.11: Premium elasticities

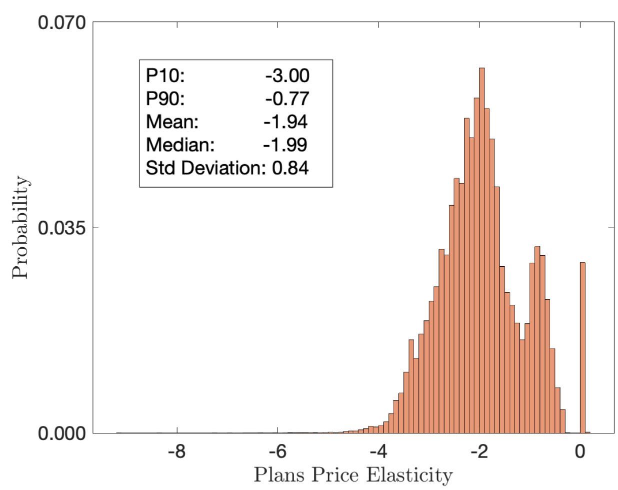

This figure shows the estimated premium elasticities for the chosen plans derived from the demand model in Table 1.4. Elasticities are calculated as $\hat{\eta}_{fjt} = \hat{\alpha}_f p_{fjt}(1 - \hat{s}_{fjt})$, where $\hat{s}_{fjt}$ is the predicted choice probability of plan j by household f in time t.

Figure D.12: Ban of gender-based pricing: advertisement

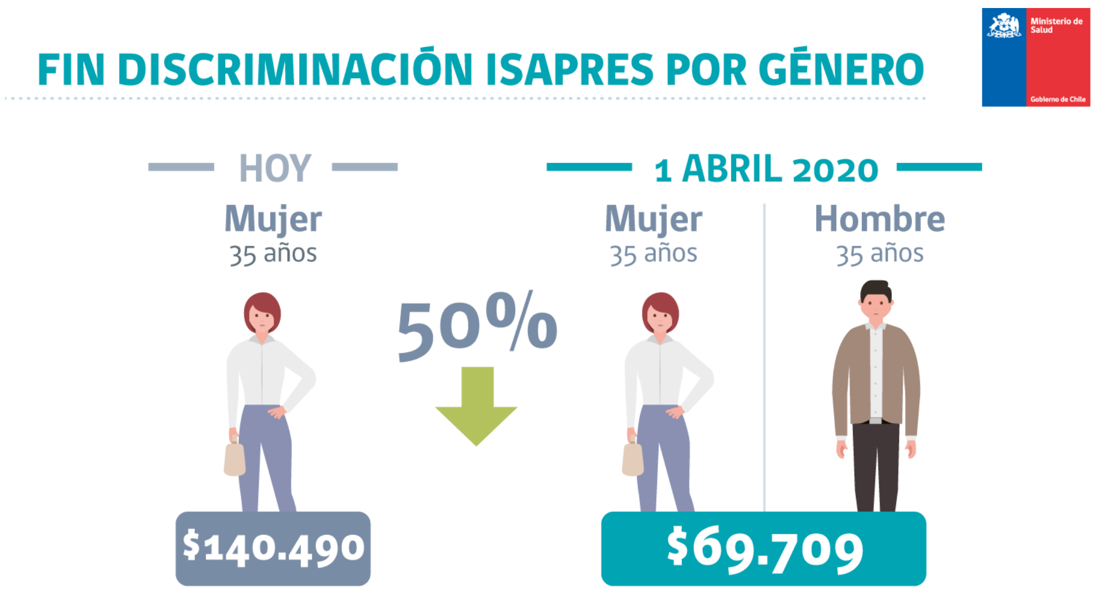

This figure shows an ad from the regulator promoting the ban of gender-based pricing to women. The ad explains that, by switching plans, a woman of 35 years old would pay 50% lower premiums (the same as her male counterpart). Prices are in Chilean pesos.

Figure D.13: Change in consumer surplus - Reclassification risk protection

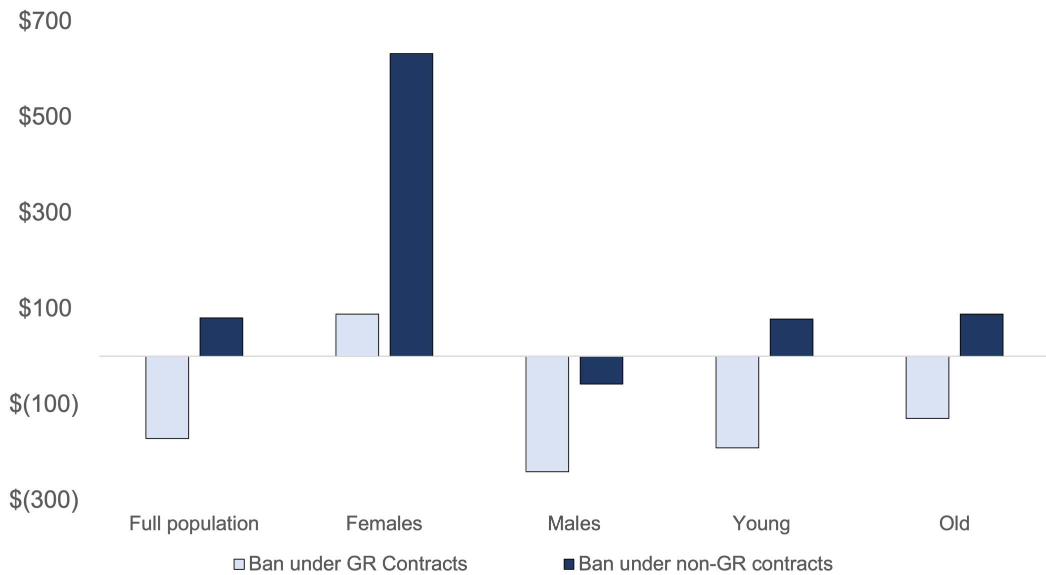

This figure shows the median change in consumer surplus after banning gender-based pricing, relative to the baseline of no ban, for multiple demographic groups. The numbers are in USD of 2016.

