Contribution-Floor Mandates: A Welfare Framework and Evidence from Chile

Mateus Dias*

May 2026

Abstract

Many health insurance systems link premiums to income through contribution floors: enrollees must spend at least a fixed share of their income on health insurance. These income-based mandates are economically distinct from the penalty-based mandates studied in the U.S. context, because they create a wedge between the premium paid and the plan's actuarial price. I develop a framework that decomposes the welfare effect of a contribution-floor mandate change into a *fiscal channel*, the mechanical transfer arising from the price floor, and an *equilibrium channel*, the behavioral response through plan switching and risk pool recomposition coupled with supply-side responses. I implement the framework using administrative data from Chile's private health insurance market (ISAPREs), where enrollees must pay at least 7% of their income in premiums. Raising the mandate from 7% to 8% of income redistributes approximately \$38.9 million annually from enrollees to insurers while generating essentially no net welfare: about 12 thousand USD, three orders of magnitude smaller than the gross transfer. The equilibrium channel is negligible. Because the mandate operates as a near-pure transfer, it is dominated as a redistributive instrument by a revenue-equivalent direct cash transfer, which delivers roughly three times the distributionally weighted welfare.

JEL: I13, I18, G22.

Keywords: contribution-floor mandate; income-based mandate; health insurance regulation; adverse selection; welfare decomposition.

*Catolica-Lisbon School of Business and Economics; e-mail: *matdias@ucp.pt*. I thank the Chilean *Superintendencia de Salud* and the *Instituto Nacional de Estadísticas* for providing access to their datasets. All the remaining errors are my own.

1 Introduction

Mandates are among the most important regulatory instruments in health insurance markets. The literature has analyzed how requiring individuals to purchase insurance affects adverse selection, risk pool composition, and welfare. Yet, this literature has focused on a specific type of mandate: the penalty-based mandate, in which individuals face a financial penalty for not purchasing coverage. The individual mandates in Massachusetts and under the Affordable Care Act are prominent examples, and the framework for analyzing their welfare effects is well established (Hackmann et al., 2015). However, the mandates used in many health insurance systems outside the United States take a structurally different form. In Chile, Germany, the Netherlands, Switzerland, and Colombia, among others, health insurance systems feature income-proportional contribution floors, where enrollees must spend at least a fixed share of their income on health insurance premiums. These contribution-floor mandates are economically distinct from penalty-based mandates, and their welfare effects have not been analyzed in the literature.

The distinction is fundamental. A penalty-based mandate imposes a cost for being uninsured, but does not alter the price of insurance conditional on purchasing it and enrollees pay the premium of their chosen plan, nothing more. A contribution-floor mandate, in contrast, can directly alter the price of insurance: the premium for individual i in plan j is $p_{ij} = \max\{\tilde{p}_{ij}, \tau y_i\}$, where $\tilde{p}_{ij}$ is the risk-adjusted premium and τy_i is the mandate floor. Whenever $\tau y_i > \tilde{p}_{ij}$, the enrollee is bound by the mandate, i.e., she pays more than the baseline premium of her plan, and a wedge arises between the premium paid and the plan's price. This wedge is a fiscal transfer that has no counterpart in penalty-based mandate designs.[1]

This paper makes two contributions. First, I develop a framework for analyzing the welfare effects of contribution-floor mandates. The framework decomposes the welfare effect of a mandate change into a *fiscal channel* and an *equilibrium channel*. The fiscal channel captures the mechanical transfer of resources that occurs when the income-based premium floor changes, holding plan choices fixed. The equilibrium channel captures the

[1]The ACA individual mandate also included an income-dependent component: the penalty was the greater of a flat amount or 2.5% of income above the filing threshold. However, this remained a penalty for non-insurance, since it did not create a price floor on insurance premiums. Conditional on purchasing a plan, ACA enrollees paid the plan's premium regardless of their income. The penalty was also capped at the cost of the cheapest available bronze plan, ensuring that no wedge could arise between the premium paid and the plan's actuarial value.

behavioral response (i.e., consumers may switch plans, potentially shifting risk pools), and the premium adjustments by insurers. This decomposition reveals that contribution-floor mandates bundle two analytically separable economic forces, and that the relative magnitude of each is an empirical question. The framework applies regardless of the specific institutional channel through which the wedge is mediated, whether it accrues to the insurer directly (as in Chile and Colombia), flows to a centrally pooled fund (as in Germany), or is offset by government subsidies (as in the Netherlands and Switzerland).

Second, I implement the framework empirically using detailed administrative data from Chile's private health insurance market, where enrollees must pay at least 7% of their income in premiums. The main finding is that the mandate operates as a near-pure fiscal transfer and, when equity considerations are taken into account, a simpler tax-and-transfer scheme can generate three times as much welfare gains. Raising the mandate from 7% to 8% of income redistributes approximately 38.9 million U.S. Dollars (USD, in 2012 values) from mandate-bound enrollees to insurers, with a net welfare effect of approximately 12,000 USD (negligible relative to the gross flows). Essentially none of the net welfare effect is attributable to selection. Although, in my counterfactual exercises, this is in part because plan choice probabilities barely move, I also show that the scope for gains via selection is virtually non-existent: the typical switchers tend to be slightly more costly than incumbents, so even a larger change would not generate strong selection gains. Thus, the mandate increase simply raises the price floor for bound enrollees, and the additional revenue flows to insurers as profit.

