Gender-Based Pricing Ban in Health Insurance: Evidence from Chile

Valentina Reig*

Toulouse School of Economics

August 2026

Abstract

Regulation in health insurance aims to reduce inefficiencies while addressing redistribution across demographic groups. We examine a regulation within Chile's health insurance market that mandates plan pricing must not be based on gender. Leveraging administrative data from Chile, we offer empirical insights into the impact of this pricing constraint on consumer welfare. Our analysis sheds light on the discourse surrounding the prohibition of observable-based pricing, driven by considerations of equity and market efficiency. While this pricing ban may benefit some demographic groups, it can also exacerbate adverse selection by increasing prices. To analyze the implications of this regulation, we develop and estimate a discrete choice model of health insurance plans that incorporates the distribution of prospective individual medical costs. We combine the insurance plan demand model with a plan pricing model that accounts for insurer competition. Using this model, we explore counterfactual scenarios to evaluate how insurers might adjust their pricing and plan offerings under alternative regulatory policies. Our results show that after the policy, consumer surplus increased for women and decreased for men. Focusing only on the policy's price effect without accounting for changes in plan characteristics would underestimate the losses experienced by men. These findings highlight the importance of considering both price effects and insurer responses when designing health insurance regulations, providing insights for policymakers seeking to balance efficiency and redistribution.

1 Introduction

Access to quality healthcare is a primary objective for governments around the world. Health insurance systems are highly regulated to ensure affordability and efficiency. Insurance markets are characterized by selection, where high-cost individuals are more willing to pay for health insurance plans. This drives up average costs and premiums, potentially exceeding the willingness to pay of low-cost individuals and leading them to exit the market. Because of this adverse selection problem, governments may mandate health insurance. On the supply side, insurers seek to attract healthy individuals to have lower costs, or determine price based on medical conditions and individuals' characteristics. These risk-rated premiums generate another inefficiency known as reclassification

*I am grateful to Pierre Dubois for the guidance on this project. I also thank Angelique Acquatella, Matt Backus, Nano Barahona, Isis Durrmeyer, Ana Gazmuri, Laura Grigolon, Ben Handel and Paula Navarro for valuable comments and discussion. I received helpful feedback from the participants at UC Berkeley IO Workshop, Mannheim ENTER Seminar, ECWE Workshop, BECCLE and EARIE Conference.

risk which may lead governments to prohibit pricing on observable characteristics. Yet, this ban on risk-based pricing increases adverse selection which implies higher welfare costs. This trade-off between adverse selection and reclassification risk is studied by Handel, Hendel and Whinston (2015), who conclude that reclassification risk from health-based pricing is about five times more costly than the adverse selection created by restricting it. The introduction of the ACA exchanges in the US put an end to price increases due to new health conditions, as (pre-subsidy) pricing can only depend on age and smoking status in some states (Tebaldi, 2025). Einav and Finkelstein (2011) show that when the risk is directly linked to gender (under no private information, and assuming one insurer), having separate prices for men and women will be efficient, whereas pooling the two genders under the same price will generate adverse selection and welfare losses. However, when there is private information not captured by gender, the result after pooling the groups is ambiguous depending on the demand groups of the pooled and the separated groups.

Bans on pricing by gender are common in insurance markets. In the United States, the Affordable Care Act banned gender rating in the individual health insurance market from 2014. In the European Union, a 2011 ruling of the Court of Justice (the *Test-Achats* case) required the same premiums for men and women in all insurance contracts from December 2012 (Huang and Salm (2020) study its effects in Germany). In Chile, the regulator fixed a single table of risk factors in 2020, with the same factors for men and women of the same age. The table applied only to new contracts. In 2022, the Supreme Court ordered insurers to also apply the 2020 table to old contracts. Insurers can respond to these bans by changing their prices and the plans they offer, and these responses change how the gains and losses of the ban are split between women and men.

This paper provides empirical evidence on the effects of restricting insurers from pricing on observable characteristics. Using a model that combines demand for health insurance plans with insurer supply-side decisions, we analyze consumer surplus changes across demographic groups and how insurers respond to the policy. Our study focuses on a policy change affecting the pricing of plans in the health insurance private sector in Chile. The regulation prohibited health insurance prices to be based on gender, what was allowed before the second quarter of 2020. We expect the ban to have two effects. First, pooling different risks under a single price should increase adverse selection. Since women have higher medical costs than men on average, a common price makes plans more expensive for low-cost men, some of whom may leave the private sector, which pushes average costs and prices up (Einav & Finkelstein, 2011). Second, the ban has redistributive effects, and women and older individuals should benefit as their prices fall relative to those of men and younger individuals. However, insurers in Chile operate in an imperfectly competitive market, with plans that differ in financial and non-financial characteristics such as coverage rates and hospital networks. They can respond to the ban by changing not only their prices but also the plans they offer, for example to attract healthier individuals. The final effects of the ban may therefore differ from these expected effects.

In Chile, individuals choosing a health-insurance plan for the first time have to contribute at least in a 7% of their income, and can choose to join the public sector or the private sector. In the private sector, they can choose a health insurance plan from a large menu with different char-

acteristics and coverage rates, and pay a monthly premium. Insurers offer plans at a base price, adjusted by risk factors based on age and gender. Before the second quarter of 2020, the risk factor weights were also determined by the insurer. In the second quarter of 2020, the regulator fixed the risk factor weights, giving the same weight for men and women, and decreasing the weights for older people. This policy only affected the purchase of new health insurance plans, but not existing contracts. As a consequence, insurers modified the menu of plans to offer and their associated prices. This setting offers a unique context to analyze how regulation for equity may inadvertently increase adverse selection and welfare costs, or potentially mitigate adverse selection by attracting healthier enrollees (Tebaldi, 2025).

Our data come from the regulatory agency of the private health insurance system in Chile (*Superintendencia de Salud*). We present descriptive evidence on the price differences between women and men, before and after the policy. Controlling for plan fixed effects, we find that women were paying around 60% more than men for a similar plan, while coinsurance rates showed little change. We study the effects of the policy for women comparing switchers before the policy and switchers after the policy, and present an event-study where the event time is the time of switching. On average, women who switched plans after the policy had lower total expenditure, premiums plus out-of-pocket costs, than women who switched before the policy. We also present medical cost across all enrollees, finding aggregate costs are higher for women, though highly variable by age.

To evaluate alternative pricing rules, we set up a model in which the regulator chooses the rules to maximize a welfare function. The objective of the regulator is a weighted sum of consumer surplus and profits, following the literature on social welfare weights (Saez & Stantcheva, 2016). We will study three rules: i) change in factor weights for new plans, ii) change in factor weights for all plans and iii) allow for risk-adjustment transfers as in the US, based only on gender. For a fixed regulatory rule, insurers sell many plans and compete in prices à la Bertrand-Nash. They choose the coverage and the price of each plan. In the demand model, individuals choose one health insurance plan from their choice set, that can be different across periods and across individuals. We model the choice of the insurance plan based on Abaluck and Gruber (2011) and Atal (2019) and papers that use a discrete choice model (Ho, 2006; Starc, 2014; Ho & Lee, 2017; Tebaldi, 2025) to estimate taste parameters and risk aversion.Individuals have heterogeneous preferences for price and also present switching costs, as a number of studies that allow for choice inconsistencies and inertia (Abaluck & Gruber, 2011; Handel, 2013; Handel & Kolstad, 2015; Abaluck & Gruber, 2016; Fleitas, 2017; Atal, 2019).

In the supply-side model, insurers compete in the health insurance market offering plans with different financial and non-financial characteristics. Insurers choose price, and also the coverage rate of a health insurance plan with some fixed characteristics. This allows us to compare what happened after the policy when insurers could change their plans and how the market will respond when facing different policies that we evaluate in counterfactuals. Cuesta, Noton and Vatter (2024) model a plan design step to decide multiple coverage rates that one plan faces in different hospitals. They need that structure to provide counterfactuals of different hospital ownership if the vertical integration structure changes. Here, we only allow for one average coverage rate which should

capture the main characteristics of the health insurance plan. Other papers in the literature allow firms to choose more than one variable in competition games, for example, in Abi-Rafeh, Dubois, Griffith and O'Connell (2026) firms choose both prices and advertising in a dynamic equilibrium.

We find that the policy increased the consumer surplus of women by 3.5 UF per month (150 US dollars), more than their average monthly premium of 2.9 UF. It lowered the consumer surplus of men by 2.0 UF. Even if expected, these results could have been different if prices had increased a lot on average, and women could have lost too. The share of women who choose the private sector increased by 11.5 percentage points, from 62% to 73%, while the share of men fell by 2.8 percentage points. The gains for women mostly come from prices. Of the 3.5 UF gain, 2.3 UF come from lower prices and 1.2 UF from the new plans insurers offered. The losses for men come almost entirely from the new plans. Prices explain only 0.03 UF of their 2.0 UF loss. Therefore, measuring the welfare effects of the policy requires accounting for how insurers changed their plans. In 2022 the Supreme Court ordered that the policy also applies to old contracts. If insurers do not react, this extension increases consumer surplus by 1.2 UF per month on average for women and men together. But insurers would set higher base prices if they could re-price their old plans. This suggests that the extension lowers the profits of insurers at current prices.

More broadly, our results show that a ban on pricing by gender redistributes from men to women, but insurers respond by changing the plans they offer. To evaluate how the gains and losses of a pricing ban are split between groups, regulators need to take these responses into account, not only the direct effect of the ban on prices.

Our paper makes three contributions. First, we evaluate a gender pricing ban in health insurance after it was introduced, with a structural model of demand and data from before and after the policy. This complements previous work on Chile that simulates the ban before it happened (Figueroa, 2023; Pardo & Sabat, 2025), and evidence from Germany showing that women moved to private insurance after a similar ban (Huang & Salm, 2020). We add the response of insurers and the welfare effects for women and men. Second, we show that insurers respond to the ban by changing the plans they offer, and that this response explains most of the welfare loss for men. This relates to work showing that menus of contracts can undo part of the redistribution of a pricing ban (Finkelstein, Poterba, & Rothschild, 2009), and to work on product design in selection markets (Veiga & Weyl, 2016; Azevedo & Gottlieb, 2017; Marone & Sabety, 2022). Third, we develop a supply model in which insurers choose both the price and the coverage of their plans under different pricing rules. The model shows how a pricing rule changes the first-order conditions of insurers, and how selection enters both the price and the coverage decisions.

Related literature. Several papers study how restrictions on pricing by gender redistribute across groups. Finkelstein, Poterba and Rothschild (2009) show that in the UK annuity market, when insurers can adjust their menu of contracts, the redistribution from men to women can fall to half of what it would be with a single contract. Huang and Salm (2020) find that the EU unisex rule made women in Germany more likely than men to move from public to private health insurance. Geruso (2017) documents that men and women, and younger and older individuals,

value insurance differently even at the same expected costs, so a single price cannot sort all groups efficiently. In Chile, Pardo and Sabat (2025) study the equity and efficiency effects of flat premiums, and Figueroa (2023) simulates the 2020 policy with a model of plan demand. We study the Chilean policy with data from before and after it was introduced. This lets us measure how insurers changed the plans they offer, and we find that this response explains most of the loss for men.

A large literature studies the efficiency costs of restricting prices in insurance markets with selection. Uniform prices create adverse selection, while pricing on risk exposes individuals to reclassification risk (Einav & Finkelstein, 2011; Bundorf, Levin, & Mahoney, 2012; Handel, Hendel, & Whinston, 2015; Dardanoni & Li Donni, 2016; Curto, 2023). These papers mostly assume competitive insurers and a fixed set of contracts. We study a pricing restriction in a market where a few insurers set prices and change the plans they offer.

Our supply model builds on work on imperfect competition between health insurers. Starc (2014) estimates that in Medigap, where two firms hold most of the market, premiums are more than 25 percent above claims, and low price sensitivity and brand preferences explain these markups. Curto, Einav, Levin and Bhattacharya (2021) find that insurers in Medicare Advantage capture about two thirds of the surplus the program creates. Tebaldi (2025) estimates an oligopoly pricing model for the California ACA exchange and finds that the design of subsidies across age groups changes premiums for all enrollees. Polyakova and Ryan (2021) study how insurer market power affects the redistribution from ACA subsidies. In these papers insurers choose prices for a given set of plans. In our model, insurers also choose the coverage of their plans, and the pricing rule changes both decisions.