Figure D.14: Front-loading in GR contracts

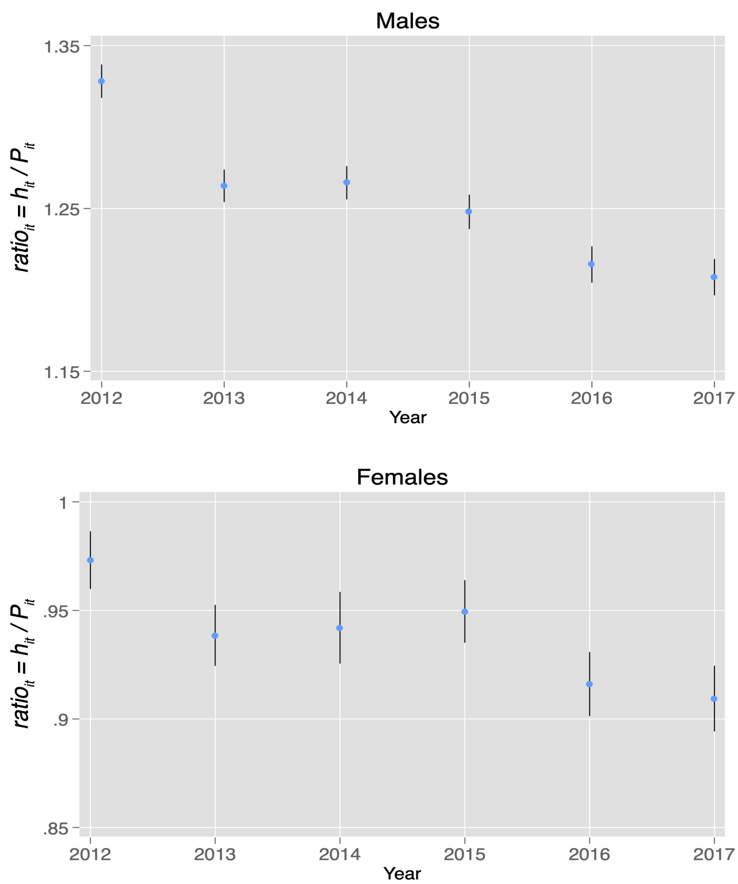

Notes: This figure shows a scatter plot of $ratio_{it} = \frac{P_{it}}{h_{it}}$, where h_{it} is the expected claims in period t by individual i, and P_{it} the corresponding premium. I use a panel of the sampled single policyholders enrolled in January 2012 and followed until December 2017. The top panel displays female policyholders. The bottom panel displays male policyholders.

Table D.5: Regression - Lapsers and probability of positive spending by income terciles

	(1)	(2)	(3)	(4)	(5)
			1st Income Tercile		
$1\{Lapsers\}$	0.072	0.109	0.110	0.103	0.080
	(0.010)	(0.010)	(0.010)	(0.010)	(0.011)
			2nd Income Tercile		
$1\{Lapsers\}$	0.032	0.047	0.048	0.038	0.038
	(0.009)	(0.009)	(0.009)	(0.009)	(0.009)
			3rd Income Tercile		
$1\{Lapsers\}$	0.031	0.032	0.031	0.019	0.020
	(0.013)	(0.013)	(0.013)	(0.013)	(0.013)
Firm FE	No	Yes	Yes	Yes	Yes
Year FE	No	No	Yes	Yes	Yes
Demographic Characteristics	No	No	No	Yes	No
Policyholder FE	No	No	No	No	Yes

Notes: This table shows the results of a regression where the dependent variable is a dummy equal to one if policyholder i fills a positive number of claims 6 months before and after the signing month, and the main independent variable is a dummy $1\{Lapsers\}$ equal to one if the policyholder lapsed her plan during the signing month in Figure 2.3. Each column is a different specification with different controls. Demographic characteristics include insurer FE, year FE, age FE, gender FE, region of residency FE and policyholder FE. I restrict the estimation sample to single policyholders that do not switch insurance companies and that do not leave the private market and re-enter in later dates. Additionally, I drop individuals with zero or missing income at any month. This exercise is done on a 10% random sample of policyholders.