To understand why contribution-floor mandates generate this decomposition, consider what happens when τ increases. Two groups of enrollees are affected. First, those who were already bound ($\tau y_i > \tilde{p}_{ij}$ at the old rate) pay more, but their plan choice is unaffected as long as the ranking of plans does not change. This is the fiscal channel, a mechanical increase in the price floor that transfers resources without altering behavior. Second, some previously unbound enrollees become bound, and the compression of effective premium differentials across plans may induce plan switching among both old and new bound enrollees. Additionally, the increase in the revenue in the form of premiums from the bound incumbents may generate some competitive pressure, lowering base premiums. This is the equilibrium channel through which the mandate could, in principle, improve welfare by correcting adverse selection, the mechanism emphasized in the penalty-based mandate literature. The empirical question is whether this second channel is quantitatively important.

The theoretical framework formalizes this intuition. I set up a two-stage model in which agents choose a health insurance plan in the first stage and make health spending decisions in the second stage, after a health shock is realized. The model embeds the contribution floor via the price floor, and I show analytically that the welfare effect of a change in the mandate decomposes into a fiscal and an equilbrium components. A key feature of the model is that the utility of consumption is strictly concave, departing from the quasi-linear specification commonly used in the literature (Einav et al., 2013; Marone and Sabety, 2022). This departure is not incidental: with quasi-linear utility, the willingness to pay for insurance is independent of income, and the mandate floor has limited interaction with the demand side. With concave utility, the marginal value of consumption varies with income, and this interacts with the willingness to pay for insurance: at low incomes, the marginal utility of consumption is too high to justify expensive coverage; at high incomes, risk aversion becomes less important and willingness to pay declines. In my setting, this further interacts with the mandate floor, which generates a potentially complex relationship between income and willingness to pay for insurance that a non-linear utility is better suited to capture. This non-linearity also lets the relationship between healthcare spending and income emerge more naturally from the model, a correlation that plays a critical role in my analysis.

For the empirical implementation, I estimate a structural model of plan choice and health care utilization using administrative data on enrollees in Chile's private health insurance market (ISAPREs) in 2012. The Chilean setting is particularly well-suited for this analysis. The mandate floor is transparent and directly observed. Risk-adjusted premiums are computed from observed base premiums and regulatory risk factor tables. The market offers meaningful plan differentiation in terms of coverage and prices, and the administrative records include individual-level health spending, income, demographics, and plan choices. I estimate the model via Markov Chain Monte Carlo, treating the moral hazard parameter as a latent variable recovered through data augmentation, and inverting the first-order condition to pin down the health shocks from individual observed spending.

The counterfactual exercises raise the mandate from 7% to 8% of income and compute the new equilibrium under a constant-wedge assumption for insurer pricing. The decomposition proceeds in two steps. The fiscal counterfactual applies the 8% floor but holds plan choices at their baseline (7%) probabilities, so that the welfare change comes exclusively from the mechanical increase in premiums for bound enrollees. The full counterfactual allows consumers to re-optimize their plan choices and insurers to adjust premiums in response to the new composition of the risk pool. The equilibrium channel is the difference between the full and fiscal counterfactuals. I find that the equilibrium channel is economically negligible: the compensating variation under the fiscal counterfactual is virtually identical to that under the full counterfactual, both in aggregate and across the income distribution.

I then compare the mandate increase to a direct transfer policy that raises the same revenue through an extra 1% income tax on people who would be bound in the 8% mandate counterfactual, but distributes it as a lump-sum cash to unbound enrollees. Since the mandate operates as a near-pure transfer, routing the redistribution through the insurance market rather than through direct transfers is costly: the insurer captures the wedge, and the incidence falls disproportionately on middle-income enrollees who are bound but receive little additional coverage value. At moderate levels of inequality aversion, the direct transfer generates greater welfare gains because it avoids insurer intermediation and can be targeted more efficiently.

These empirical findings are specific to the Chilean setting, but the framework is general. In the Chilean market, where premiums are risk-rated and the mandate operates on the intensive margin of plan choice (all ISAPRE enrollees are already insured), the equilibrium channel is small because the mass of switchers is low and the *de facto* risk-rating that insurers are able to implement already internalizes much of the cost variation with income. In pure community-rated markets, or in settings where the contribution floor operates on the extensive margin of insurance take-up, the equilibrium channel could be larger. The decomposition framework is designed to quantify this in any specific setting, and the magnitudes of the fiscal and equilibrium channels are empirical questions that depend on market structure, the degree of risk-rating, and the price sensitivity of plan choices.