Our results also relate to work on how insurers design their plans. Theoretical work studies how firms choose product features to attract profitable consumers when consumers differ in costs and preferences (Veiga & Weyl, 2016; Azevedo & Gottlieb, 2017), and Marone and Sabety (2022) study when offering several coverage levels is efficient. There is also evidence that insurers use plan characteristics to select enrollees. Geruso, Layton and Prinz (2019) find that insurers in the ACA exchanges design drug coverage to avoid unprofitable enrollees. In Shepard (2022), covering expensive star hospitals attracts costly patients, which gives insurers an incentive to drop them from their networks, and Serna (2024) finds that insurers in Colombia exclude providers in services used by unprofitable patients. Ericson and Starc (2016) find that a regulated change in the plans available in Massachusetts moved consumers to more generous plans. We show that insurers also change their plans in response to a gender pricing rule, and that this change explains most of the loss for men.

Finally, we use the same administrative data as other structural studies of private health insurance in Chile. Atal (2019) studies lock-in under long-term contracts, and Cuesta, Noton and Vatter (2024) study how vertical integration between insurers and hospitals affects plan design and welfare. Our demand model follows Atal (2019) and Abaluck and Gruber (2011).

The rest of the paper is organized as follows. Section 2 describes the health insurance market

in Chile, and Section 3 describes the data. Section 4 shows descriptive evidence on price differences between women and men and on medical costs. Section 5 presents the model of supply and demand. Section 6 explains identification and estimation, and Section 7 shows the estimates. Section 8 presents the welfare effects of the policy and of extending it to old contracts. Section 9 concludes.

2 Institutional Framework

The health insurance system in Chile combines a public sector (*FONASA*), that offers 4 levels of plans depending on the income of the individual, and a private sector, with approximately 6 insurers that compete offering different health care plans with distinct coinsurance rates, characteristics and premiums. Approximately 17% of the population is enrolled in the private sector and the rest in the public sector. Every employee has to contribute at least 7% of their income to a health insurance plan in the public or the private sector. In the public sector it is exactly 7%, and in the private sector the premium can be higher than 7% of income to opt for better coverage and additional benefits, or lower than 7% of income and receive back money that can be spent on health services.

In the private sector, each plan has a base price determined by the insurer and the final premium paid by the consumer is the base price weighted by a risk factor, which depends on age, and also on gender before the second quarter of 2020. Before the second quarter of 2020, the risk factors were determined by each insurer and each plan could have different risk factors. After the second quarter of 2020, the regulatory agency fixed the risk factors for all plans and all insurers[1]. This "non-discrimination" policy gave the same weight to men and women, and decreased the weights for older people. It only affected the purchase of new health insurance plans, but not existing contracts. After the policy, insurers modified the menu of plans they offer and their associated prices. For more details on the private health insurance sector in Chile see Atal (2019) or Cuesta, Noton and Vatter (2024).

Before being enrolled in a plan, the individual must declare all pre-existing conditions, and the individual may be denied insurance altogether or part of coverage for some months. In the event that multiple insurers deny coverage to a person, this person will end up in the public system. However, once an individual is enrolled in a plan, there is no reclassification, in the sense that new preexisting conditions cannot allow the insurer to give less coverage, pay a higher premium, or let the individual out of the plan. Pre-existing conditions will play a role in coverage if the person wants to switch to another insurer.

Risk-rated premiums are not allowed but insurers can change the base price of a health insurance plan, and this change affects all enrollees in that plan. However, the yearly price changes are regulated and there is a maximum adjustment value (see Atal (2019) for more details). The maximum adjustment allowed varies throughout the period we observe. Every year, enrollees receive a letter

[1]The weights are 0.6 for ages 0–19, 0.9 for 20–24, 1 for 25–34, 1.3 for 35–44, 1.4 for 45–54, 2 for 55–64, and 2.4 for 65 or older.

if there will be a change in price for next year. The enrollee may accept the change or change plans.

In general, the contracts do not have deductibles, and individuals pay a coinsurance for each medical service that is around 25% of the total cost, but there is a cap to the total cost the insurer will cover which is different for every medical service. Everyone enrolled in the health insurance system, private or public, has access to guaranteed services (GES) to treat specific conditions (oncologic treatment, depression, diabetes, etc), with special generosity rates (usually 80% or more)[2]. In other cases, part of the paid premium is insurance for catastrophic illnesses that will be activated if the treatment of the illness represents very high costs, and also has a deductible.

Insurers are regulated by *Superintendencia de Salud*. They compete in the private health insurance, in which they offer a menu of plans, which is numerous, to attract consumers and charge a premium. The plans include generosity rates for different medical services and usually there is some discount in drugs for a specific pharmacy. Insurers also cover sick-leave payments. A number of insurers are vertically integrated with hospitals through holdings (Cuesta, Noton, & Vatter, 2024). In our time period there will be 6 insurers competing.

To summarize, the private health insurance contracts have the following features. i) They are long-term contracts, because individuals cannot be forced out of a plan they have if they continue paying for it, and increases in prices should be justified by increases in costs. ii) There is guaranteed renewability (an increase in premium cannot be larger than any other individual increase in premium). iii) There is a certain degree of community rating (premiums cannot be a function of individuals' observable characteristics), with the exception of age and gender before the second quarter of 2020, and only age ranges after 2020. iv) There is no guaranteed issue, because there can be reduced coverage or even no contract if the individual has preexisting conditions when choosing a plan for the first time, or if they want to switch to another insurer. v) There is no medical underwriting or reclassification risk once the individual has a plan. vi) There is an individual mandate, but no regulation on the contracts insurers can offer, no subsidies from the government to private insurers, and no general risk-adjustment transfers. vii) Only for some specific health conditions, there are some risk-adjustment transfers across insurers at some points in time. viii) There is reinsurance for "catastrophic illnesses", which are very high-cost illnesses.

3 Data and Sample Selection

This paper uses administrative data from the health insurance regulatory agency in Chile (*Superintendencia de Salud*). It consists of the universe of enrollees on the health insurance private sector that have a contract with one of the insurers. The data provides details about each insurer's health insurance plans, including features such as generosity rates for inpatient and outpatient services (rate covered by the insurer), maximum reimbursement amounts, coverage exclusions (typically bariatric surgery), fixed copayment rates, and the type of healthcare providers associated with each plan. The provider types are categorized as follows: open (all providers have same generos-

[2]There are some risk-adjustment transfers between insurers for these specific illnesses.

ity rates), preferred provider (better generosity rates within a certain network of providers), and closed providers (you can only go to designated providers to get the specified coinsurance rates). It also includes all the enrollees and their characteristics, such as income, age, employment status (working for a company or self-employed), and the contracts with prices paid and medical claims. In addition, there is information about the sick leave payments made by the insurer in cases of illness. The data covers quarterly observations spanning from 2016 to the second semester of 2021. We also have access to data from the public sector, *FONASA* (insurer from the public sector). We include only the people that are in the most expensive level of the public system, and are policyholders. We have information when they joined the system for the first time. Wedding these two databases allows us to model the outside option for the discrete choice model, so it depends on the characteristics of the service offered by the public sector.

In the whole sample there are 3,528,095 enrollees in the private sector, of whom 45% are women. Policyholders account for 42% of the total enrollees, and the rest are dependents. There are around 7000 plans per quarter being sold before 2020, and after 2020, the number of plans is reduced to 2500. There are 13 insurers in the whole sample, although we will focus only on 6 that are those selling to the general public instead of selling to specific firms. Plans are offered throughout Chile. The analysis focuses on the capital city, Santiago.

Sample selection. To evaluate the welfare effects of the non-discrimination policy, we select a subsample of individuals that satisfy some restrictions: individuals over 21 years old, that choose an individual plan (in contrast to plans offered to groups of employees of a firm). They might switch plans in our sample within the same insurer or across insurers. We only keep plans from insurers that sell plans to the general public, as opposed to insurers that only manage health insurance plans for individuals that belong to one or a few companies, resulting in six insurers. Descriptive statistics on the characteristics of the sampled individuals are shown in Table A1. In Appendix B.1 we give more details of the selected sample for the estimation of the discrete choice model.

4 Descriptive evidence

As a first step we look at how plan prices change after the policy. For this, we estimate an event study in which we compare the prices paid by women and men through time. We normalize the last period of the observed data to zero, in which the price difference between women and men for the same plan should be zero[3]. We expect the difference in prices between women and men to be positive before the policy.

The estimated equation is:

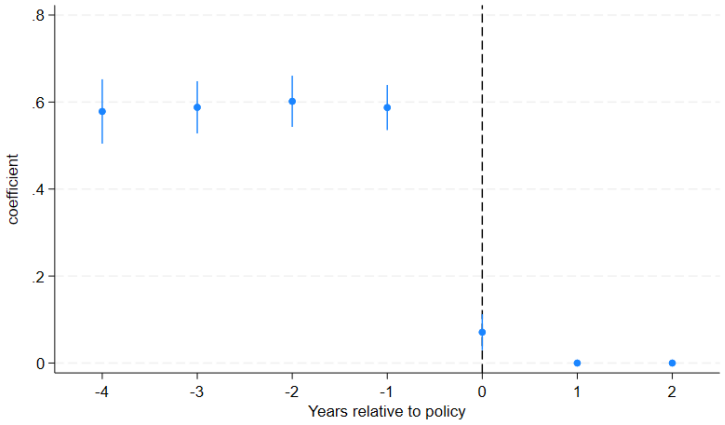
$$\log(price_{ij(s)t}) = \sum_{k=-4}^{2} \beta_k female_i * \mathbb{1}\{k = t\} + \Theta \mathbf{x_{ij(s)t}} + \delta_s + \varepsilon_{ij(s)t}$$

The price is at the plan level j but we control for aggregated plan fixed effects δ_s, which are similar

[3]We aggregate different health insurance plans under different names under the same plan for the analysis using k-means, matching on observable characteristics

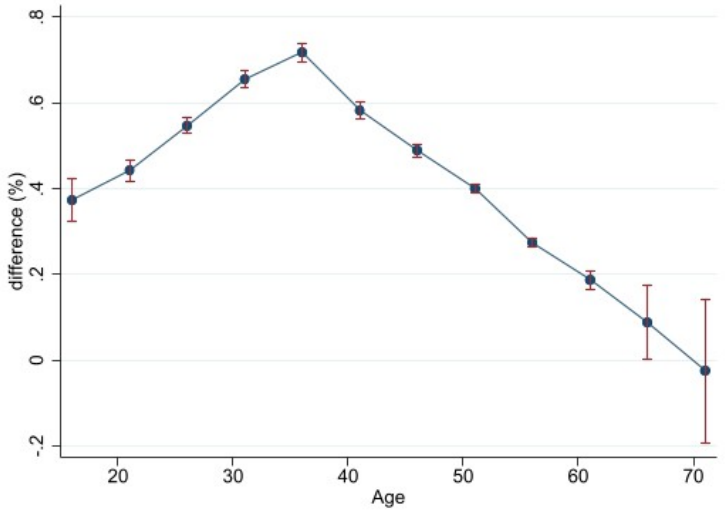
plans before and after the policy. The policy is introduced at $t = 0$, and the treatment is being a woman, to compare with prices faced by men. Control variables $\mathbf{x_{ij(s)t}}$ include characteristics of the individual: income, age and characteristics of the plan: inpatient and outpatient generosity rates, maximum reimbursement amount, provider type: includes preferred hospital or not, and the number of restrictions on medical use in the plan (for example, childbirth, bariatric surgery or psychology). Standard errors are clustered at aggregated plan level.

Figure 1: Event study of difference in prices between men and women



Notes: Coefficients β_k of the regression of log prices on being a woman interacted with each period. The policy is introduced at $t = 0$. The regression includes aggregated plan fixed effects and controls for income, age and plan characteristics. Standard errors are clustered at the aggregated plan level.