Figure D.15: Expected actuarial profit - Year-by-year

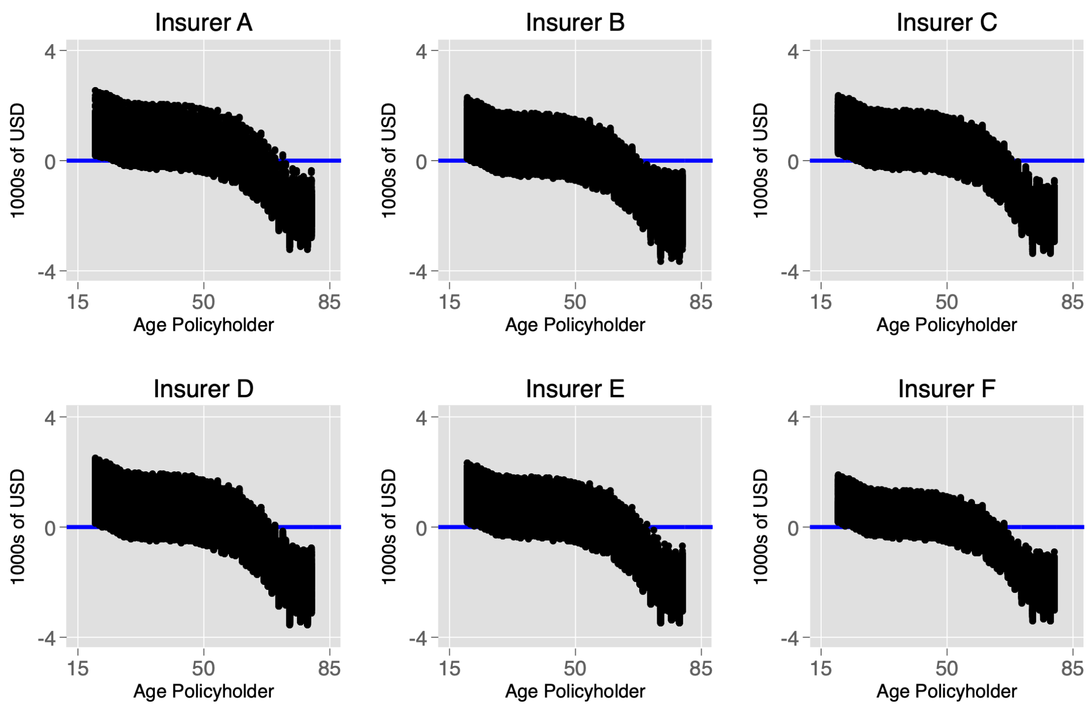

Notes: Each dot in the figure represents the expected actuarial profit for a particular plan of enrolling a male policyholder of a particular age. Data come from the six insurance companies. The price is calculated for a 20 years old male that signed a new contract in 2013. The cost is the expected medical spending by age across active policyholders in 2013 multiplied by the coverage rate of the plan. Profits are measured in U.S. dollars using the exchange rate on December 2013.

Figure D.16: Expected actuarial profit - Cumulative

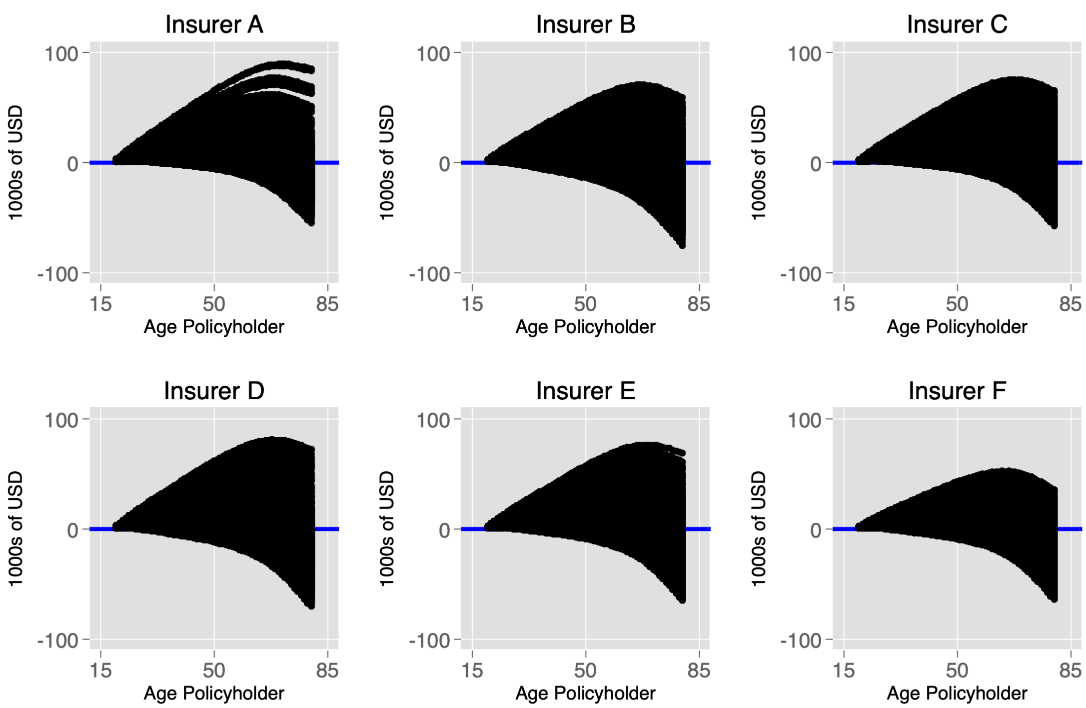

Notes: Each dot in the figure represents the expected actuarial cumulative profit for a particular plan of enrolling a male policyholder of a particular age. Data come from the six insurance companies. The price is calculated for a 20 years old male that signed a new contract in 2013. The cost is the expected medical spending by age across active policyholders in 2013 multiplied by the coverage rate of the plan. Profits are measured in U.S. dollars using the exchange rate on December 2013.