This paper contributes to several strands of the literature. First, it contributes to the literature on mandates and welfare in insurance markets. The conceptual foundations for welfare analysis in insurance markets with adverse selection are provided by Einav et al. (2010) and Einav and Finkelstein (2011), who develop the cost-curve and willingness-to-pay framework that has become the standard tool for evaluating regulatory interventions. Hackmann et al. (2015) apply this framework directly, using pre- and post-mandate equilibria in Massachusetts to trace out demand and average cost curves and decompose the welfare effect into an adverse selection correction and a markup reduction, finding large welfare gains from bringing healthier individuals into the risk pool. Despite being mainly methodological, Azevedo and Gottlieb (2017) apply their method to analyze mandates within a competitive equilibrium framework, characterizing how mandates shift equilibrium allocations under adverse selection and the potential unintended consequences of them. Bundorf et al. (2012) study welfare in employer-sponsored plan choice using a revealed-preference approach, quantifying the efficiency costs of adverse selection in a setting with multiple plans. All of these papers, and the broader literature on welfare in insurance markets, analyze settings in which the regulatory instrument operates through channels that do not directly alter the price of insurance conditional on purchase. I study a fundamentally different instrument. Contribution-floor mandates create a wedge between the premium paid and the plan's actuarial value, generating a fiscal transfer channel that is absent under penalty-based mandates. The decomposition into fiscal and equilibrium channels is new, and the empirical finding that the equilibrium channel is negligible in Chile's risk-rated market contrasts with the penalty-based mandate literature, where the selection correction is the primary source of welfare gains. This contrast is not a theoretical prediction, however: it is an empirical result that depends on the degree of risk-rating and the structure of the market.

Second, it contributes to the growing literature on market design in regulated health insurance. A range of regulatory instruments and their effects on market outcomes was previously studied in the literature: subsidies (Finkelstein et al., 2019; Tebaldi, 2025), risk adjustment (Glazer and McGuire, 2000; Brown et al., 2014; Einav et al., 2025), and market structure and pricing regulation (Curto et al., 2021; Marone and Sabety, 2022). Einav et al. (2025) represents a recent push toward studying these instruments jointly, showing that subsidies and risk adjustment interact through their effects on equilibrium pricing and markups. This paper adds contribution floors to this toolkit, a distinct and internationally prevalent instrument. The finding that Chile's contribution floor functions primarily as a fiscal transfer, with negligible selection effects, has direct implications for how policymakers in countries with income-based mandates should evaluate changes in contribution rates. It suggests that, at least in risk-rated markets, the mandated contribution is better understood as a tax-and-transfer scheme than as a tool for improving market efficiency.

Third, it contributes to the literature on structural estimation of insurance demand and moral hazard. I build on the framework of Einav et al. (2013), departing from their quasi-linear consumption utility by assuming strict concavity. This departure is motivated by the empirical setting: with a contribution-floor mandate, the interaction between income and the willingness to pay for insurance is central to the analysis, which is naturally generated by a concave utility of consumption. The use of concave utility is not without precedent (Cardon and Hendel (2001), for instance, use CARA preferences) but it is uncommon in recent work on moral hazard and plan choice, which has largely followed the quasi-linear tradition. The model also relates to Marone and Sabety (2022), who estimate a model of vertical plan choice with moral hazard in the ACA marketplace, and to Handel (2013), who combines adverse selection with inertia in a structural framework used for counterfactual market design. My setting differs in that the contribution floor creates a price floor that varies continuously with income, generating a source of premium variation that is absent in community-rated markets and that is the object of the welfare analysis.

2 Institutional Setting

2.1 The Chilean Health Insurance System

Chile's health system is characterized by the coexistence of a public and a private system.[2] Workers in the formal sector and retirees are required to acquire health insurance, choosing between the public option, administered by an institution called FONASA (*Fondo Nacional de Salud*), and the private plans offered by insurers known as ISAPREs (*Instituciones de Salud Previsional*). In 2012, 76.5% of the Chilean population was enrolled in FONASA, whereas 17.5% had acquired a private health plan.[3] These shares were relatively stable throughout the decade.

A central feature of the system is the mandatory contribution: irrespective of specific choices, each person must pay at least 7% of her income as a health insurance premium.[4] This income-based mandate is the policy instrument at the center of this paper. It dif-

[2]For more detailed descriptions of the Chilean health system, see Atal (2019), Duarte (2012), Palmucci and Dague (2014), and Cuesta et al. (2025).

[3]The rest of the population was either in a separate system, such as the armed forces' health system, or was uninsured (Palmucci and Dague, 2014).

[4]There is a cap on contribution income (*tope imponible*) of approximately 2,500 US dollars per month (in 2012 values), which is binding for a non-negligible fraction of higher-income enrollees. I discuss how this cap is handled in the empirical model in section 3.

fers fundamentally from the binary mandates studied in the U.S. context — such as the Massachusetts individual mandate analyzed by Hackmann et al. (2015) — because the contribution is proportional to income rather than a fixed penalty. As I show in the theoretical framework, this distinction has important implications for the welfare effects of mandate changes.

2.2 The Public Option: FONASA

Enrollees in the public option are classified into one of four income-based groups that determine their coinsurance rates for public healthcare services. Groups A and B (the poorest enrollees) face zero coinsurance, while groups C and D face coinsurance rates of 10% and 20%, respectively. People in the public plan have access to public providers, and those in groups B through D can also choose to visit private providers in the FONASA network, albeit at substantially higher out-of-pocket costs. In practice, most public option enrollees rely on public healthcare services (Cuesta et al., 2025).

Movement between the FONASA and ISAPRE systems is limited. The two systems serve substantially different populations: private plan enrollees are richer, younger, healthier, and more likely to be male than public option enrollees, as illustrated in Appendix Figures A.1–A.3. While there is meaningful overlap in the income distributions of the two systems (Appendix Figure A.1), the differences in age composition (Appendix Figure A.2) and health status (Appendix Figure A.3) are stark. This segmentation, combined with the very different nature of the products offered — public providers with low copays versus private providers with risk-rated premiums — means that the two markets are largely separate.[5]

2.3 The Private Market: ISAPREs

The private market is the focus of this paper. In 2012, there were seven open ISAPREs (whose plans are available to anyone) and a small number of closed ISAPREs (which offer plans only to the employees of specific affiliated firms). The five largest open ISAPREs cover approximately 95% of the private market.[6]

[5]This is consistent with the treatment of the two systems in the Chilean health economics literature (Duarte, 2012; Pardo, 2019; Cuesta et al., 2025). Sapelli (2004) documents the risk segmentation between systems and show that adverse selection operates primarily at the FONASA-ISAPRE boundary rather than within either system.