Figure 2: Differences in prices between men and women by age

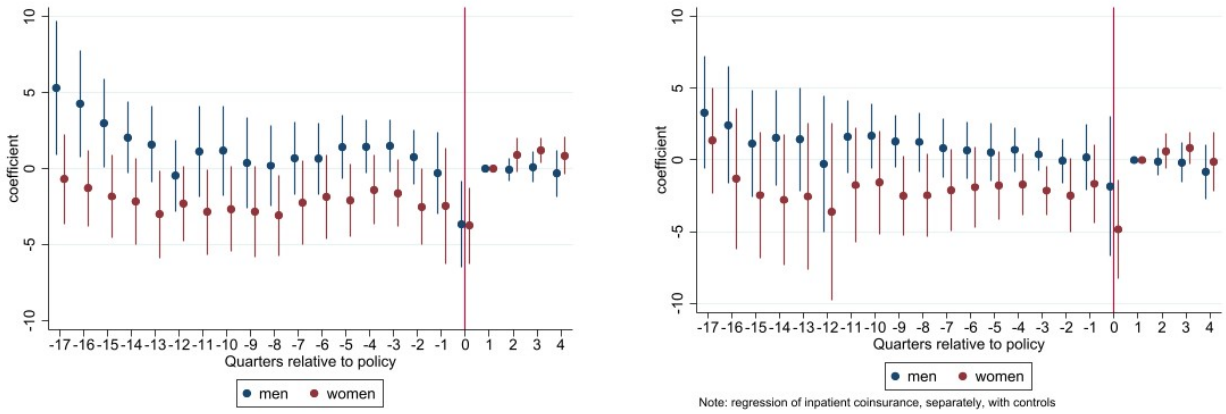


Notes: The graph presents for each age the estimated coefficient of female of different regressions with plan fixed effects. The regression for each age group is $\log(price_{ijt}) = \delta_j + \tau_t + \beta female_i + \Theta \mathbf{x_{ijt}} + \varepsilon_{ijt}$.

Figure 1 shows women are paying around 60% more than men for similar plans before the policy.

It is around 50% if we include insurer and age fixed effects instead of aggregated plan fixed effects (Figure 18 in the Appendix). We also show how the effect looks for women in maternity ages (Figure 4a) and outside maternity ages (Figure 4b). Although the difference is larger for women who could go through maternity, there is a positive difference for ages outside the maternity years. In Figure 2 we see the coefficients related to the price difference between women and men for different ages, highlighting again that the difference is higher for maternity years. Figures 3a and 3b show regressions for outpatient and inpatient coverage (in %) separately for men and women controlling for age, income, provider type, and out of pocket medical costs (OOP) cap[4]. We observe the mean coinsurance rate is lower for women before the policy, but it is not statistically different from 0. This suggests that the price differences are not explained by better coinsurance rates for women.

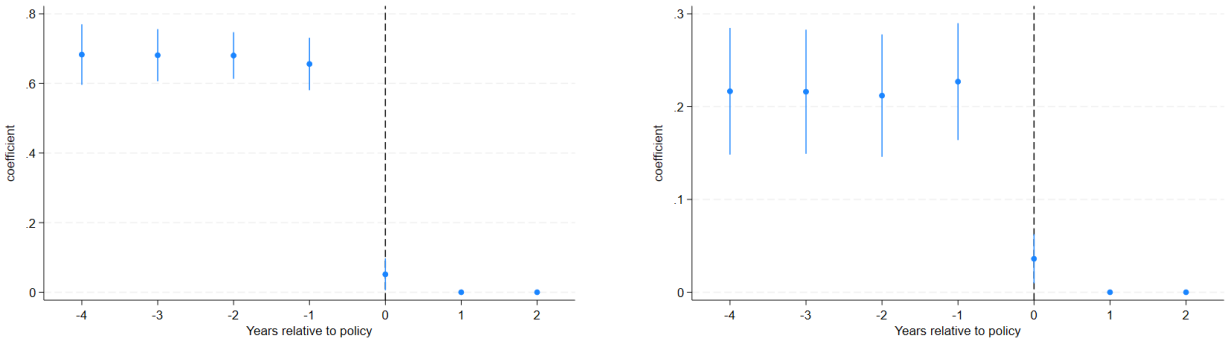
Figure 3: Event study of coinsurance rates by gender



(a) Outpatient (b) Inpatient

Notes: Coefficients β_k of the regression of the outpatient (panel a) and inpatient (panel b) coinsurance rate on period indicators, estimated separately for women and men. Controls are age, income, provider type and the cap on out-of-pocket costs. Standard errors are clustered at the insurer level.

Figure 4: Event study by maternity age group



(a) During Maternity Age (21–50) (b) Outside Maternity Age (> 50)

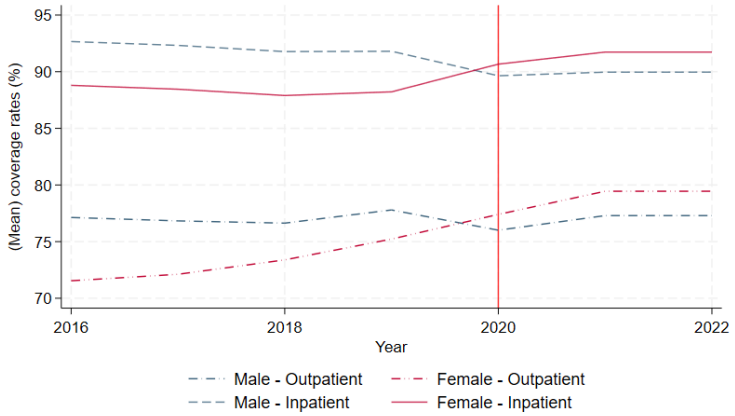
Notes: Same regression as in Figure 1, for individuals aged 21 to 50 (panel a) and over 50 (panel b).

Figure 5 shows the mean coinsurance rates of the health insurance plans purchased over time. On average, women purchased plans with lower coinsurance rates for both inpatient and outpa-

[4]The equation is $coins_rate = \delta_s + \tau_t + \sum_{k=-17}^{4} \beta_k \mathbb{1}\{k = t\} + \Theta \mathbf{x2_{ist}} + \varepsilon_{ist}$, clustered at insurer level

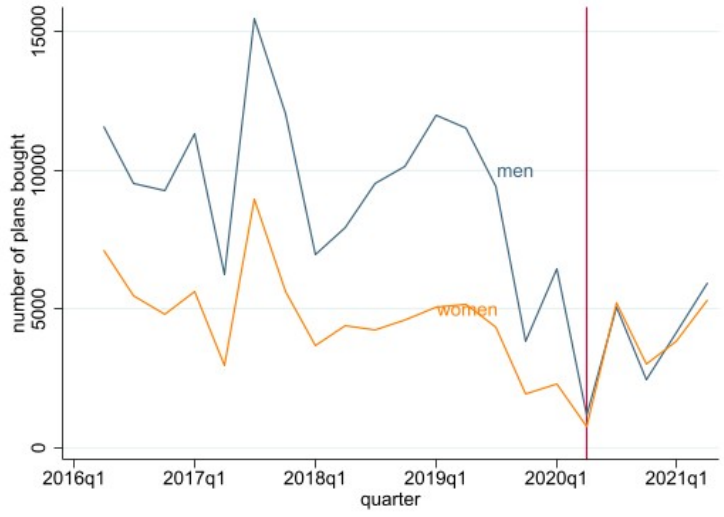
tient, but these rates reach the men's coinsurance rates after the policy, which decrease a little. The behavior is similar for the number of plans purchased by men and women in the private sector, as shown in Figure 6. Men were buying more plans than women before the policy, but the numbers are similar after the policy. Figure 7 shows thedistribution of prices paid and the 7% of income, which is the mandatory amount that should be paid for health insurance.

Figure 5: Mean coinsurance rates for offered plans



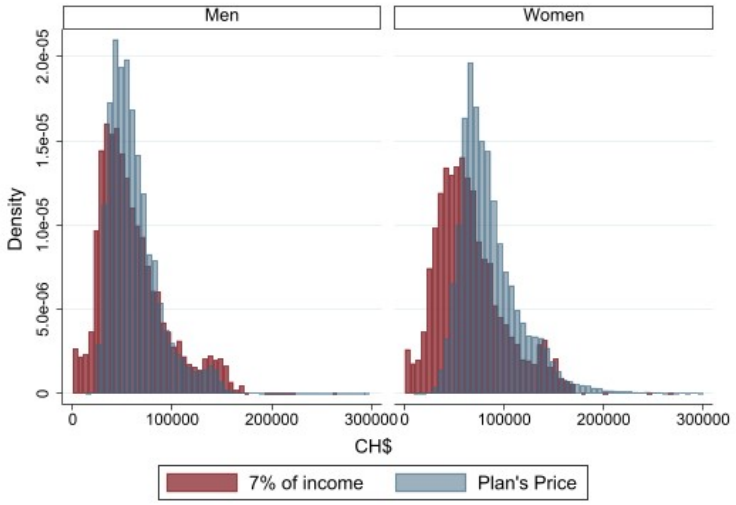
Notes: Mean inpatient and outpatient coinsurance rates of the plans purchased, by gender and period.

Figure 6: Number of new plans bought



Notes: Number of new plans bought by women and men in the private sector, by period.

Figure 7: Distribution of Premiums paid

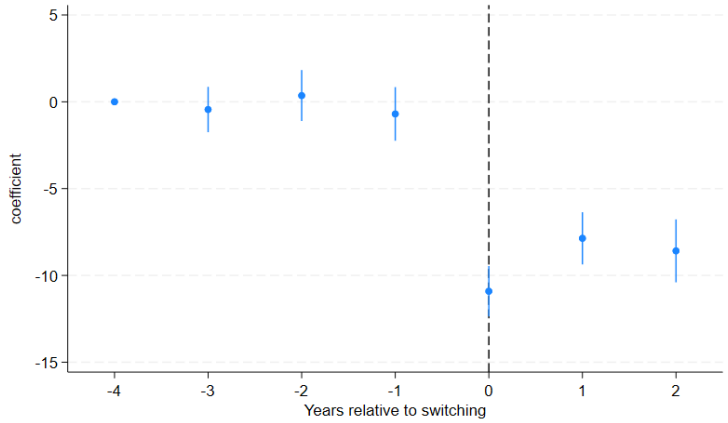


Notes: Distribution of the premiums paid. The graph also shows 7% of income, the mandatory contribution to health insurance.

4.1 Switchers analysis

We provide evidence on the effect of the policy in total expenditure for individuals that are switching, accounting for yearly premiums plus out-of-pocket costs. We present an event study comparing switchers after the policy with switchers before the policy. We find that the total yearly payments for women who switched after the policy are 10 UF lower than those who switched before the policy, in which prices were higher for women. We control for individual fixed effects to capture health status. The policy started in the COVID year, in which people went less to the doctor, and had less out-of-pocket expenditures. We partially control for this with year fixed effects.

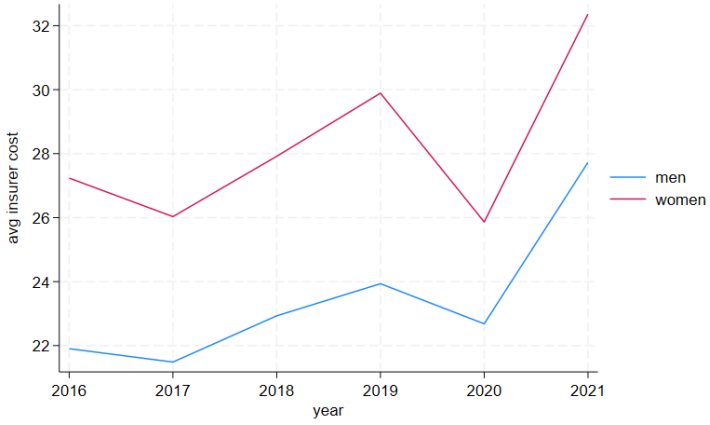
Figure 8: Total individual's expenditure for switchers compared to switchers before policy



Notes: Event study for total expenditure (out-of-pocket costs plus yearly price) at the time individual switched, controlling for plan characteristics, individual fixed effects, insurer fixed effects and year fixed effects. Standard errors are clustered at the individual level.