Figure D.17: Expected actuarial profit - Females

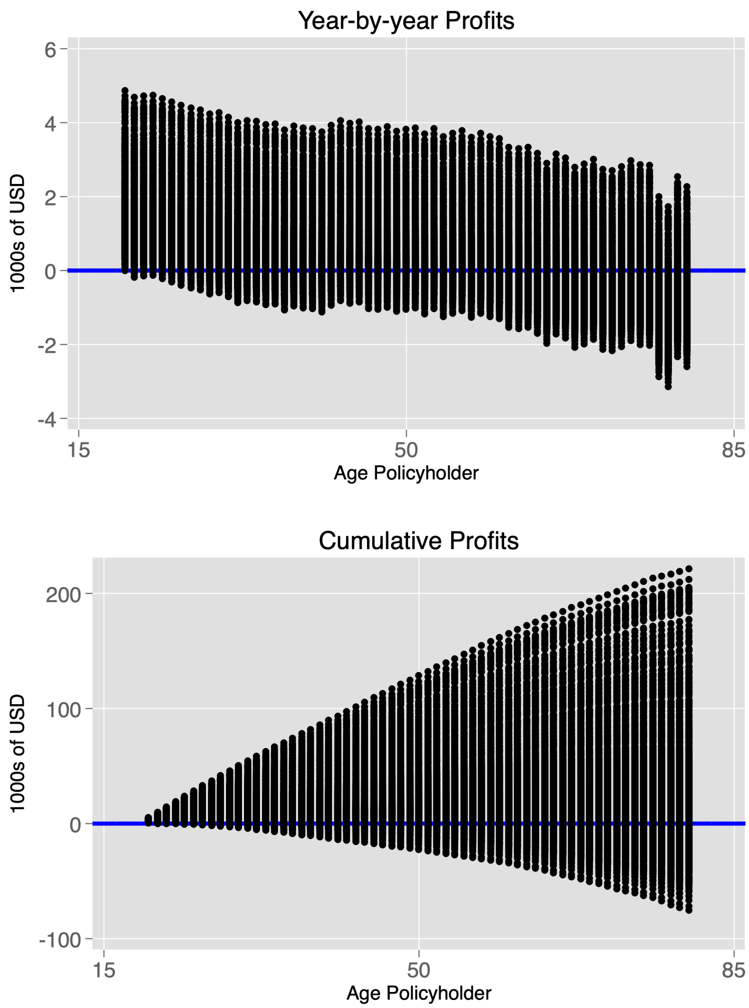

Notes: Each dot in the figure represents the expected actuarial profit for a particular plan of enrolling a female policyholder of a particular age. The upper panel computes annual profits for each age and each plan. The lower panel computes cumulative profits for each age and each plan. Data come from one representative insurance company. The price is calculated for a 20 years old female that signed a new contract in 2013. The cost is the expected medical spending by age across active policyholders in 2013 multiplied by the coverage rate of the plan. Profits are measured in U.S. dollars using the exchange rate on December 2013.

Figure D.18: Realized profit - By demographics

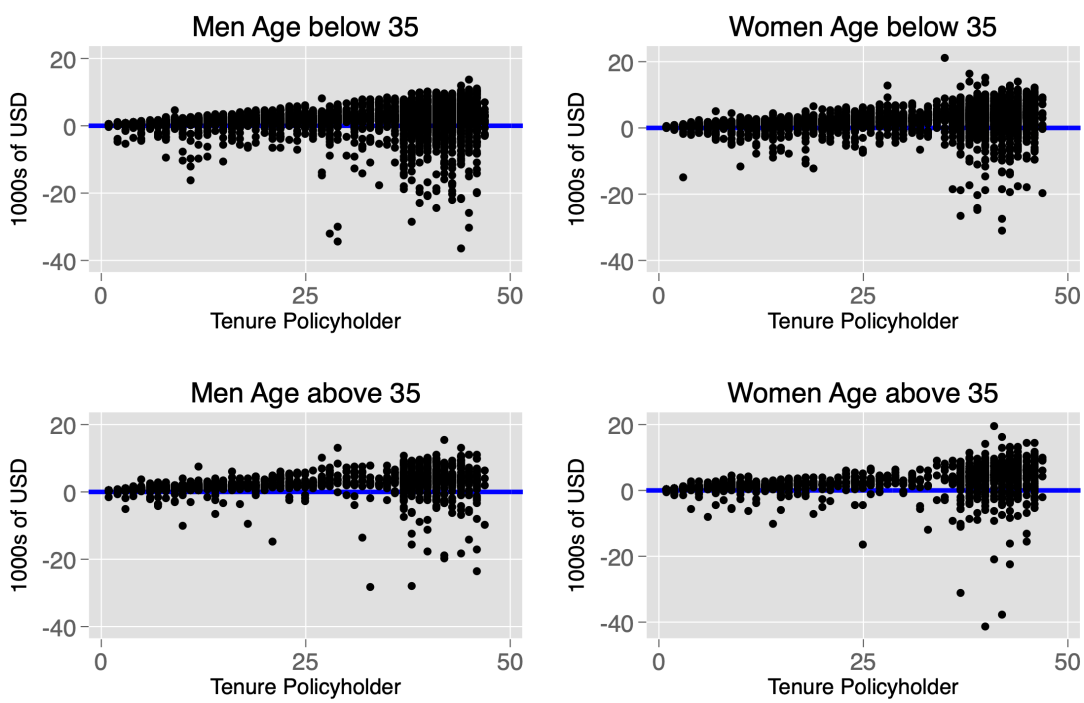

Notes: Each dot in the figure represents the cumulative realized actuarial profit for a particular male policyholder of age less than 35 that signed a new contract in 2013 until he lapses. Tenure measures how many months the policyholder remained in the contract. Data come from one representative insurance company. Profits are measured in U.S. dollars using the exchange rate on December 2013.

Figure D.19: Realized profit

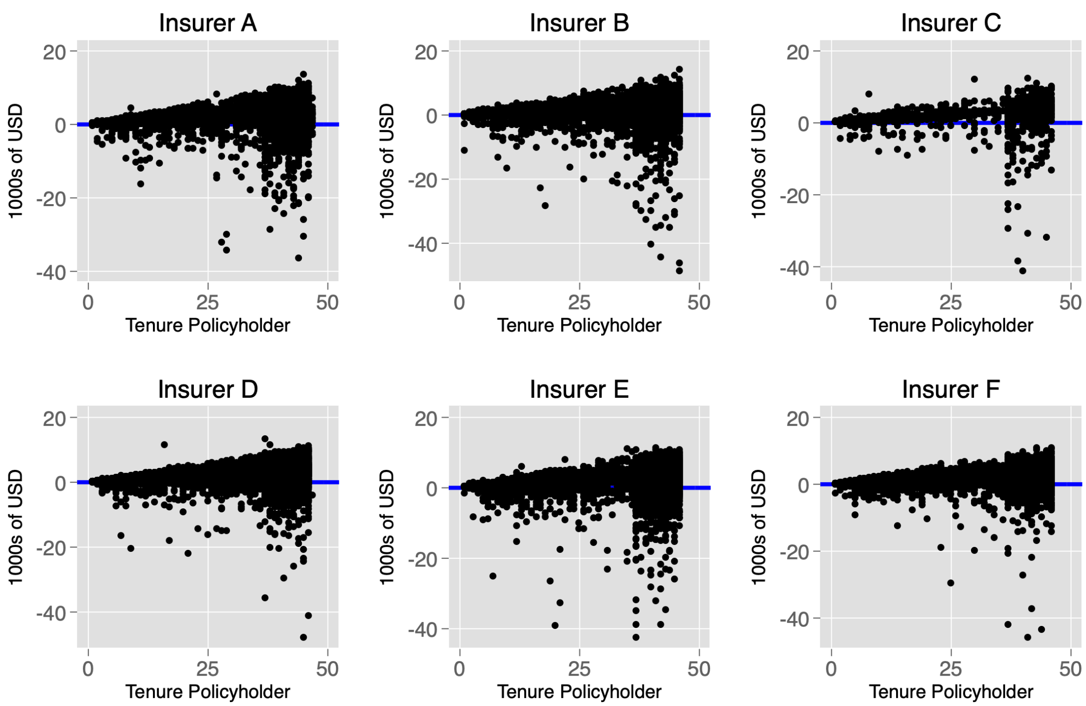

Notes: Each dot in the figure represents the cumulative realized actuarial profit for a particular male policyholder of age less than 35 that signed a new contract in 2013 until he lapses. Tenure measures how many months the policyholder remained in the contract. Data come from the six insurance companies. Profits are measured in U.S. dollars using the exchange rate on December 2013.