[6]The closed insurers had only 3.7% of total private market enrollees in 2012.

Premiums and the mandate. Each private plan has a base premium, which is then adjusted by a demographic-specific risk factor. ISAPREs are allowed to maintain up to two risk-pricing tables per insurer, establishing the multipliers applied to the base premium according to the enrollee's age and gender.[7] The risk-adjusted premium determines the plan's actuarial price for a given individual. However, the mandate imposes a floor: the premium an enrollee pays is the maximum of the risk-adjusted premium and 7% of her income. That is, the premium for individual i in plan j is

$$p_{ij} = \max\{\tilde{p}_{ij},\ 0.07 \times y_i\}$$

where $\tilde{p}_{ij}$ is the risk-adjusted premium and y_i is the enrollee's income. When $0.07 \times y_i > \tilde{p}_{ij}$, the enrollee is *bound* by the mandate: the insurer receives the full mandatory contribution, which exceeds the actuarially determined premium.

This is the mechanism that distinguishes income-based mandates from binary mandates. Under a binary mandate, the penalty for not purchasing insurance is a fixed amount — the effective price of insurance is reduced uniformly by π for everyone. Under an income-based mandate, the effective price of insurance is $\max\{\tilde{p}_{ij}, \tau y_i\}$, which depends on the enrollee's income. For enrollees whose risk-adjusted premium exceeds the mandate floor, the mandate has no effect on the price they pay. For bound enrollees, the mandate acts as an income-proportional tax whose revenue accrues to the insurer. This creates a fiscal transfer channel that is absent in binary mandate designs.

Excedentes. When the mandatory contribution exceeds the plan price, the difference, known as the *excedente*, is deposited into an individual account managed by the ISAPRE. These funds can be used for copayments, uncovered health services, and other health-related expenses.[8] However, ISAPREs routinely adjust plan prices to absorb the contribution through the annual *adecuación* process. Specifically, when excedentes exceed 10% of the legal contribution in successive annual adjustments, ISAPREs are required to offer

[7] Recently, and not directly relevant to this paper, gender-based pricing was abolished. See also Cuesta et al. (2025), who document that insurers effectively achieve finer demographic price discrimination by creating multiple plan codes with identical coverage but different base premiums targeted at different demographic segments. In section 4, I describe how I handle this in the data.

[8] See Superintendencia de Salud (2013) for details. Since 2009, excedentes are irrenunciable: enrollees cannot waive them in exchange for plan upgrades (Ley N° 20.317).

an alternative plan whose price better matches the contribution.[9] Given their negligible prevalence, I abstract from excedentes in the model.

Coverage and plan characteristics. Private plans in Chile do not have deductibles. They do, however, impose limits on insurer liability: an annual cap on total covered expenses (generally high and rarely binding) and a per-procedure cap on the amount the insurer covers each time a procedure is performed. These per-procedure caps mean that the effective coverage of a plan can differ from the coinsurance rates used to market it. Regarding provider networks, unlike the U.S., plan networks in Chile are usually open. They are either fully open or have tiered networks with preferred providers (as in a PPO), with a negligible fraction of enrollees in plans with closed networks.

Preexisting conditions. Private plans can partially or completely deny coverage to newly enrolled individuals based on a declaration of preexisting conditions.[10] Enrollees who switch insurers may lose guaranteed renewability protections, which creates a degree of lock-in within the private system. This institutional feature reinforces the segmentation between FONASA and ISAPREs, as switching from a private plan to the public option (or vice versa) involves re-entering a system with very different rules.[11]

2.4 Relevance to Other Income-Based Mandate Systems

Many health insurance systems worldwide link the price of insurance to income. The specific institutional mechanisms vary. In some systems, such as Chile and Colombia, the mandatory contribution is proportional to income and flows directly to the insurer (Guerrero et al., 2011). In Germany's statutory health insurance, income-dependent contributions account for the bulk of system financing and are pooled centrally before be-

[9]As of December 2012, while 48.2% of ISAPRE affiliates had an active excedente account from past accumulation, the average per-capita balance was approximately 84,000 CLP ($160 USD) (Superintendencia de Salud, 2013). In the full *cotizantes* data for 2012, only 1.3% of enrollees have a total premium below their 7% contribution. This confirms that the *adecuación* mechanism effectively eliminates most excedentes.

[10]Atal (2019) studies this issue in more detail, finding that about 20% of people with a preexisting condition in the private market are affected by this feature.