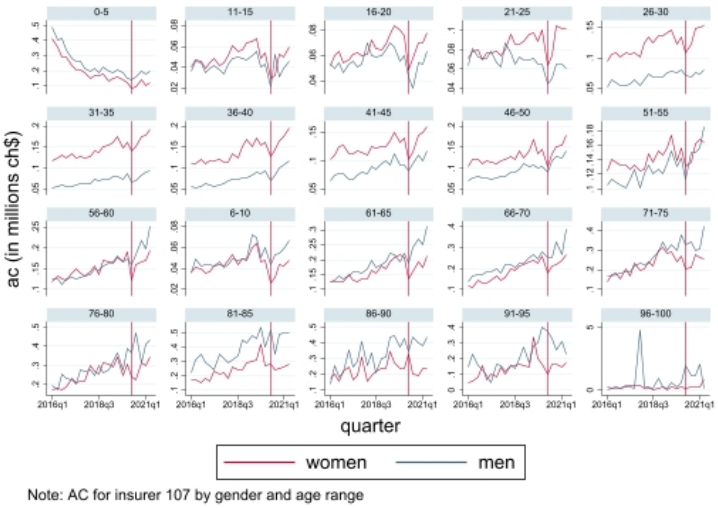
Are gender price differences justified by the expected costs in healthcare? We show some descriptive statistics of total amount spent by insurers for each gender, and across age. Figure 9 shows total insurer's cost per capita for the whole industry, revealing higher costs for women across all periods. While these elevated costs might suggest that women should pay more for their health insurance plans, this graph masks large differences in costs among insurers and across age ranges. In addition, these aggregate statistics do not control for differences in plans' characteristics. Figure 10 provides a more nuanced analysis focusing on one particular insurer and disaggregating the cost heterogeneity by age ranges. The graph shows average costs (AC) by age ranges for men and women. Average costs for women are notably higher than for men between the ages of 26 and 50, while the trend reverses after the age of 61. These findings highlight the need for a better understanding of gender-based pricing in healthcare that considers the underlying complexities and characteristics of the plan.

Figure 9: Average insurer cost by gender



Notes: Average insurer cost per enrollee by year and gender, for the whole industry.

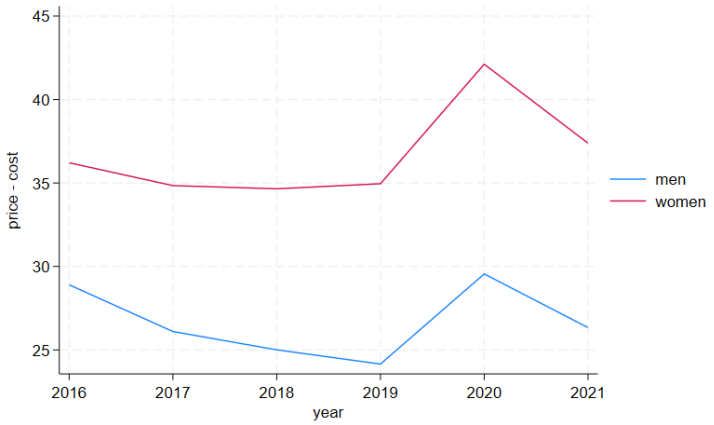
Figure 10: AC by age-range



Notes: Average costs of one insurer by age range, for women and men.

In addition to the previous analysis, we present the insurers' markup (premium-costs). The markup is calculated by individual, as yearly premium minus yearly average cost incurred by the insurer for each person. Figure 11 shows that markups in levels are higher for women than for men, suggesting insurers profit more from women than from men before the policy.

Figure 11: Markups



Notes: Markup computed as yearly price minus cost paid by insurer by gender, age and plan. Average taken over all insurers and plans for each year.

5 Model

We build a model of the private health insurance market in Chile to study how insurers respond to pricing rules, and how these responses affect consumers. The model has two sets of players. Insurers choose the coverage and the prices of their plans. Individuals choose one plan from the private sector, or the public option. The pricing rule set by the regulator is taken as given. We use the model to compare the rule before the policy, the rule after the policy, and alternative rules the regulator could adopt.

The supply side has 6 insurers that compete to attract individuals. Each insurer sells many plans. Plans differ in coverage, provider network, maximum reimbursement amounts and restrictions on medical use. In the model, insurers choose the coverage rate and the price of each plan. The other characteristics are fixed. We take the 30 aggregated plans observed before the policy, which have different combinations of hospitals, maximum reimbursement amounts and restrictions on medical use.

The demand side is a discrete choice model. Individuals choose a health insurance plan from the 6 insurers, or the public option. Their indirect utility depends on the characteristics of the plans and on their demographics. It also depends on the expected out-of-pocket costs they will have in a year, and on the switching costs they face if they move to another plan. Switching costs capture inattention to other plans, and the cost of studying many plans and going through the administrative process.

5.1 Supply side

We model the supply side as a two-stage game. In the first stage, each insurer f chooses the coverage rate c_j of each of its plans $j \in J_f$. In the second stage, insurers observe the coverage of all plans and set prices. They compete in prices à la Bertrand-Nash. We solve the game backwards, and the equilibrium is subgame perfect.[5]

Prices and the pricing rule. Each plan has a base price P_j set by the insurer. The price paid by an individual in group g, defined by age and gender, is the base price times a risk factor. Before the policy, insurers chose the risk factor of each group for each plan, so the price is $P_{gj} = \phi_{jg} P_j$. After the policy, the regulator fixes the risk factors, and they are the same for men and women of the same age. The price is then $P_{gj} = \phi_g P_j$. Choosing P_j and ϕ_{jg} freely is the same as choosing a separate price P_{gj} for each group. So before the policy the insurer sets one price per group, and after the policy it only sets the base price.

Costs. Let $cost_i$ be the expected yearly medical cost of individual i. The insurer pays the share c_j of this cost. The cost of individual i in plan j is

$$\tilde{c}_{ij} = c_j \, cost_i. \tag{1}$$

We assume that the total medical cost of an individual does not depend on the plan. This is the same no moral hazard assumption we use in the demand side to build out-of-pocket costs (Section 6). We also assume insurers have rational expectations. They know the demand for their plans and the expected cost of each individual. Prices in the demand model are monthly premiums, so we convert yearly medical costs to monthly costs.

Profits. The profit of insurer f is

$$\pi_f(\mathbf{P}, \mathbf{c}) = \sum_{j \in J_f} \sum_g \sum_{i \in g} \left(P_{gj} - \tilde{c}_{ij}\right) s_{ij}(\mathbf{P}, \mathbf{c}), \tag{2}$$

where s_{ij} is the probability that individual i chooses plan j, given by the demand model in Section 5.2. This probability depends on the prices $\mathbf{P}$ and the coverage rates $\mathbf{c}$ of all plans in the market. Profits add up, over plans and individuals, the price minus the cost, weighted by the probability of choosing the plan. As in the demand model, a market is a period t, and we drop the index t to simplify notation.

5.1.1 Second stage: Pricing decision

Take the coverage rates $\mathbf{c}$ as given. Before the policy, insurer f chooses the price P_{gj} of each plan $j \in J_f$ for each group g. The first-order condition is

$$\sum_{i \in g} \left[s_{ij} + \sum_{k \in J_f} \left(P_{gk} - \tilde{c}_{ik}\right) \frac{\partial s_{ik}}{\partial P_{gj}} \right] = 0 \qquad \text{for all } g \text{ and } j \in J_f. \tag{3}$$

[5] We do not allow plans to exit. As long as a plan has one enrollee, the plan exists.

The first term is the gain from charging a higher price to the individuals who stay in plan j. The second term is the profit lost from the individuals who leave plan j, net of the profit gained when some of them move to other plans of the same insurer.

After the policy, the risk factors ϕ_g are fixed and the insurer only chooses the base price P_j. Since $\partial P_{gj}/\partial P_j = \phi_g$, the first-order condition is

$$\sum_g \phi_g \sum_{i \in g} \Big[s_{ij} + \sum_{k \in J_f} \big(\phi_g P_k - \tilde{c}_{ik}\big) \frac{\partial s_{ik}}{\partial P_{gj}} \Big] = 0 \qquad \text{for all } j \in J_f. \tag{4}$$

Equations (3) and (4) show how the policy works. Before the policy, there is one condition for each group. The price for women depends only on the demand and the costs of women. After the policy, there is one condition for each plan. It adds up the conditions of all groups, weighted by ϕ_g. So the base price balances the demand and the costs of men and women. If women have higher costs, the insurer can no longer charge them a higher price. Men of the same age pay the same price, so they pay part of the higher costs of women.

Demand derivatives. From the logit model in Section 5.2, for an individual i in group g,

$$\frac{\partial s_{ij}}{\partial P_{gj}} = \alpha_i \, s_{ij}(1 - s_{ij}), \qquad \frac{\partial s_{ik}}{\partial P_{gj}} = -\alpha_i \, s_{ij} \, s_{ik} \quad \text{for } k \neq j, \tag{5}$$

where $\alpha_i = \alpha_0 + \alpha_1 age_i + \alpha_2 income_i$ is the price coefficient of individual i. Individuals in other groups do not pay P_{gj}, so their choices do not respond to it.

Selection. Health insurance is a selection market. The cost of a plan depends on who buys it. The first-order conditions above already take this into account, because they are written for each individual. When the price of plan j goes up, the individuals who leave are not random. How much each individual responds depends on the price coefficient α_i and on the probability s_{ij}, which depends on the expected costs of the individual. To see this, take an insurer with a single plan j and a single price P_j. Let $s_j = \sum_i s_{ij}$ be the enrollment of the plan and $AC_j = \sum_i s_{ij} \, \tilde{c}_{ij}/s_j$ its average cost. Adding up the first-order condition over individuals gives

$$P_j - AC_j = \frac{s_j}{-\partial s_j / \partial P_j} \Big(1 - \frac{\partial AC_j}{\partial P_j}\Big). \tag{6}$$

We show the derivation in Appendix C. The first part is the usual markup of a differentiated products model. The second part is the selection term of Mahoney and Weyl (2017). With adverse selection, the individuals who leave when the price goes up are the ones with low costs. The average cost goes up with the price ($\partial AC_j/\partial P_j > 0$) and the markup is lower. With advantageous selection ($\partial AC_j/\partial P_j < 0$) the markup is higher. In our model we do not need to estimate the selection term separately. It comes from the demand model and the individual costs.

5.1.2 First stage: Coverage decision

In the first stage, insurer f chooses the coverage rate c_j of each plan $j \in J_f$. It anticipates that the prices of the second stage depend on coverage, $\mathbf{P}^*(\mathbf{c})$. The insurer maximizes $\pi_f(\mathbf{P}^*(\mathbf{c}), \mathbf{c})$. The first-order condition for c_j is

$$\frac{d\pi_f}{dc_j} = \underbrace{\frac{\partial \pi_f}{\partial c_j}}_{\text{direct effect}} + \underbrace{\sum_{k \in J_{-f}} \sum_g \frac{\partial \pi_f}{\partial P_{gk}} \frac{\partial P^*_{gk}}{\partial c_j}}_{\text{effect through rival prices}} = 0. \tag{7}$$

J_{-f} is the set of plans of rival insurers. A change in coverage also changes the insurer's own prices. These terms are zero by the envelope theorem, because own prices already satisfy the first-order conditions (3) or (4). After the policy, the effect of each own group price is not zero on its own, but their sum weighted by ϕ_g is zero by (4). The effect through rival prices does not vanish. When insurer f raises coverage, rivals adjust their prices, and this changes the demand for the plans of insurer f. Equation (7) holds before and after the policy, because after the policy $P^*_{gk} = \phi_g P^*_k$.

The direct effect is

$$\frac{\partial \pi_f}{\partial c_j} = \sum_g \sum_{i \in g} \Big[- cost_i \, s_{ij} + \sum_{k \in J_f} \big(P_{gk} - \tilde{c}_{ik}\big) \frac{\partial s_{ik}}{\partial c_j} \Big]. \tag{8}$$

The first term is the cost effect. With higher coverage, the insurer pays a larger share of the medical costs of its enrollees. The second term is the demand effect. Higher coverage makes plan j more attractive, so more individuals choose it, and some of them come from other plans of the same insurer. From the logit model,

$$\frac{\partial s_{ij}}{\partial c_j} = s_{ij}(1 - s_{ij}) \frac{\partial V_{ij}}{\partial c_j}, \qquad \frac{\partial s_{ik}}{\partial c_j} = -s_{ij} \, s_{ik} \frac{\partial V_{ij}}{\partial c_j} \quad \text{for } k \neq j, \tag{9}$$

with

$$\frac{\partial V_{ij}}{\partial c_j} = \beta^1 + \omega_i \frac{\partial EOP_{ij}}{\partial c_j} + \omega_3 \frac{\partial VOP_{ij}}{\partial c_j}. \tag{10}$$

Coverage enters utility directly through β^1, the coefficient on the coverage rate in Section 6. It also lowers the expected out-of-pocket costs EOP_{ij} and their variance VOP_{ij}. The fall in out-of-pocket costs is larger for individuals with high expected costs. So higher coverage attracts high-cost individuals the most. This is selection on coverage, and it limits the incentive of insurers to offer generous plans. The term $\partial P^*_{gk}/\partial c_j$ comes from the implicit function theorem applied to the price first-order conditions of all insurers. We give the details in Appendix D.