Figure D.20: Realized profit - 2014

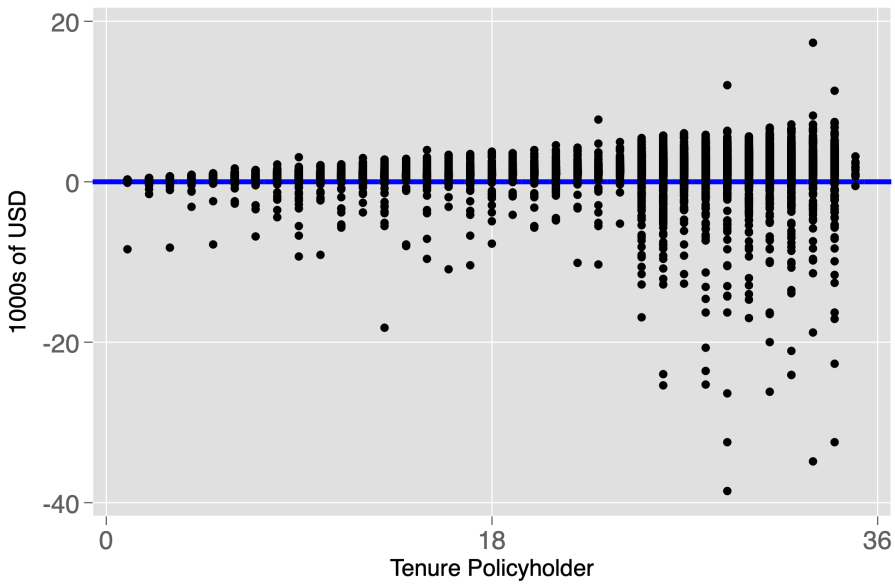

Notes: Each dot in the figure represents the cumulative realized actuarial profit for a particular male policyholder of age less than 35 that signed a new contract in 2014 until he lapses. Tenure measures how many months the policyholder remained in the contract. Data come from one representative insurance company. Profits are measured in U.S. dollars using the exchange rate on December 2014.

Table D.6: Summary statistics for the full sample

Panel A: Leading Brands	Average Price (US Dollars) (1)	Trips (2)	Market Share (Pre-Treatment) (3)	Market Share (Post-Treatment) (4)
Cristal (1L bottle)	1.26	4.3%	4.33%	2.93%
Cristal (350cc can)	0.47	9.9%	8.44%	9.74%
Escudo (1L bottle)	1.25	5.3%	4.90%	3.45%
Escudo (350cc can)	0.47	12.0%	10.98%	12.66%
Other Cristal	0.59	3.7%	4.51%	3.62%
Other Escudo	0.63	4.2%	5.59%	6.19%
All Cristal and Escudo			38.75%	38,59%

Panel B: Small Brands	Average Price (US Dollars) (1)	Trips (2)	Market Share (Pre-Treatment) (3)	Market Share (Post-Treatment) (4)
Baltica (350cc can)	0.39	5.4%	3.40%	4.20%
Becker (350cc can)	0.41	4.7%	3.37%	5.13%
Stella Artois (354cc can)	0.66	2.7%	3.49%	3.76%
Heineken (350cc can)	0.70	6.3%	8.48%	8.34%
Royal Guard (350cc can)	0.62	2.5%	3.00%	3.19%
Other Beers	0.89	39.0%	39.50%	36.78%
All Non Cristal and Escudo			61.25%	61.41

Av. Trips per household	23.0		No. of households	28,005
Av. Leading Brands per household	10.9		No. of Stores	64

Notes: Column (1) shows the average price for each product, Column (2) shows the percentage of purchases for each product, conditional on a beer purchase. Column (3) and (4) are the sales market shares before the Treatment period and after the Treatment period respectively. We consider 28,005 households that have at least ten beer transactions within the initial 21 weeks of data.

Table D.7: Summary statistics of prices and market shares (full sample)

Panel A: Prices	Pre Treatment			Treatment			Post-Treatment		
	Mean	p5	p95	Mean	p5	p95	Mean	p5	p95
Cristal (1L bottle)	1.30	1.15	1.62	1.24	1.16	1.55	1.30	1.15	1.53
Cristal (350cc can)	0.50	0.44	0.59	0.48	0.44	0.56	0.47	0.43	0.55
Escudo (1L bottle)	1.32	1.15	1.71	1.32	1.17	1.62	1.27	1.13	1.59
Escudo (350cc can)	0.49	0.44	0.59	0.49	0.44	0.59	0.48	0.43	0.58
Other Cristal	0.63	0.50	0.77	0.65	0.52	0.73	0.65	0.51	0.74
Other Escudo	0.71	0.52	0.99	0.69	0.53	0.96	0.70	0.52	0.94
Baltica (350cc can)	0.41	0.36	0.49	0.41	0.37	0.47	0.40	0.34	0.47
Becker (350cc can)	0.45	0.40	0.56	0.46	0.41	0.54	0.41	0.37	0.52
Stella Artois (340cc can)	0.71	0.62	0.93	0.73	0.68	0.89	0.63	0.55	0.86
Heineken (350cc can)	0.72	0.67	0.78	0.72	0.68	0.81	0.71	0.65	0.80
Royal Guard (350cc can)	0.66	0.57	0.82	0.65	0.60	0.78	0.64	0.58	0.77
Other Brands/Formats	1.07	0.43	2.16	1.07	0.42	2.17	1.04	0.40	2.07