[11]This is an additional reason to treat the two systems as largely separate markets beyond the differences in populations and product characteristics. Moreover, the preexisting condition mechanism primarily affects switching between insurers, not plan choices within an insurer. Since my counterfactual exercises involve a 1 percentage point change in the mandate rate, a perturbation too small to plausibly induce insurer switching, I abstract from this feature in the model. If anything, the resulting lock-in further limits the scope for plan resorting, reinforcing the finding of a small equilibrium channel.

ing redistributed to sickness funds via risk adjustment (Blumel and Busse, 2020). In the Netherlands, roughly half of total system funding comes from income-related contributions, alongside community-rated premiums (Kroneman et al., 2016). In Switzerland, premiums are community-rated per plan and region, and the income-linked component of insurance pricing enters through cantonal and federal subsidies that reduce the effective price for lower-income households (De Pietro et al., 2015). Kauer et al. (2026) provide a recent comparative analysis of regulated competition in health insurance markets across Germany, the Netherlands, Switzerland, and the U.S. Marketplaces, documenting how income-based financing interacts with risk adjustment and competition in these systems. Despite these institutional differences, all of these systems share a common feature: the effective price of insurance depends on income, whether through the premium itself, through centrally pooled contributions, or through subsidies.

This common feature is what makes the welfare decomposition developed in this paper broadly relevant. In any system where the income-linked component of insurance pricing changes, whether by raising a contribution rate, adjusting a subsidy schedule, or modifying a tax, two forces are at play. The first is a mechanical fiscal transfer: resources move between consumers, insurers, or the government as a direct consequence of the price change. The second is a potential behavioral response: consumers may alter their plan choices, changing risk pool composition and potentially affecting market efficiency. The framework I develop separates these two forces and quantifies the relative importance of each.

While the Chilean ISAPRE market is the setting for my empirical analysis, the conceptual contribution, i.e., that income-based mandates operate through a fiscal channel that is absent in binary mandates, and that this channel can dominate the equilibrium channel, applies whenever the price of insurance is tied to income. The relative magnitudes of the fiscal and equilibrium channels in any specific country, however, are empirical questions that depend on the degree of risk-rating, the structure of the plan menu, and the price sensitivity of plan choices. The Chilean market provides a particularly clean setting for this analysis because the mandate floor is transparent, risk-adjusted premiums are observed in administrative data, and the within-system plan choice allows for a detailed structural estimation of the demand for insurance coverage.

3 Model

Here, I present a two-stage model of health plan and utilization to guide the analysis. The empirical specification used in the estimation is based on this model, but with a richer environment that more closely matches the empirical setting.

Assume agents are in a market where two health plans exist: plan H, with coverage $c_H = 1$ (i.e., full coverage), and plan L, with coverage $c_L \in (0, 1)$. Agent i is characterized by (y_i, ω_i, F_i), where y_i is her income, ω_i a parameter that will control her *ex-post* moral hazard *conditional on income*, and F_i is the cdf of the distribution of her health shock λ_i, where $\lambda_i \geq 0$ and higher values denote worse health shocks.

In the first stage, agents must choose their health plan without knowing the realization of λ. Each agent i chooses her plan j according to

$$j = \underset{j'}{\text{argmax}}\ V_{ij'}$$

where

$$V_{ij'} = \mathbb{E}_{\lambda_i}[U_{ij'}]$$

In the equation above, $U_{ij'}$ denotes the second stage utility of plan j' for agent i and the expectation is taken with respect to the distribution of the health shock λ_i.

In the second stage, the health shock is realized and the agent chooses consumption z_{ij} and health spending h_{ij} conditional on the realization λ_i and the plan choice j. The second stage utility is given by

$$U_{ij} = u(z_{ij}) + f(h_{ij}, \lambda_i, \omega_i)$$

$$\text{s.t. } z_{ij} = y_i - p_{ij} - (1 - c_j)h_{ij}$$
$$z_{ij} > 0\,, h_{ij} \geq 0$$

where the utility of consumption and the utility derived from health spending are assumed to be additively separable and p_{ij} is the premium of plan j to agent i. Assuming the mandate is such that enrollees must pay in premiums at least τy_i,

$$p_{ij} = \max\{\tilde{p}_j, \tau y_i\}$$

where $\tilde{p}_j$ is the premium floor set by the insurer. Moreover, we make the following assumptions on the functions u and f:

$$u' > 0\,,\, u'' < 0\,,\, u(0) = -\infty\,,\, \lim_{z\to\infty} u(z) = \infty$$

$$\frac{\partial f}{\partial h} > 0\,,\, \frac{\partial f}{\partial \lambda} < 0\,,\, \frac{\partial^2 f}{\partial h^2} < 0\,,\, \frac{\partial^2 f}{\partial h \partial \lambda} > 0$$

Assume also that, under the healthiest possible state $\lambda = 0$, there is no value of spending anything in health:[12]

$$\frac{\partial f}{\partial h}(h, 0, \omega_i) = 0 \quad \forall h \geq 0$$

We can then substitute the budget constraint into the second stage utility function; the agent's problem in the second stage then becomes

$$\max_{0 \leq h < y_i - p_{ij}} u(y_i - p_{ij} - (1 - c_j)h) + f(h, \lambda_i, \omega_i)$$

Second-stage solution. Taking first-order conditions,

$$-(1 - c_j)u'(y_i - p_{ij} - (1 - c_j)h_{ij}^{\lambda_i}) + \frac{\partial f}{\partial h}(h_{ij}^{\lambda_i}, \lambda_i, \omega_i) = 0$$