Equilibrium and assumptions. An equilibrium is a set of coverage rates and prices such that, given coverage, prices satisfy (3) before the policy or (4) after the policy, and coverage satisfies (7) for every insurer. The model rests on four assumptions. First, it is static. Insurers do not take into account that enrollees stay in their plans for many years because of switching costs. Second, insurers have rational expectations about demand and costs. Third, there is no moral hazard, so total medical costs do not depend on the plan. Fourth, besides price, coverage is the only characteristic insurers choose, for a given set of plans.

Because we observe the claims of each enrollee, we measure $cost_i$ directly. So the first-order conditions can be used in two ways. They test whether observed prices are consistent with Bertrand-Nash competition, and they recover the part of costs we do not observe, such as administrative costs. In Section 8 we use a simple version of the pricing condition to study how insurers would re-price their old plans. The demand model below gives the choice probabilities s_{ij} and their derivatives that enter the supply side.

5.2 Demand-side Model

Individuals have to choose a health insurance plan when they start to work or when they cannot depend anymore on their parents' plan. We model the choice of a health insurance plan in the private health insurance sector. Individuals have to choose one plan from six insurers, who offer a large menu of plans with different coverage rates, maximum reimbursement amounts, and other benefits. As in Cardon and Hendel (2001), we assume individuals choose a health insurance plan in a first stage, and in a second stage they decide how much to spend on medical services (out-of-pocket costs), which will depend on the chosen plan and the health status. Therefore, when deciding what plan to choose, the individual has already decided how much they will spend in the second stage, or at least knows what the expected costs are. To capture this, we follow Abaluck and Gruber (2011) to derive a linear indirect utility function for each individual from a constant absolute risk aversion (CARA) utility function over expenditures[6]. The indirect utility is a function of the distribution of future expenses, captured by the expectation and the variance of out-of-pocket (OOP) costs.

Each individual i in group g (defined by age and gender) has to choose a health insurance plan j from an insurer, at each period t. Each market is a period t. The indirect utility depends on price, observed and unobserved characteristics of the plan, and the average and the variance of expected out-of-pocket costs (oop). We allow individuals to switch health insurance plans once they have chosen one. Given the structure of the Chilean market, we specify some switching costs that the individual might face when choosing other plans. The switching costs capture inertia that individuals might have after choosing one plan, or the actual administrative costs of looking for new plans. We specify the price coefficient and the expected oop as a function of observed demographics $z_i = (age_i, income_i)$. This is:

$$u_{ijt} = \alpha_i P_{gjt} + \beta x_j + \xi_{jt} + \omega_i EOP_{ijt} + \omega_3 VOP_{ijt} + \gamma_1 \mathbf{1}_{ij',t-1} + \gamma_2 \mathbf{1}_{ij',t-1} \mathbf{1}\{j \in post\} + \epsilon_{ijt} \quad (11)$$

$$u_{ijt} = \underbrace{\delta_{gjt} + \mu_{ijt}(z_i)}_{V_{ijt}} + \epsilon_{ijt}$$

where P_{gjt} is the price of plan j for group g, $\alpha_i = \alpha_0 + \alpha_1 age_i + \alpha_2 income_i$ is the price coefficient and $\omega_i = \omega_0 + \omega_1 age_i + \omega_2 income_i$ is the coefficient on expected out-of-pocket costs. δ_{gjt} is the part of

[6]The CARA utility function is $U(C) = -\exp(-\rho(w - C))$ and costs C are normally distributed. It is possible to take the Taylor expansion of the expected value of this utility to get to an indirect utility.

utility common to all individuals in group g, and $\mu_{ijt}(z_i)$ is the part that varies with demographics. We call V_{ijt} the deterministic part of utility.

The plan characteristics x_j are the outpatient and inpatient coverage rates and the providers. $\mathbf{1}_{ij',t-1}$ is one if individual i chose a different plan j' in period $t-1$, so choosing j means switching plans. The interaction $\mathbf{1}_{ij',t-1}\,\mathbf{1}\{j \in post\}$ captures the extra utility of switching to a plan sold after the policy. It reflects that individuals may pay more attention to the new plans because of the policy. The error ϵ_{ijt} follows a type 1 extreme value distribution. The choice probabilities are

$$s_{ijt} = \frac{\exp(V_{ijt}(z_i))}{\sum_{k\in t}\exp(V_{ikt}(z_i))}, \tag{12}$$

and the total enrollment of plan j is

$$s_{jt} = \int \frac{\exp(V_{ijt}(z_i))}{\sum_{k\in t}\exp(V_{ikt}(z_i))} dF(z_i). \tag{13}$$

Choice sets. There is a large literature studying choice inconsistencies when choosing a health insurance contract. Many document inertia or study its reasons (Handel, 2013; Handel & Kolstad, 2015; Heiss, McFadden, Winter, Wuppermann, & Zhou, 2021). Individuals may not know how to make the best decision given these complex financial contracts (Abaluck & Gruber, 2011), or may not study all the plans being offered in a market (Lucarelli, Prince, & Simon, 2008; Abaluck & Adams-Prassl, 2021; Coughlin, 2023). Atal (2019) defines a choice set for different individuals that already have a plan, which depends on individuals' characteristics and health status. Atal (2019) assumes individuals that choose for the first time face the whole choice set. In this paper, we define the choice set of the individual as all the plans that were chosen in a year by someone in their group, defined by age and gender. Although we do not model choice inconsistencies as Heiss, McFadden, Winter, Wuppermann and Zhou (2021) do, we construct the choice set for each individual depending on their age, gender and income[7].

6 Identification and Estimation

For the identification of the parameters related to characteristics of the products, there are multiple plans each period with different financial characteristics: outpatient and inpatient coverage rates. Also, different choice sets for individuals help with identification (this is similar to variation of plans offered by market). We can identify the price coefficient from variation of prices through time for the same plan and price variation across plans. The prices across groups are also different for the same plan. Some unobserved characteristics of plans are correlated with price. To address this endogeneity problem we use instrumental variables for price, which we describe in the second step below.

We use a two step method to estimate the coefficients of the model. In the first step we recover δ_{gjt} and the coefficients associated with individual heterogeneity using Maximum Likelihood. In

[7] A number of methodological papers study how to overcome the problem of unobserved choice sets (Abaluck & Compiani, 2020; Barseghyan, Coughlin, Molinari, & Teitelbaum, 2021; Crawford, Griffith, & Iaria, 2021)

the second step, we estimate the rest of the parameters using a linear IV method to tackle the endogeneity of prices due to unobserved characteristics of the plan.

$$
\begin{aligned}
u_{ijt} &= (\alpha_0 + \alpha_1 age_i + \alpha_2 income_i)P_{gjt} + \beta x_j + (\omega_0 + \omega_1 age_i + \omega_2 income_i)EOP_{ijt} + \omega_3 VOP_{ijt} \\
&\quad + \gamma_1 \mathbf{1}_{ij',t-1} + \gamma_2 \mathbf{1}_{ij',t-1}\mathbf{1}\{j \in post\} + \xi_{jt} + \epsilon_{ijt} \\
u_{ijt} &= \delta_{gjt} + \mu_{ijt} + \epsilon_{ijt}, \qquad \text{with} \\
\delta_{gjt} &= \beta x_j + \xi_{jt} + \alpha_0 P_{gjt}, \\
\mu_{ijt} &= (\alpha_1 age_i + \alpha_2 income_i)P_{gjt} + (\omega_0 + \omega_1 age_i + \omega_2 income_i)EOP_{ijt} + \omega_3 VOP_{ijt} + \gamma_1 \mathbf{1}_{ij',t-1} + \gamma_2 \mathbf{1}_{ij',t-1}\mathbf{1}\{j \in
\end{aligned}
$$

$$
\beta x_j = \beta^1 coverage_j + \beta^2 preferedprov_j + \sum_{k=1}^{K} \beta_k^3 \mathbf{1}\{provider_k \in j\} + Ins_j
$$

$coverage_j$ corresponds to the average rate of reimbursement from the insurer in plan j. It is the coverage rate c_j that insurers choose in the supply model (Section 5.1). $preferedprov_j$ is one if plan j has preferred hospitals.

In the **first step** of the estimation we recover α_1, α_2, $\omega_0, \ldots, \omega_3$, γ_1, γ_2 and δ_{gjt} maximizing the following likelihood:

$$
L(\alpha, \delta_{gjt}) = \sum_t \sum_i \sum_j \mathbf{1}_{ij,t} \log s_{ijt},
$$

where $\mathbf{1}_{ij,t}$ is one if individual i chooses plan j in period t.

To recover δ_{gjt} we match the observed market shares of the products by market with the theoretical market shares given by s_{jt} and for a fixed α we compute δ_{gjt} through a contraction mapping function. This is done separately by age group and market (year).

In the **second step**, we estimate the following equation by two-stage least squares:

$$
\delta_{gjt} = \beta x_j + \alpha_0 P_{gjt} + \xi_{jt}. \tag{14}
$$

Price is endogenous because insurers set prices knowing the unobserved characteristics of their plans, ξ_{jt}. We use three types of instruments for price. Table A5 shows the first stage.

First, $n_{r,g,j}^{(2017)}$ is the number of people in group g who live outside the region of study (Santiago) and whose income allows them to buy plan j, in 2017. Here a group is defined by age, gender and income. We take this number from the 2017 CASEN survey, so it is not related to demand shocks in the years we study. Second, we interact this number with the national increase in medical costs, measured by the yearly change in the Health CPI from the Chilean Central Bank:

$$
Z_{gjt} \;=\; n_{r,g,j}^{(2017)} \, \frac{\Delta \text{MedicalCPI}_{t+1}}{\text{MedicalCPI}_t}. \tag{15}
$$

When national medical costs go up, the costs of insurers go up, and so do their prices. The national increase is not related to the unobserved characteristics of a given plan. The interaction with $n_{r,g,j}^{(2017)}$ gives variation across groups and plans, not only across years. We also multiply Z_{gjt} by the coverage rate of the plan, because a plan that pays a larger share of medical costs faces a larger increase in costs. Third, we use differentiation instruments (Gandhi & Houde, 2019): the number of rival plans, and the number of rival plans that include each hospital h.[8]

Construction of variables. We have information on ages by 5-year bins for all enrollees. Income is available for some individuals, both for the private sector and the public sector. In the public sector, we impute the missing income as the lower bound. To define the expected out-of-pocket (OOP) expenditures we build on Abaluck and Gruber (2016) and Atal (2019), who assume rational expectations over the expected out-of-pocket costs that the individual will have in the following year from the start of the contract. We have information of the sum of out-of-pocket costs in the following four quarters after the contract started, but only for the chosen plan. Therefore, to determine the out-of-pocket costs that the individual would have had in other plans, we assume the claims (medical services) would have been the same (like Abaluck and Gruber (2016)) but also that the total costs are the same, because we cannot compute what the costs would be in other plans[9]. However, we calculate the out-of-pocket as the total cost the person had by the coinsurance rate (separating on inpatient and outpatient rate) of the alternative plan. To compute the expected out of pocket costs, we regress the out of pocket costs in a year for each enrollee, on many variables that might define the out of pocket costs: previous costs, chronic conditions and demographics. The predicted value from the regression represents the expected out of pocket for each enrollee. With this prediction we can also compute what the expected out-of-pocket costs would be if they had chosen other plans. We approximate the conditional variance as the square of the residuals of the regression.

Outside option specification. We rely on data from the public sector of the health insurance system to construct the outside option. Individuals that are enrolled in the private sector can also choose the outside option, which is to be enrolled in Fonasa but only in the highest level of income, called level D. The premium paid in the public sector is 7% of the income. We assume the inpatient and outpatient coverage rates are 60%. In practice these rates vary, but this value captures that the quality is very different from the private sector. We also include individuals that are enrolled in Fonasa as policyholders, at level D. We have information on income, age, and region.