Panel B: Market Shares	Pre Treatment			Treatment			Post-Treatment		
	Mean	p5	p95	Mean	p5	p95	Mean	p5	p95
Cristal (1L bottle)	5.77	2.04	11.20	5.91	2.64	9.00	4.08	1.76	9.70
Cristal (350cc can)	9.39	5.30	15.72	18.09	11.17	30.48	11.21	6.01	18.68
Escudo (1L bottle)	6.13	2.14	10.55	1.92	0.77	3.60	4.33	1.95	7.27
Escudo (350cc can)	11.74	8.36	17.69	9.46	4.47	14.82	13.61	9.01	19.77
Other Cristal	6.02	1.33	12.20	3.98	1.28	7.83	4.89	1.22	7.99
Other Escudo	7.32	1.30	15.49	6.16	2.74	9.76	8.19	1.80	15.39
Baltica (350cc can)	3.98	2.01	7.56	3.81	1.77	5.36	4.99	2.52	9.08
Becker (350cc can)	3.68	2.05	5.99	4.89	2.22	8.90	5.75	3.06	10.25
Stella Artois (354cc can)	4.17	1.37	7.47	4.88	1.94	7.17	4.25	1.66	6.48
Heineken (350cc can)	8.91	5.35	12.60	6.89	4.71	10.28	8.69	4.44	12.82
Royal Guard (350cc can)	3.38	1.70	5.49	5.01	2.16	8.54	3.60	1.71	6.45
Other Brands/Formats	41.50	24.69	56.94	37.94	20.90	51.44	37.77	20.88	53.34

Panel C: Incidence	Pre Treatment			Treatment			Post-Treatment		
	Mean	p5	p95	Mean	p5	p95	Mean	p5	p95
Cristal (1L bottle)	6.11	2.13	12.16	6.45	2.98	11.96	4.34	1.87	9.32
Cristal (350cc can)	10.38	6.77	15.30	18.85	12.31	29.77	11.42	7.41	17.89
Escudo (1L bottle)	6.88	2.86	12.06	2.24	0.97	4.21	5.03	2.31	8.08
Escudo (350cc can)	12.21	8.76	18.10	9.68	5.52	15.08	13.67	8.68	20.71
Other Cristal	5.10	1.25	8.44	3.52	1.07	6.35	4.39	1.35	7.67
Other Escudo	5.19	1.32	11.23	4.61	2.38	7.76	6.39	1.64	11.14
Baltica (350cc can)	5.62	2.93	10.33	5.48	2.87	8.51	7.32	3.65	12.32
Becker (350cc can)	4.43	2.68	6.50	5.82	2.53	8.68	6.50	3.84	9.94
Stella Artois (354cc can)	3.19	1.16	5.62	3.89	1.78	5.49	3.45	1.33	5.26
Heineken (350cc can)	6.79	3.66	9.38	5.46	3.81	8.09	6.65	3.66	10.20
Royal Guard (350cc can)	2.75	1.30	4.55	4.12	1.89	7.17	2.87	1.53	4.91
Other Brands/Formats	42.40	28.65	54.36	38.25	22.58	49.28	38.77	23.83	51.75

Notes: The table shows the mean prices across transactions (top Panel A), the average value market shares calculated across stores (middle Panel B), and the incidence rate calculated as the average presence in consumer's trip across stores (bottom Panel C). For each period described in Figure 3.1, we report the mean and the percentiles 5 and 95 of the corresponding distribution. The statistics above consider the full sample of 28,005 households with at least ten beer transactions within the pre-treatment period.

Table D.8: Empirical results: estimated posterior mean, standard deviation, 2.5% and 97.5% quantiles for the variance covariance-matrix of the random coefficients Λ.

Posterior Mean

	Cristal B.	Cristal C.	Escudo B.	Escudo C.	Cristal O.	Escudo O.	Baltica	Becker	Stella A.	Heineken	Royal G.	Other Brands	Temp.	State Dep.	ln(Price)
Cristal Bottle	20.92	13.68	13.53	8.78	13.53	8.84	11.36	12.57	14.80	11.38	10.33	13.50	0.32	0.04	-2.38
Cristal Can	13.68	13.69	7.88	8.24	12.03	7.25	9.62	11.38	13.09	10.71	9.55	10.48	0.14	-0.12	-1.87
Escudo Bottle	13.53	7.88	15.71	10.81	7.67	10.87	11.09	10.62	13.04	9.53	8.59	11.93	-0.13	-0.09	-1.91
Escudo Can	8.78	8.24	10.81	12.62	7.07	11.15	10.94	11.00	13.05	10.39	9.33	10.58	-0.24	-0.21	-1.67
Cristal Other	13.53	12.03	7.67	7.07	12.68	7.64	9.21	10.39	11.67	9.41	8.42	9.57	0.06	-0.12	-1.66
Escudo Other	8.84	7.25	10.87	11.15	7.64	11.90	10.32	10.10	11.96	9.26	8.36	9.82	-0.33	-0.22	-1.50
Baltica	11.36	9.62	11.09	10.94	9.21	10.32	18.38	14.11	12.70	9.74	9.02	12.11	-0.60	-0.22	-1.61
Becker	12.57	11.38	10.62	11.00	10.39	10.10	14.11	16.55	14.26	11.24	11.15	12.34	-0.23	-0.18	-1.83
Stella Artois	14.80	13.09	13.04	13.05	11.67	11.96	12.70	14.26	19.36	14.48	13.14	14.48	-0.16	-0.06	-2.26
Heineken	11.38	10.71	9.53	10.39	9.41	9.26	9.74	11.24	14.48	12.79	10.33	11.09	-0.11	-0.10	-1.78
Royal Guard	10.33	9.55	8.59	9.33	8.42	8.36	9.02	11.15	13.14	10.33	11.58	10.27	-0.15	-0.14	-1.57
Other Brands	13.50	10.48	11.93	10.58	9.57	9.82	12.11	12.34	14.48	11.09	10.27	12.83	-0.07	-0.08	-1.92
Temperature	0.32	0.14	-0.13	-0.24	0.06	-0.33	-0.60	-0.23	-0.16	-0.11	-0.15	-0.07	1.53	-0.01	-0.33
State Dependence	0.04	-0.12	-0.09	-0.21	-0.12	-0.22	-0.22	-0.18	-0.06	-0.10	-0.14	-0.08	-0.01	0.27	0.00
ln(Price)	-2.38	-1.87	-1.91	-1.67	-1.66	-1.50	-1.61	-1.83	-2.26	-1.78	-1.57	-1.92	-0.33	0.00	0.42