Where $h_{ij}^{\lambda_i}$ denotes the optimal value of medical spending. Given the assumptions made on u and f, the solution $h_{ij}^{\lambda_i}(y_i - p_{ij}, c_j, \lambda_i, \omega_i)$ exists, is unique, and

$$h_{ij}^{\lambda_i} \in [0, y_i - p_{ij})$$

One point is worth noticing. As previously mentioned, ω_i controls the degree of moral hazard conditional on income, which differs from models that are linear in consumption, typically used in the literature. The distinction becomes clearer when we look at the

[12] I define the partial derivative of f with respect to h at zero as the right derivative.

first-order condition: here, income itself impacts how much agents respond to changes in coinsurance rates (or the price of acquiring health care) due to the concavity (or, more generally, the non-linearity) of the consumption utility function. Given that low-income individuals have higher marginal utility of consumption, this direct income effect makes lower-income agents more sensitive to the price of healthcare (Acquatella, 2024). The parameter ω_i controls the degree of the price response conditional on the income level.[13]

We can also use the first-order condition to characterize how changes in λ_i, y_i, and c_j affect the solution through the implicit function theorem. Applying that, we get that

$$\frac{dh_{ij}^{\lambda_i}}{d\lambda_i} = -\frac{\frac{\partial^2 f}{\partial h \partial \lambda}(h_{ij}^{\lambda_i}, \lambda_i, \omega_i)}{(1-c_j)^2 u''(y_i - p_{ij} - (1-c_j)h_{ij}^{\lambda_i}) + \frac{\partial^2 f}{\partial h^2}(h_{ij}^{\lambda_i}, \lambda_i, \omega_i)} > 0$$

$$\frac{dh_{ij}^{\lambda_i}}{dy_i} = \frac{(1-c_j)u''(y_i - p_{ij} - (1-c_j)h_{ij}^{\lambda_i})}{(1-c_j)^2 u''(y_i - p_{ij} - (1-c_j)h_{ij}^{\lambda_i}) + \frac{\partial^2 f}{\partial h^2}(h_{ij}^{\lambda_i}, \lambda_i, \omega_i)} > 0$$

$$\frac{dh_{ij}^{\lambda_i}}{dc_j} = -\frac{u'(y_i - p_{ij} - (1-c_j)h_{ij}^{\lambda_i}) - (1-c_j)u''(y_i - p_{ij} - (1-c_j)h_{ij}^{\lambda_i})}{(1-c_j)^2 u''(y_i - p_{ij} - (1-c_j)h_{ij}^{*}) + \frac{\partial^2 f}{\partial h^2}(h_{ij}^{\lambda_i}, \lambda_i, \omega_i)} > 0$$

First-stage solution. For simplicity, let us assume the following:

- There are two health states: $\lambda_i \in \{0, \lambda\}$, with $\lambda > 0$.
- For agent i, the bad health state (λ) happens with probability π_i.

Then, agent i chooses plan L if

$$\begin{aligned}(1-\pi_i)[u(y_i - p_{iL}) + f(0,0,\omega_i)] + \pi_i[u(y_i - p_{iL} - (1-c_L)h_{iL}^{\lambda}) + f(h_{iL}^{\lambda}, \lambda, \omega_i)] > \\ (1-\pi_i)[u(y_i - p_{iH}) + f(0,0,\omega_i)] + \pi_i[u(y_i - p_{iH}) + f(h_{iH}^{\lambda}, \lambda, \omega_i)]\end{aligned}$$

[13]Because of this nuance, I refrain from assuming that changes in spending associated with changes in healthcare coverage are socially wasteful when talking about welfare. However, I still refer to it as moral hazard, due to the fact that the term is already well-established in the literature.

Rearranging

$$\left[(1-\pi_i)u(y_i - p_{iL} - (1-c_L)h_{iL}^0) + \pi_i u(y_i - p_{iL} - (1-c_L)h_{iL}^\lambda)\right] - u(y_i - p_{iH}) > \pi_i \left[f(h_{iH}^\lambda, \lambda, \omega_i) - f(h_{iL}^\lambda, \lambda, \omega_i)\right]$$

Given the assumptions, the right-hand side is strictly positive, since

$$\frac{dh_{ij}^{\lambda_i}}{dc_j} > 0 \implies h_{iH}^\lambda > h_{iL}^\lambda$$

This means that, for plan L to be chosen, it must be true that

$$p_{iL} < p_{iH}$$

This means that, for L to be chosen,

$$\tilde{p}_L < \tilde{p}_H \ \& \ \tau y_i < \tilde{p}_H$$

In other words, for L to be chosen by agent i, the mandate cannot be binding for both plans, and can only be binding at most for plan L.

Effect of health on plan choice. Notice that, since $h_{iL}^\lambda > 0$ and given the assumptions, we have that

$$u(y_i - p_{iL}) > u(y_i - p_{iL} - (1-c_L)h_{iL}^\lambda)$$

Hence, if the consumer becomes healthier — i.e., if π_i decreases —, the condition holds for more agents, since the left-hand side of the inequality grows and the right-hand side decreases with everything else held constant. Therefore, plan L tends to have healthier people.

Effect of income on plan choice. For small changes in income, there is ambiguity on how plan choice changes:

- If the mandate binds for plan L, the premium difference decreases for an increase in income, which makes plan H more attractive.
- However, given the concavity assumptions, an increase in income also means that

the value of insurance decreases.

For large increases, the mandate also starts to bind for H, which pushes agents above a certain income threshold to the plan with better coverage, H, since both plans cost them the same.