7 Estimation Results

The estimated parameters of the demand model are shown in Table A4. The parameters are estimated separately for men and women, before the policy. The table shows the estimates for women. The first stage of the price instrument is in Table A5. We find women dislike high prices. The price coefficient decreases with age ($\alpha_1 < 0$), so older women are more sensitive to prices. It

[8] Cuesta, Noton and Vatter (2024) address endogeneity by instrumenting each insurer's negotiated hospital prices with off-sample "closed" insurer rates, and plan premiums with the rivals' average "actuarially fair" premiums.

[9] We are assuming no moral hazard because they spend the same in any plan.

increases with income ($\alpha_2 > 0$), so women with higher income are less sensitive to prices.

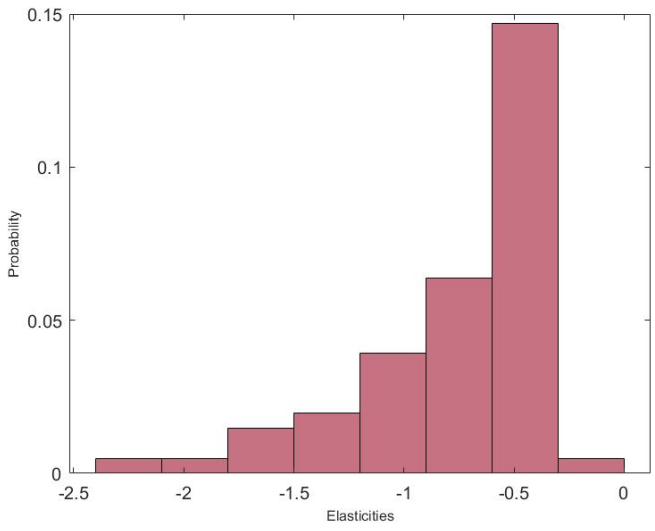
Willingness to pay. The willingness to pay of individual i for a higher coverage rate is $wtp_i = -\beta^1/\alpha_i$. It is the money value of one unit of coverage. Figure 14 shows the willingness to pay for an increase of 10 percentage points in the coverage rate, which is $0.1\, wtp_i$, in UF per month. The distribution is over women before the policy. On average, women would pay 5 UF each month for a coverage rate 10 percentage points higher, which is large compared to their mean premium of 2.9 UF.

Elasticities. We compute the own-price elasticity of each plan as the average of the individual elasticities over individuals:

$$\eta_{jt} = \int \alpha_i P_{gjt}(1 - s_{ijt})\, dF(z_i).$$

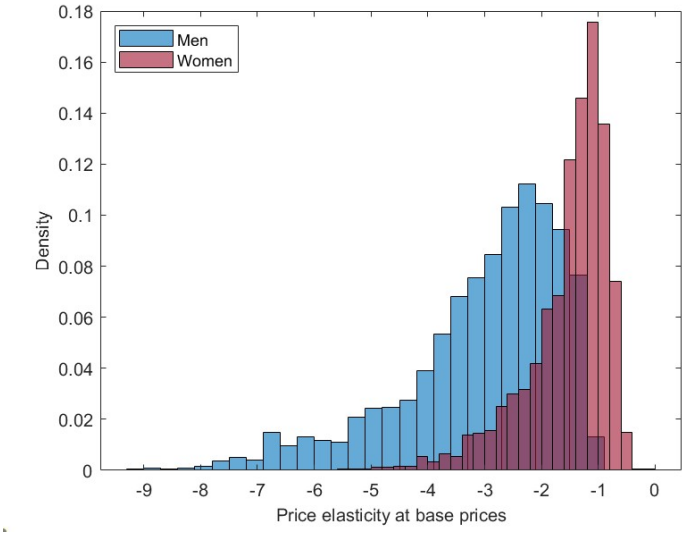
Figure 12 shows the distribution of own-price elasticities across plans for women before the policy. Before the policy, the average own-price elasticity is -1.6 for women and -2.2 for men. These elasticities are at the lower end of those found in studies of plan choice in the US. Decarolis, Polyakova and Ryan (2020), who study subsidy design in Medicare Part D, find average elasticities between -12 and -5.9 depending on the risk group, while Starc and Town (2019) find elasticities between -6.3 and -5. Tebaldi (2025) computes the price elasticity in the category of Silver plans, when all Silver plans raise their premiums by 1%, and finds that young individuals have an elasticity around -2.

Figure 12: Own-price elasticities: distribution over plans (before the policy), for women



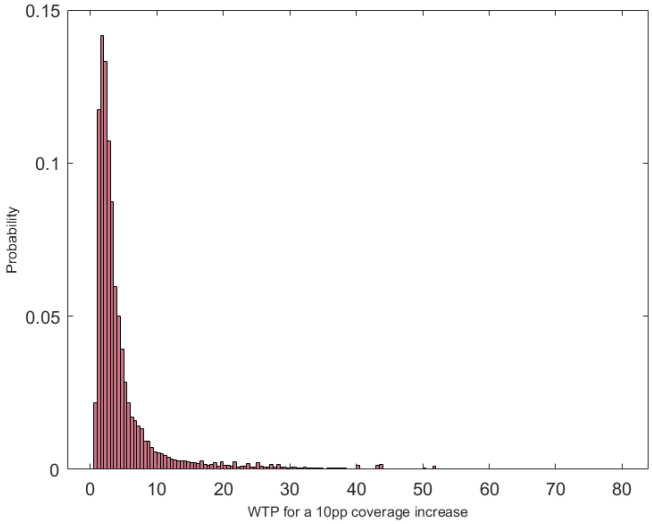
Notes: Price elasticities for plans, only women before the policy.

Figure 13: Price elasticities computed at base prices



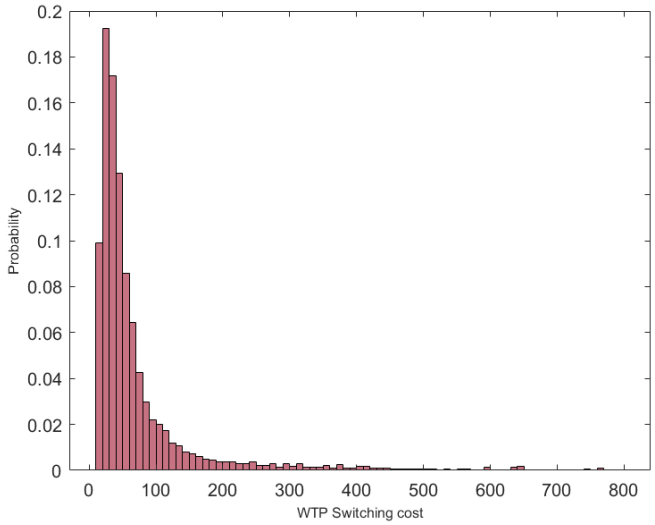
Notes: Own-price elasticities at base prices to compare price sensitivity between men and women.

Figure 14: Willingness to pay for coverage, for women before the policy



Notes: Willingness to pay for a 10 percentage point increase in the coverage rate, in UF per month. The histogram is for women before the policy.

Figure 15: Willingness to pay to avoid switching costs



Notes: Willingness to pay to have no switching costs, for women before the policy.

8 Welfare effects

In this section we measure how the gender pricing ban changed the welfare of women and men.

8.1 Consumer surplus

We assume that the utility that drives the choice of a plan is also the utility relevant for welfare, with one exception. Switching costs partly capture inattention and the cost of studying many plans. They affect choices, but we do not count them in consumer surplus. So when we compute consumer surplus, V_{ij} does not include the switching cost terms γ_1 and γ_2 of equation (11) (see Ericson and Starc (2016) for other welfare measures). We measure the change in consumer surplus as the equivalent variation, following Nevo (2001). For individual i facing a choice set Ω, consumer surplus is the logsum

$$CS_i(\Omega) = \frac{1}{-\alpha_i} \log\Big(\sum_{j\in\Omega} \exp(V_{ij})\Big), \tag{16}$$

up to a constant that cancels when we take differences.[10] Dividing by the price coefficient turns utility into money. Since prices are monthly premiums, consumer surplus is measured in UF per month. The change in consumer surplus after the policy is

$$\Delta CS_i = CS_i(\Omega_a) - CS_i(\Omega_b), \tag{17}$$

where Ω_b is the choice set before the policy and Ω_a the choice set after the policy. Each choice set includes the plans, their prices for the group of individual i, and the outside option. We report the

[10] McFadden (1999) shows that this formula is valid if there are no income effects. Our price coefficient depends on income. Train (2009) explains that the measure is still a good approximation when the changes are small. In the current computations we divide by the average price coefficient $\bar{\alpha}$ instead of α_i.

average of ΔCS_i over women and over men. The number of plans changes after the policy, and choice sets differ a lot across groups defined by gender and age. Since the logsum increases with the number of plans, the change in the number of plans also affects the measure.

8.2 Changes in consumer surplus after the policy

Table 1 shows the results. After the policy, the consumer surplus of women increases by 3.51 UF per month, around 151 US dollars.[11] This is more than the average monthly premium of women, which is 2.9 UF. The consumer surplus of men falls by 1.98 UF, around 85 US dollars, compared to an average monthly premium of 1.7 UF. The policy also changes enrollment. The share of women who choose the private sector increases by 11.5 percentage points, from 62% to 73%. For men, the share falls by 2.8 percentage points, from 63% to 60%.

Table 1: Consumer surplus and enrollment

	Women			Men		
	Before	After	Change	Before	After	Change
Share of enrollment in private sector	0.618	0.733	0.115	0.626	0.598	−0.028
Consumer surplus (average over individuals)	−33.609	−30.099	3.509	−4.414	−6.389	−1.975

Notes: The share of enrollment is the sum of the predicted choice probabilities of each individual over the plans of the private sector (inside options), averaged over individuals. Consumer surplus is in UF per month. Only the changes in consumer surplus have a meaning, not the levels.

Decomposition. Is consumer surplus increasing for women because of lower prices or because of the new plans? The policy changed the prices of plans, and insurers also changed the plans they offer. To separate the two, let $CS(P, X)$ be the average consumer surplus when individuals face plans X at prices P. X^{pre} and X^{post} are the plans sold before and after the policy. Before the policy, the price of plan j for group g is $\phi_{jg}P_j^{pre}$. We build an intermediate scenario in which individuals face the old plans X^{pre}, with their old base prices P_j^{pre}, but with the risk factors ϕ_g fixed by the policy. The change in consumer surplus is then

$$\Delta CS = \underbrace{CS(P^{post}, X^{post}) - CS(\phi_g P^{pre}, X^{pre})}_{\text{due to the new plans}} + \underbrace{CS(\phi_g P^{pre}, X^{pre}) - CS(\phi_{jg} P^{pre}, X^{pre})}_{\text{due to prices}}. \quad (18)$$

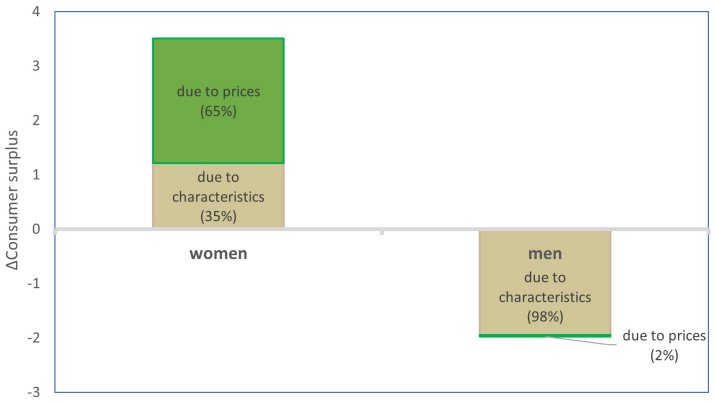
The second term is the effect of the new risk factors on the old plans. It measures what the policy would have done if insurers had not reacted. The first term is the effect of the change in plans after the policy. It includes both the new characteristics of the plans and their new base prices.

Figure 16 shows the results. For women, 2.29 UF of the 3.51 UF gain comes from prices and

[11] The UF (*Unidad de Fomento*) is a Chilean unit of account indexed to inflation. We convert UF to US dollars using the 2018 exchange rate, 1 UF ≈ 43 US dollars.