Posterior Standard Deviation

	Cristal B.	Cristal C.	Escudo B.	Escudo C.	Cristal O.	Escudo O.	Baltica	Becker	Stella A.	Heineken	Royal G.	Other Brands	Temp.	State Dep.	ln(Price)
Cristal Bottle	0.47	0.32	0.34	0.29	0.32	0.32	0.43	0.43	0.49	0.33	0.35	0.31	0.12	0.06	0.06
Cristal Can	0.32	0.29	0.27	0.24	0.27	0.25	0.31	0.32	0.39	0.26	0.29	0.24	0.09	0.05	0.04
Escudo Bottle	0.34	0.27	0.36	0.28	0.29	0.28	0.36	0.34	0.44	0.30	0.29	0.28	0.11	0.05	0.05
Escudo Can	0.29	0.24	0.28	0.28	0.26	0.27	0.31	0.30	0.41	0.26	0.29	0.25	0.09	0.04	0.04
Cristal Other	0.32	0.27	0.29	0.26	0.33	0.28	0.34	0.34	0.40	0.26	0.27	0.24	0.10	0.05	0.04
Escudo Other	0.32	0.25	0.28	0.27	0.28	0.30	0.31	0.33	0.41	0.29	0.27	0.26	0.09	0.05	0.04
Baltica	0.43	0.31	0.36	0.31	0.34	0.31	0.51	0.46	0.51	0.34	0.36	0.34	0.11	0.07	0.05
Becker	0.43	0.32	0.34	0.30	0.34	0.33	0.46	0.46	0.45	0.32	0.33	0.32	0.13	0.06	0.06
Stella Artois	0.49	0.39	0.44	0.41	0.40	0.41	0.51	0.45	0.65	0.40	0.40	0.38	0.14	0.06	0.06
Heineken	0.33	0.26	0.30	0.26	0.26	0.29	0.34	0.32	0.40	0.34	0.28	0.27	0.12	0.05	0.05
Royal Guard	0.35	0.29	0.29	0.29	0.27	0.27	0.36	0.33	0.40	0.28	0.36	0.28	0.10	0.06	0.05
Other Brands	0.31	0.24	0.28	0.25	0.24	0.26	0.34	0.32	0.38	0.27	0.28	0.27	0.09	0.04	0.04
Temperature	0.12	0.09	0.11	0.09	0.10	0.09	0.11	0.13	0.14	0.12	0.10	0.09	0.06	0.02	0.02
State Dependence	0.06	0.05	0.05	0.04	0.05	0.05	0.07	0.06	0.06	0.05	0.06	0.04	0.02	0.01	0.01
ln(Price)	0.06	0.04	0.05	0.04	0.04	0.04	0.05	0.06	0.06	0.05	0.05	0.04	0.02	0.01	0.01

Table D.9: Discounts that offset the average stockout effect

	Av. Log(price) $\log(p_j)$ (1)	Stockout Eff. ρ_j (2)	Av. Stockout $\overline{ST}_j$ (3)	Discount [%] d_j^* (4)
Cristal (1L bottle)	6.51	-0.44	0.55	20.46
Cristal (350cc can)	5.49	-0.48	0.09	4.13
Escudo (1L bottle)	6.49	-0.19	2.01	30.49
Escudo (350cc can)	5.51	-0.06	1.16	6.39
Cristal Other	5.73	-0.04	0.97	3.62

Notes: The price discount d_j^* that offsets the stockout effect, is such that $\overline{\eta}\ln((1-d_j^*)p_j) = \overline{\eta}\ln(p_j) - \rho_j\overline{ST}_j$, where p_j is the average price over time for product j; the estimates of the stockout effects, ρ_j and the average price coefficient, $\overline{\eta} = -1.05$ are taken from Table 3.8; and $\overline{ST}_j$ is the observed average stockout treatment across consumers. The treatments were normalized by the overall average stockout treatment (6.34). We do not include "Escudo Other" products that displayed a positive treatment effect.

Table D.10: Effects of pre-treatment stockouts on the first-time purchase probability of small brands

	Pre-Treatment Period (February)		Treatment Period (7 weeks)		Post-Treatment Period (16 weeks)	
Panel A: OLS						
Pre-Treatment Stockouts	0.129	0.011	-0.151	0.019	0.144	0.179
	(0.106)	(0.161)	(0.128)	(0.163)	(0.134)	(0.149)
Constant	0.109***	0.115***	0.292***	0.336***	0.131***	0.120***
	(0.016)	(0.012)	(0.024)	(0.020)	(0.015)	(0.009)
R-squared	0.01	0.05	0.00	0.05	0.00	0.06
Panel B: Logit						
Pre-Treatment Stockouts	0.948	0.124	-0.728	0.113	1.222	2.011
	(0.690)	(1.127)	(0.649)	(0.822)	(0.968)	(1.337)
Constant	-2.021***	-1.987***	-0.876***	-0.686***	-1.815***	-1.901***
	(0.108)	(0.079)	(0.006)	(0.092)	(0.183)	(0.137)
Log-Likelihood	-584.38	-549.49	-892.91	-860.01	-520.98	-472.76
Store FE	N	Y	N	Y	N	Y
Number of Observations	1,430	1,367	1,430	1,425	1,430	1,261

Notes: The table shows the logit estimates of stockouts during the pre-treatment period on the probability of a first-time purchase in small brand products (any brand different from Cristal and Escudo). We us the stockout variable as the sum of visits with unavailable leading brands from October 2009 to January 2010 (pre-treatment). The stockout measure is consumer-specific, as described in section 3.2.2. As a normalization, we divide the stockout variable by the maximum value. Columns (2) and (4) add store fixed effects. As some stores have no variation in stockouts, we must drop some observations when including store fixed effects. Cluster-robust standard errors (at the store level) in parenthesis. P-values notation: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table D.11: Summary statistics of first-time consumers of small brands - Low treatment

	Pre-Treatment (21 weeks) (1)	Treatment (7 weeks) (2)	Post-Treatment (16 weeks) (3)
Panel A: Baltica (350cc can)			
# 1st Timers	-	4	0
# Total Consumers	29	33	33
Fraction of 1st timers	-	0.12	0.00
# of Potential 1st timers	1,319	1,290	1,286
# 1st Timers per Week	-	0.57	0.00
Panel B: Becker (350cc can)			
# 1st Timers	-	1	5
# Total Consumers	41	42	47
Fraction of 1st timers	-	0.02	0.11
# of Potential 1st timers	1,319	1,278	1,277
# 1st Timers per Week	-	0.14	0.31
Panel C: Stella Artois (354cc can)			
# 1st Timers	-	1	1
# Total Consumers	13	14	15
Fraction of 1st timers	-	0.08	0.07
# of Potential 1st timers	1,319	1,306	1,305
# 1st Timers per Week	-	0.14	0.06
Panel D: Heineken (350cc can)			
# 1st Timers	-	3	3
# Total Consumers	46	49	52
Fraction of 1st timers	-	0.06	0.06
# of Potential 1st timers	1,319	1,273	1,270
# 1st Timers per Week	-	0.43	0.19
Panel E: Royal Guard (350cc can)			
# 1st Timers	-	0	0
# Total Consumers	22	22	22
Fraction of 1st timers	-	0.00	0.00
# of Potential 1st timers	1,319	1,297	1,297
# 1st Timers per Week	-	0.00	0.00
Panel F: Other Brands/Formats			
# 1st Timers	-	5	4
# Total Consumers	122	127	131
Fraction of 1st timers	-	0.04	0.03
# of Potential 1st timers	1,319	1,197	1,192
# 1st Timers per Week	-	0.71	0.25

Notes: The table shows the placebo consumers purchasing small brand products in each period for the first time. Each Panel presents the data for different small brands, and each column shows a specific period. We define new consumers (1st timers) as those who have not purchased the corresponding SKU in our data. We label placebo consumers to those facing a sum of stockouts at the bottom five percentile of the aggregate exposure to stockouts, ie., consumers who experienced at most three episodes of leading brand unavailability across all products and visits during the treatment period. Column (1) shows the records for the 21 weeks of pre-treatment period; Column (2) for the 7 weeks of the treatment period; and Column (3) for the 16 weeks of the post-treatment period.

Table D.12: Structural model - Market share estimates conditional on buying

	Baseline	Post Treatment	% Change
	(1)	(2)	(3)
Cristal Bottle	10.56	9.37	-11.26
Cristal Can	22.34	22.94	2.68
Escudo Bottle	11.88	9.65	-18.75
Escudo Can	26.34	26.58	0.94
Cristal Other	3.90	4.06	4.05
Escudo Other	4.30	4.90	14.14
Leading Brands	79.32	77.51	-2.27
Baltica	2.19	2.37	7.99
Becker	2.65	2.85	7.53
Stella Artois	0.88	0.95	7.52
Heineken	2.47	2.65	7.21
Royal Guard	1.23	1.33	7.71
Other Brands	11.25	12.34	9.68
Small Brands	20.68	22.49	8.72

Notes: Market shares are calculated using demand estimates and average explanatory variables observed in the data. Column (1) shows the posterior mean of the market shares evaluating the estimated demand function at the average price, state dependence and temperature observed in the Pre-treatment period, imposing full availability. Column (2) shows the same calculation of Column (1) but assuming that all consumers faced the average stockout exposure observed during the Post-Treatment period. Column (3) shows the relative change in the market shares caused by the presence of stockouts.

Table D.13: Marginal changes in market shares conditional on buying due to an extra week of stockouts

		Additional Stockout for					
	Baseline	Cristal Bottle	Cristal Can	Escudo Bottle	Escudo Can	Cristal Other	Escudo Other
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Cristal Bottle	9.37	**6.78**	10.21	9.50	9.44	9.39	9.36
Cristal Can	22.94	23.70	**16.85**	23.15	23.16	22.99	22.92
Escudo Bottle	9.65	9.96	10.16	**8.40**	9.78	9.66	9.63
Escudo Can	26.59	27.10	28.30	26.97	**25.67**	26.61	26.52
Cristal Other	4.06	4.25	4.71	4.11	4.10	**3.92**	4.05
Escudo Other	4.90	5.01	5.21	5.01	5.02	4.91	**5.06**
Baltica	2.37	2.43	2.55	2.41	2.41	2.37	2.37
Becker	2.85	2.93	3.15	2.88	2.89	2.85	2.84
Stella Artois	0.95	0.97	1.05	0.96	0.97	0.95	0.95
Heineken	2.65	2.72	2.94	2.68	2.70	2.66	2.65
Royal Guard	1.33	1.36	1.47	1.34	1.35	1.33	1.32
Other Small Brands	12.34	12.79	13.38	12.58	12.52	12.36	12.32

Notes: The matrix shows the market shares for each product resulting in the demand function when using baseline parameters plus the post-treatment estimates at the average observed price and state dependence, but adding an additional week of stockout to the observed average stockout treatment. The difference between the baseline market share (in the first row) is the marginal effect in market shares of an extra week of stockout. The expected effect is a reduction in the same product market share and a weekly increasing in competitor and outside good, that includes not buying beer.