1.22 UF from the new plans. For men, almost all of the loss comes from the new plans. Prices explain only −0.03 UF of the −1.98 UF loss, and the new plans explain −1.95 UF. Therefore, the effect of the policy on prices alone explains only 0.03 UF of the 1.98 UF loss for men. The reaction of insurers matters for how the gains and losses of the policy are split between women and men.

Figure 16: Changes in consumer surplus



Notes: Decomposition of the change in consumer surplus in equation (18), in UF, for women and men.

8.3 Counterfactual: the policy applied to old contracts

The 2020 policy only applied to new contracts. Individuals who bought a plan before the policy kept paying the risk factors chosen by their insurer. In the counterfactual, old contracts also pay the risk factors set by the regulator, as the Supreme Court ordered in 2022.

Insurers do not react. Insurers keep the base prices of old plans and cannot introduce new plans, so individuals can only choose among the old plans. Only the risk factors change, from ϕ_{jg} to ϕ_g. The change in consumer surplus, averaged over all individuals with old contracts, women and men together, is

$$CS(\phi_g P^{pre}, X^{pre}) - CS(\phi_{jg} P^{pre}, X^{pre}) = 1.22 \text{ UF}. \tag{19}$$

This expression has the same form as the price term in equation (18). The difference is the population. Here we average over all individuals with old contracts, women and men together, while the decomposition reports women and men separately.

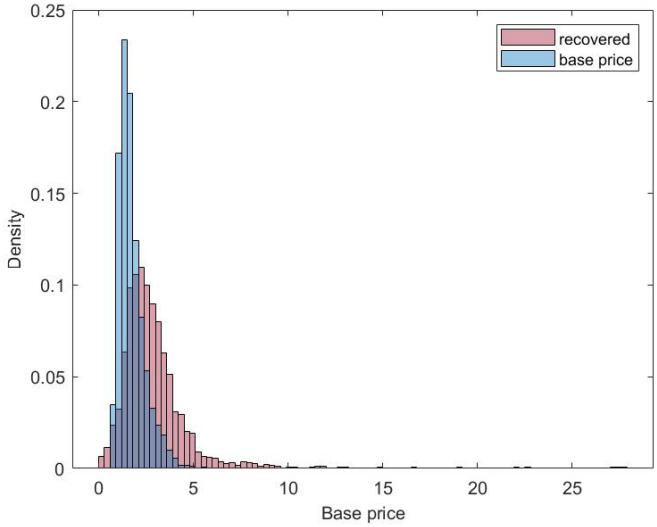
Insurers re-price old plans. Insurers choose new base prices for their old plans, with the risk factors fixed at ϕ_g. For this exercise we use a simpler version of the supply model in Section 5.1. Each plan is priced on its own, and costs are average costs by group, AC_{gj}, which we measure from claims. The profit of plan j is $\sum_g (\phi_g P_j - AC_{gj})\, s_{gj}$, where s_{gj} is the enrollment of group g in plan j. The first-order condition for the base price is

$$\sum_g \phi_g \left[s_{gj} \left(1 - \frac{\partial AC_{gj}}{\partial P_{gj}} \right) + (\phi_g P_j - AC_{gj}) \frac{\partial s_{gj}}{\partial P_{gj}} \right] = 0. \tag{20}$$

This is equation (6) applied to each group and added up with the weights ϕ_g, as in equation (4). We take $\partial s_{gj}/\partial P_{gj}$ from the demand estimates. We estimate $\partial AC_{gj}/\partial P_{gj}$ with a regression of average costs on price with insurer fixed effects. Prices respond to costs, so this estimate may be biased. In the full model, this term comes directly from the demand estimates and the individual costs from claims, as in equation (6). With adverse selection this term is positive and lowers the markup. With advantageous selection it is negative and raises the markup.

Figure 17 compares the observed base prices with the prices that solve equation (20). Insurers would set higher base prices if they could re-price their old plans. This suggests that applying the policy to old contracts lowers the profits of insurers at current prices.

Figure 17: Observed and recovered base prices



Notes: Observed base prices of plans in blue. The recovered price solves equation (20) with the risk factors fixed by the policy, using average costs from claims.

8.4 Welfare function

To compare pricing rules, we write the objective of the regulator following the public finance literature on social welfare weights (Saez & Stantcheva, 2016):

$$W = \sum_g \lambda_g \sum_{i \in g} CS_i + \sum_f \pi_f, \tag{21}$$

where CS_i is the consumer surplus of individual i in equation (16), λ_g is the weight on the consumer surplus of group g, and π_f are the profits of insurer f from equation (2). When $\lambda_g = 1$ for all groups, W is total surplus. Weights that differ across groups put more weight on the surplus of some demographic groups, such as women or older individuals. This is how redistribution enters the objective of the regulator, and it is the trade-off between efficiency and redistribution behind policies that ban pricing on observable characteristics.

Equation (21) can be used in two ways. First, it compares policies for given weights. Second, it gives the relative weights λ_g under which the 2020 policy is the best one among the alternatives. A higher weight for one group does not mean that the regulator does not value the surplus of the other groups. It shows how the policy trades off the surplus of different demographic groups.

9 Conclusion and Discussion

We study a policy that banned pricing by gender in the private health insurance market in Chile. Since the second quarter of 2020, men and women of the same age pay the same price for the same new plan. We estimate a discrete choice model of plan demand, in which individuals choose a plan based on its price, its characteristics, their expected out-of-pocket costs and their switching costs. We use the model to measure how the consumer surplus of women and men changed after the policy.

The policy increased the consumer surplus of women by 3.5 UF and lowered the consumer surplus of men by 2.0 UF. The share of women who choose the private sector increased from 62% to 73%, while the share of men fell from 63% to 60%. Most of the gain for women comes from lower prices. Almost all of the loss for men comes from the new plans that insurers offered after the policy, not from prices.

These results have three implications for policy. First, a ban on pricing by gender redistributes from men to women, but insurers respond by changing the plans they offer. The effect of the ban on prices alone explains only a small part of the loss for men. Therefore, regulators should also follow how plan design changes after a pricing rule. Second, the Chilean policy applied only to new contracts. Because switching costs are large, many individuals stay in their old plans and keep paying the old risk factors. Applying the policy to old contracts, as the Supreme Court ordered in 2022, raises consumer surplus if insurers do not react. But insurers would set higher base prices if they could re-price these plans, which suggests a loss in profits. This trade-off is at the center of the current debate in Chile. Third, the predicted share of women who choose the private sector rises from 62% to 73%. This may change the risk pool of the public insurer, FONASA, which matters for the design of the whole system.

Our next steps include estimating how insurers choose the coverage rates of their plans. The coverage decision captures a trade-off. More coverage makes a plan more attractive, but it attracts individuals with high expected costs the most, because their out-of-pocket costs fall the most. The model then gives the full welfare function and the relative welfare weights across groups under which the 2020 policy is the best one among the alternatives. We also plan to use it to study subsidies based on risk scores as an alternative to the gender pricing rule.

10 References

Abaluck, J., & Adams-Prassl, A. (2021). What do consumers consider before they choose? identification from asymmetric demand responses*. *The Quarterly Journal of Economics*, *136*(3), 1611–1663.

Abaluck, J., & Compiani, G. (2020). *A method to estimate discrete choice models that is robust to consumer search* (Working Paper No. 26849). National Bureau of Economic Research.

Abaluck, J., & Gruber, J. (2011). Choice inconsistencies among the elderly: Evidence from plan choice in the medicare part d program. *American Economic Review*, *101*(4), 1180–1210.

Abaluck, J., & Gruber, J. (2016). Evolving choice inconsistencies in choice of prescription drug insurance. *American Economic Review*, *106*(8), 2145–84.

Abi-Rafeh, R., Dubois, P., Griffith, R., & O'Connell, M. (2026). The effects of sin taxes and advertising restrictions in a dynamic equilibrium. *American Economic Journal: Microeconomics*, *18*(3), 39–76.

Atal, J. P. (2019). Lock-in in dynamic health insurance contracts: Evidence from chile. *Working Paper*.

Azevedo, E. M., & Gottlieb, D. (2017). Perfect competition in markets with adverse selection. *Econometrica*, *85*(1), 67–105.

Barseghyan, L., Coughlin, M., Molinari, F., & Teitelbaum, J. C. (2021). Heterogeneous choice sets and preferences. *Econometrica*, *89*(5), 2015–2048.

Bundorf, M. K., Levin, J., & Mahoney, N. (2012). Pricing and welfare in health plan choice. *American Economic Review*, *102*(7), 3214–48.

Cardon, J. H., & Hendel, I. (2001). Asymmetric information in health insurance: Evidence from the national medical expenditure survey. *The RAND Journal of Economics*, *32*(3), 408–427.

Coughlin, M. (2023). Insurance choice with non-monetary plan attributes: Limited consideration in medicare part d. *Working Paper*.

Crawford, G. S., Griffith, R., & Iaria, A. (2021). A survey of preference estimation with unobserved choice set heterogeneity [Annals Issue: Structural Econometrics Honoring Daniel McFadden]. *Journal of Econometrics*, *222*(1, Part A), 4–43.

Cuesta, J. I., Noton, C., & Vatter, B. (2024). Vertical integration and plan design in healthcare markets. *NBER Working Paper*, (32833).

Curto, V. (2023). Pricing regulations in individual health insurance: Evidence from medigap. *Journal of Health Economics*, *91*, 102785.

Curto, V., Einav, L., Levin, J., & Bhattacharya, J. (2021). Can health insurance competition work? evidence from medicare advantage. *Journal of Political Economy*, *129*(2), 570–606.

Dardanoni, V., & Li Donni, P. (2016). The welfare cost of unpriced heterogeneity in insurance markets. *The RAND Journal of Economics*, *47*(4), 998–1028.

Decarolis, F., Polyakova, M., & Ryan, S. P. (2020). Subsidy design in privately provided social insurance: Lessons from medicare part d. *Journal of Political Economy*, *128*(5), 1712–1752.

Einav, L., & Finkelstein, A. (2011). Selection in insurance markets: Theory and empirics in pictures. *Journal of Economic Perspectives*, *25*(1), 115–38.

Ericson, K. M. M., & Starc, A. (2016). How product standardization affects choice: Evidence from the massachusetts health insurance exchange. *Journal of Health Economics*, *50*, 71–85.

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

Finkelstein, A., Poterba, J., & Rothschild, C. (2009). Redistribution by insurance market regulation: Analyzing a ban on gender-based retirement annuities. *Journal of Financial Economics*, *91*(1), 38–58.

Fleitas, S. (2017). Dynamic competition and price regulation when consumers have inertia: Evidence from medicare part d. *Working Paper*.

Gandhi, A., & Houde, J.-F. (2019). *Measuring substitution patterns in differentiated-products industries* (NBER Working Paper No. 26375). National Bureau of Economic Research.

Geruso, M. (2017). Demand heterogeneity in insurance markets: Implications for equity and efficiency. *Quantitative Economics*, *8*(3), 929–975.

Geruso, M., Layton, T., & Prinz, D. (2019). Screening in contract design: Evidence from the ACA health insurance exchanges. *American Economic Journal: Economic Policy*, *11*(2), 64–107.

Handel, B., Hendel, I., & Whinston, M. D. (2015). Equilibria in health exchanges: Adverse selection versus reclassification risk. *Econometrica*, *83*(4), 1261–1313.

Handel, B. (2013). Adverse selection and inertia in health insurance markets: When nudging hurts. *American Economic Review*, *103*(7), 2643–82.

Handel, B., & Kolstad, J. (2015). Health insurance for "humans": Information frictions, plan choice, and consumer welfare. *American Economic Review*, *105*(8), 2449–2500.

Heiss, F., McFadden, D., Winter, J., Wuppermann, A., & Zhou, B. (2021). Inattention and switching costs as sources of inertia in medicare part d. *American Economic Review*, *111*(9), 2737–81.

Ho, K., & Lee, R. S. (2017). Insurer competition in health care markets. *Econometrica*, *85*(2), 379–417.

Ho, K. (2006). The welfare effects of restricted hospital choice in the us medical care market. *Journal of Applied Econometrics*, *21*(7), 1039–1079.

Huang, S., & Salm, M. (2020). The effect of a ban on gender-based pricing on risk selection in the german health insurance market. *Health Economics*, *29*(1), 3–17.

Lucarelli, C., Prince, J., & Simon, K. (2008). *The welfare impact of reducing choice in medicare part d: A comparison of two regulation strategies* (NBER Working Papers No. 14296). National Bureau of Economic Research, Inc.

Mahoney, N., & Weyl, E. G. (2017). Imperfect competition in selection markets. *The Review of Economics and Statistics*, *99*(4), pp. 637–651.

Marone, V. R., & Sabety, A. (2022). When should there be vertical choice in health insurance markets? *American Economic Review*, *112*(1), 304–342.

McFadden, D. (1999). *Computing Willingness-to-Pay in Random Utility Models* (tech. rep. $_0$11). in J.Moore, R.Riezman, J.Melvin, eds, 'Trade, Theory, and Econometrics: Essays in Honour of John S. Chipman', Routledge, London, pp. 253–274.

Nevo, A. (2001). Measuring market power in the ready-to-eat cereal industry. *Econometrica*, *69*(2), 307–342.

Pardo, C., & Sabat, J. (2025). Equity and efficiency effects of flat premiums. *International Journal of Health Economics and Management*, *25*(1), 27–49.

Polyakova, M., & Ryan, S. P. (2021). Market power and redistribution: Evidence from the affordable care act. *NBER Working Paper*, (26367).

Saez, E., & Stantcheva, S. (2016). Generalized social marginal welfare weights for optimal tax theory. *American Economic Review*, *106*(1), 24–45.

Serna, N. (2024). Determinants of healthcare provider networks: Risk selection vs. administrative costs. *Working Paper*.

Shepard, M. (2022). Hospital network competition and adverse selection: Evidence from the Massachusetts health insurance exchange. *American Economic Review*, *112*(2), 578–615.

Starc, A. (2014). Insurer pricing and consumer welfare: Evidence from medigap. *The RAND Journal of Economics*, *45*(1), 198–220.

Starc, A., & Town, R. J. (2019). Externalities and Benefit Design in Health Insurance. *The Review of Economic Studies*, *87*(6), 2827–2858.

Tebaldi, P. (2025). Estimating equilibrium in health insurance exchanges: Price competition and subsidy design under the aca. *Review of Economic Studies*, *92*(1), 586–620.

Train, K. E. (2009). *Discrete choice methods with simulation* (2nd ed.). Cambridge University Press.

Veiga, A., & Weyl, E. G. (2016). Product design in selection markets. *Quarterly Journal of Economics*, *131*(2), 1007–1056.

Appendix

A Descriptive statistics

We show some descriptive statistics on the health insurance plans in Table A1 and the number of visits and costs of the healthcare providers in Table A2.

Table A1: Statistics for subsample 2016–2021

	# traded plans	Mean	SD	p25	p50	p90	Min	Max
Men								
Outpatient coverage	179165	78.66	11.67	70	80	90	0	100
Inpatient coverage	179165	93.99	9.16	90	100	100	1	100
OOP cap	179165	4,248.07	1,660.68	3000	4200	6500	0	9000
Premium	179165	62,219.57	27,324.59	42940.93	55689.36	97596.79	9750.074	446191.4
Income	165768	689,564.54	547,086.57	329444	623911	1433371	0	4800000
Women								
Outpatient coverage	97896	74.47	13.40	70	80	90	0	100
Inpatient coverage	97896	91.23	10.53	90	90	100	1	100
OOP cap	97896	4,295.73	1,591.89	3000	4500	6500	0	9000
Premium	97896	86,996.27	34,035.88	64326.69	78960.16	131934.9	9768.107	500108
Income	92183	823,834.02	550,852.66	466781	782528	1582954	0	7000000
Total								
Outpatient coverage	277061	77.18	12.47	70	80	90	0	100
Inpatient coverage	277061	93.02	9.76	90	100	100	1	100
OOP cap	277061	4,264.91	1,636.86	3000	4200	6500	0	9000
Premium	277061	70,974.11	32,131.04	48548.84	64155.04	113800	9750.074	500108
Income	257951	737,547.93	552,196.11	375000	681317	1500000	0	7000000

Notes: Plans traded by women and men in the selected sample, 2016 to 2021. Coverage rates are in percent. Premium and income are monthly, in Chilean pesos.

Table A2: Statistics for providers, whole sample

	# providers	p25	p50	p90	p95	p99	Min	Max
2018								
# visits for provider	16710	2.000	6.000	335.500	1523.000	76895.000	1	1.48e+10
$ spent in provider (in MM)	16710	0.080	0.220	2.381	5.052	52.212	.005	1.18e+05
2019								
# visits for provider	15041	2.000	6.000	265.000	1208.000	1.37e+05	1	1.76e+10
$ spent in provider (in MM)	15041	0.075	0.210	2.318	5.004	66.719	.00025	1.40e+05
2020								
# visits for provider	12154	2.000	4.000	159.000	760.000	3.18e+05	1	1.11e+10
$ spent in provider (in MM)	12154	0.066	0.180	1.950	4.682	110.057	.002496	1.34e+05
2021								
# visits for provider	8674	1.000	4.000	132.000	871.000	1.37e+06	1	6.65e+09
$ spent in provider (in MM)	8674	0.060	0.156	1.370	3.478	152.570	.00472	93445.984
Total								
# visits for provider	52579	2.000	5.000	234.000	1154.000	1.88e+05	1	1.76e+10
$ spent in provider (in MM)	52579	0.070	0.200	2.101	4.667	81.899	.00025	1.40e+05

Notes: Taken from claims, and matched with all plans.

B Demand estimation details

B.1 More details on sample selection for the estimation of the model

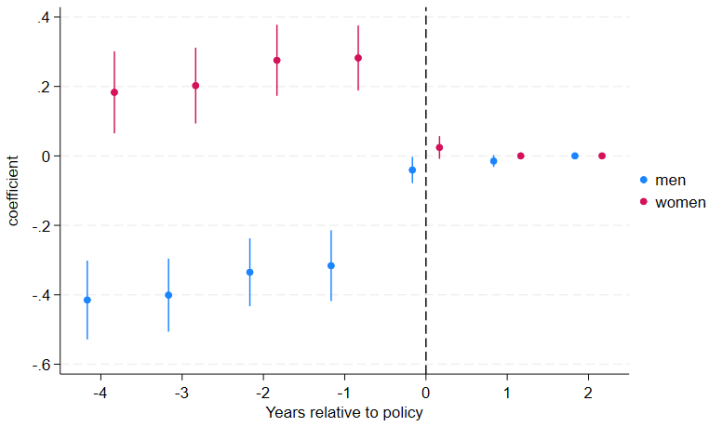
Besides taking the sample from 2018-2021 because we have information of the outside option from that year, we keep individuals with information on positive income, and we lose around 22% of the sample for enrollees of the private sector. This rate is lower for the years after the policy. We also eliminate individuals with plans that have coinsurance rates of 0 and 1. We provide a comparison between the two subsamples in Table A3.

Table A3: Sample statistics for women

	count	Mean	Sd	Min	Max
age_lb	8104473	39.87099	10.77658	21	61
id_insurer	8104473	81.17396	15.54728	0	107
id_plan	8104473	138.641	91.92032	0	336
ownplan	8104473	.012943	.1130286	0	1
outpatient	8104473	75.20697	11.56491	40	100
inpatient	8104473	88.93342	13.63584	30	100
income	8104473	97.37176	47.28596	.0036277	580.5469
e_oop	7915057	3.04266	1.873238	-1.306433	13.95044
v_oop	8104473	1934.285	353.322	7.65e-10	2000

Notes: Sample of women used to estimate the demand model, from 2018.

Figure 18: Event study for prices between women and men, with insurer and age fixed effects



Notes: Same regression as in Figure 1, showing coefficients for both genders, with insurer and age fixed effects instead of aggregated plan fixed effects. The regression controls for the other characteristics of plans.

B.2 Construction of variables

Section 6 describes the instruments for price. The number of people $n_{r,g,j}^{(2017)}$ comes from the 2017 CASEN survey. It counts the people of group g outside the region of study whose income allows them to buy plan j. The interaction instrument Z_{gjt} in equation (15) multiplies this number by the yearly increase in the Health CPI of the Chilean Central Bank. The differentiation instruments count the rival plans that are in the same choice set as plan j for a group, in total and for each hospital h.

B.3 Demand Estimation Results

Table A4: Demand estimates for women before the policy

	Parameter	Estimate	Std. error
Panel A. First step (maximum likelihood)			
Price × age	α_1	−0.584	
Price × income	α_2	0.202	
Expected out-of-pocket costs	ω_0	1.100	
Expected out-of-pocket costs × age	ω_1	−2.197	
Expected out-of-pocket costs × income	ω_2	0.468	
Variance of out-of-pocket costs	ω_3	−0.041	
Switching plans	γ_1	−5.808	
Switching to a plan sold after the policy	γ_2	1.003	
Panel B. Second step (two-stage least squares)			
Price	α_0	−0.059	0.014
Coverage rate	β^1	3.249	0.351
Restriction fixed effects			Yes
Insurer fixed effects			Yes

Notes: Estimates of the demand model in equation (11). Panel A shows the parameters of individual heterogeneity, estimated by maximum likelihood together with δ_{gjt}. Panel B shows the parameters of equation (14), estimated by two-stage least squares with the instruments in Table A5. Prices are monthly premiums in UF. Age is in years divided by 100, and income is monthly income in UF divided by 100.

Table A5: First stage: price on instruments

	Coefficient	Std. error
Number of individuals in the group in 2017, $n_{r,g,j}^{(2017)}$	0.248	0.020
Interaction instrument, Z_{gjt}	−0.125	0.010
Z_{gjt} × coverage rate	−0.002	0.003
Number of rival plans	0.121	0.292
Number of rival plans with hospital h (one for each hospital)		Yes
Restriction fixed effects		Yes
Insurer fixed effects		Yes

Notes: First stage of the two-stage least squares estimation of equation (14). The dependent variable is the price P_{gjt}. The instruments are described in Section 6 and Appendix B.

C Selection term in the pricing first-order condition

Take an insurer with a single plan j and a single price P_j. Its profit is $\sum_i (P_j - \tilde{c}_{ij})\, s_{ij}$, and the first-order condition is

$$\sum_i s_{ij} + \sum_i \big(P_j - \tilde{c}_{ij}\big) \frac{\partial s_{ij}}{\partial P_j} = 0.$$

Let $s_j = \sum_i s_{ij}$ and $AC_j = \sum_i s_{ij}\, \tilde{c}_{ij} / s_j$. Taking the derivative of $s_j AC_j = \sum_i s_{ij}\, \tilde{c}_{ij}$ with respect to P_j gives

$$\sum_i \tilde{c}_{ij} \frac{\partial s_{ij}}{\partial P_j} = s_j \frac{\partial AC_j}{\partial P_j} + AC_j \frac{\partial s_j}{\partial P_j}.$$

Replacing this in the first-order condition gives

$$s_j + P_j \frac{\partial s_j}{\partial P_j} - s_j \frac{\partial AC_j}{\partial P_j} - AC_j \frac{\partial s_j}{\partial P_j} = 0,$$

which we can write as equation (6). So the individual first-order condition already contains the selection term of Mahoney and Weyl (2017). The term $\partial AC_j / \partial P_j$ is not zero when individuals with different costs respond differently to the price.

D Coverage first-order condition

The effect through rival prices in equation (7) has two parts. The first part is the effect of the price of a rival plan $k \in J_{-f}$ on the profits of insurer f. For group g,

$$\frac{\partial \pi_f}{\partial P_{gk}} = \sum_{i \in g} \sum_{l \in J_f} \big(P_{gl} - \tilde{c}_{il}\big) \frac{\partial s_{il}}{\partial P_{gk}}, \qquad \frac{\partial s_{il}}{\partial P_{gk}} = -\alpha_i\, s_{il}\, s_{ik}.$$

When a rival raises its price, some of its enrollees move to the plans of insurer f, and insurer f earns its margin on them. The second part is the response of rival prices to the coverage of plan j, $\partial P^*_{gk}/\partial c_j$. We obtain it from the implicit function theorem. We stack the price first-order conditions (3) or (4) of all insurers. The vector of price responses is minus the inverse of the matrix of derivatives of these conditions with respect to prices, times the vector of derivatives of these conditions with respect to c_j.