Welfare when mandate changes. Let us consider what happens with welfare, including both consumer welfare and profits, when the income mandate changes from τ_0 to $\tau_1 > \tau_0$.

First, let us define some objects:

- Let $\tilde{\mathbf{p}}$ be a vector of base premiums for all plans. In equilibrium, these prices depend on the mandate level, so we have that $\tilde{\mathbf{p}} = \tilde{\mathbf{p}}(\tau)$. For simplicity, we denote

$$\tilde{\mathbf{p}}_N := \tilde{\mathbf{p}}(\tau_N)$$

- Let $D = \{j^*(i)\}_i$ denote the assignment of enrollees to plans. D depends on the mandate and the base premiums, so $D = D(\tau, \tilde{\mathbf{p}})$. In equilibrium, we have that

$$D = D(\tau, \tilde{\mathbf{p}}(\tau))$$

  For simplicity, denote

$$D_N := D(\tau_N, \tilde{\mathbf{p}}(\tau_N))$$

Denote by $W(\tau, \mathbf{p}, D)$ the total welfare when the mandate is given by τ, the vector of base premiums is $\mathbf{p}$, and the plan choices are given by D. Then, when the mandate changes from τ_0 to τ_1, we have

$$\Delta W(\tau_1, \tilde{\mathbf{p}}_1, D_1; \tau_0, \tilde{\mathbf{p}}_0, D_0) = W(\tau_1, \tilde{\mathbf{p}}_1, D_1) - W(\tau_0, \tilde{\mathbf{p}}_0, D_0)$$

This can be rewritten as

$$\Delta W(\tau_1, \tilde{\mathbf{p}}_1, D_1; \tau_0, \tilde{\mathbf{p}}_0, D_0) = \underbrace{W(\tau_1, \tilde{\mathbf{p}}_0, D_0) - W(\tau_0, \tilde{\mathbf{p}}_0, D_0)}_{\text{Fiscal channel - } \Delta W^{Fiscal}(\tau_1, \tau_0)} + \underbrace{W(\tau_1, \tilde{\mathbf{p}}_1, D_1) - W(\tau_1, \tilde{\mathbf{p}}_0, D_0)}_{\text{Equilibrium channel - } \Delta W^{Equil}(\tau_1, \tau_0)}$$

The first term on the right-hand side above is the change in welfare due to the change in the mandate only, while holding base premiums and choices fixed: the *fiscal channel*. The second term is the change in welfare from holding fixed the mandate at the new level, but letting base premiums and choices adjust to the new equilibrium level, which I call the *equilibrium channel*.

Changes in consumer surplus will be measured using the compensating variation. In other words, for consumer i, assuming the mandate is changing from τ_A to τ_B, $\Delta CS_i = CV_i$, where CV_i is defined by the value that satisfies

$$V_{ij}(y_i - CV_i \mid \tau_B, \tilde{\mathbf{p}}_B, D_B) = V_{ij}(y_i \mid \tau_A, \tilde{\mathbf{p}}_A, D_A)$$

This means that, in general, $CV_i = CV_i(\tau_B, \tilde{\mathbf{p}}_B, D_B; \tau_A, \tilde{\mathbf{p}}_A, D_A)$ Denote by $\pi_{ij}(\tau, \tilde{\mathbf{p}})$ the profit generated by consumer i choosing plan j under mandate τ and base premiums $\mathbf{p}$. I measure changes in welfare by the sum of all the individual compensating variations and the sum of the differences in total profits. In other words,

$$\Delta W(\tau_1, \tilde{\mathbf{p}}_1, D_1; \tau_0, \tilde{\mathbf{p}}_0, D_0) = \sum_i CV_i(\tau_1, \tilde{\mathbf{p}}_1, D_1; \tau_0, \tilde{\mathbf{p}}_0, D_0) + \sum_i \left[\pi_{iD_1^i}(\tau_1, \tilde{\mathbf{p}}_1) - \pi_{iD_0^i}(\tau_0, \tilde{\mathbf{p}}_0)\right]$$

where D_N^i is the i-th component of D_N — thus, the choice of the i-th agent according to D_N. We can then write the fiscal and equilibrium channels in terms of these objects:

$$\begin{aligned}
\Delta W^{Fiscal}(\tau_1, \tau_0) &= \sum_i \left[CV_i(\tau_1, \tilde{\mathbf{p}}_0, D_0; \tau_0, \tilde{\mathbf{p}}_0, D_0) + \pi_{iD_0^i}(\tau_1, \tilde{\mathbf{p}}_0) - \pi_{iD_0^i}(\tau_0, \tilde{\mathbf{p}}_0)\right] \\
\Delta W^{Equil}(\tau_1, \tau_0) &= \sum_i \left[CV_i(\tau_1, \tilde{\mathbf{p}}_1, D_1; \tau_1, \tilde{\mathbf{p}}_0, D_0) + \pi_{iD_1^i}(\tau_1, \tilde{\mathbf{p}}_1) - \pi_{iD_0^i}(\tau_1, \tilde{\mathbf{p}}_0)\right]
\end{aligned}$$

Focusing on the fiscal channel, remember that $p_{ij}(\tau) = \max(\tilde{p}_j, \tau y_i)$. Then, when the mandate changes from τ_0 to τ_1, there are three cases to consider from an individual point of view:

- $\tau_1 y_i \leq \tilde{p}_j$ — In this case, the agent is never bound by the mandate. Since we fix premiums and choices in the fiscal channel, nothing changes for these consumers, so

$$CV_i = \pi_{iD_0^i}(\tau_1, \tilde{\mathbf{p}}_0) - \pi_{iD_0^i}(\tau_0, \tilde{\mathbf{p}}_0) = 0$$

- $\tau_0 y_i \geq \tilde{p}_j$ — When this is true, the agent is already bound by the mandate to begin with. Then, the extra premium she pays is given by $\Delta p_{ij} = (\tau_1 - \tau_0) y_i$. Hence, the compensating variation satisfies

  $$V_{ij}(y_i - \tau_1 y_i - CV_i) = V_{ij}(y_i - \tau_0 y_i)$$

  Hence,
  $$CV_i = -(\tau_1 - \tau_0) y_i$$

  For the insurer, the gain in profits is given by

  $$\pi_{iD_0^i}(\tau_1, \tilde{\mathbf{p}}_0) - \pi_{iD_0^i}(\tau_0, \tilde{\mathbf{p}}_0) = (\tau_1 - \tau_0) y_i$$

  This means that the fiscal channel is essentially a transfer from consumers to insurers that is net neutral.

- $\tau_0 y_i < \tilde{p}_j < \tau_1 y_i$ — These are the people who were not bound initially but, keeping base premiums fixed, would become bound under the new mandate. In this case, $\Delta p_{ij} = \tau_1 y_i - \tilde{p}_j$ and the same logic from the previous case applies.

This means that only people who, under the new mandate but the initial prices and choices, contribute to the fiscal channel. Hence, we can write

$$\Delta W^{Fiscal}(\tau_1, \tau_0) = \sum_{i \in \tau_1\text{-bound}} \left[CV_i^0(\Delta p_i) + \Delta p_i \right]$$

where Δp_i is the difference in premium for agent i fixing base premiums and choices at the initial state and $CV_i^0(\Delta p_i)$ is the compensating variation for the same consumer when premiums change by Δp_i and the other variables are held at their initial values. Then, $CV_i^0(\Delta p_i) = -\Delta p_i$ and

$$\Delta W^{Fiscal}(\tau_1, \tau_0) = 0$$

In the empirical exercise later, I also consider the changes with different distributional weights $\varepsilon > 0$. In this case, I weight the compensating variation by $w_i = (y_i/\bar{y})^{-\varepsilon}$, following the inequality-aversion framework of Atkinson (1970) and Saez and Stantcheva (2016):[14]

$$\Delta W^{Fiscal}(\tau_1, \tau_0) = \sum_{i \in \tau_1\text{-bound}} \left[w_i CV_i^0(\Delta p_i) + \Delta p_i \right]$$

where $\bar{y}$ is the average income. In this case, we may have $|CV_i^0(\Delta p_i)| \neq \Delta p_i$, with $|CV_i^0(\Delta p_i)| > \Delta p_i$ for poorer individuals and $|CV_i^0(\Delta p_i)| < \Delta p_i$ for richer ones.

The equilibrium channel is significantly more complex since it involves equilibrium responses. But before we turn to a qualitative analysis of the equilibrium channel, we need some extra assumptions. Denote by k_{ij} the expected cost generated by consumer i for plan j, while $\bar{K}_j$ denotes the average expected cost per enrollee for plan j. Similarly, let r_{ij} be the mandate wedge for agent i who acquires plan j (i.e., the difference between the premium effectively paid and the base premium, positive if the mandate binds for agent i) with $\bar{R}_j$ denoting the average mandate wedge per enrollee in plan j. Then, I assume that the equilibrium base premium for plan j, $\tilde{p}_j$, is such that

$$\frac{d\tilde{p}_j}{d\bar{K}_j} > 0 \ \ , \frac{d\tilde{p}_j}{d\bar{R}_j} \leq 0 \ \ \forall j$$

In other words, with the first assumption, the equilibrium base premiums increase (decrease) when average costs increase (decrease). The second assumption reflects the fact that, everything else constant, an above-normal revenue from mandate wedges generated by enrollees puts downward pressure in base premiums due to competitive pressure.

The equilibrium channel. We can understand the equilibrium channel as working through two different mechanisms when the mandate goes from τ_0 to $\tau_1 > \tau_0$. The first is the fact that people in a particular plan who are bound by the mandate have now to pay a higher price for the plan. This increases how much the insurer gets in mandate wedges which, per the assumptions above, may generate competitive pressure and decrease base premiums for the plan. The second channel is related to the change in the enrollee composition caused by changes in choices. Some people who were not bound by the mandate will become bound, together with people who were bound for plan L but not for plan H, will see the difference in the effective prices they pay for each plan shrink. This will increase the incentives of those of this groups of people who were acquiring plan L to switch to plan

[14] I discuss the empirical implementation and the choice of ε values reported in subsection 6.2.

H. This changes the enrollee composition of each plan, which can push base premiums up or down depending on whether the switchers are more or less costly than the average of each plan.

To fix ideas, especially in what concerns the effect of switchers, it is useful to consider a slightly simpler scenario that can be graphed. Consider that we only have plan H, while plan L is the same as uninsurance, so we just have one plan to think about. Assume the mandate is rising, again, from τ_0 to τ_1. Let us first consider the case where income and spending are negatively correlated, so that richer people are less costly to insure. Figure 1a shows the initial equilibrium.

Figure 1: Selection channel when income-spending correlation is negative

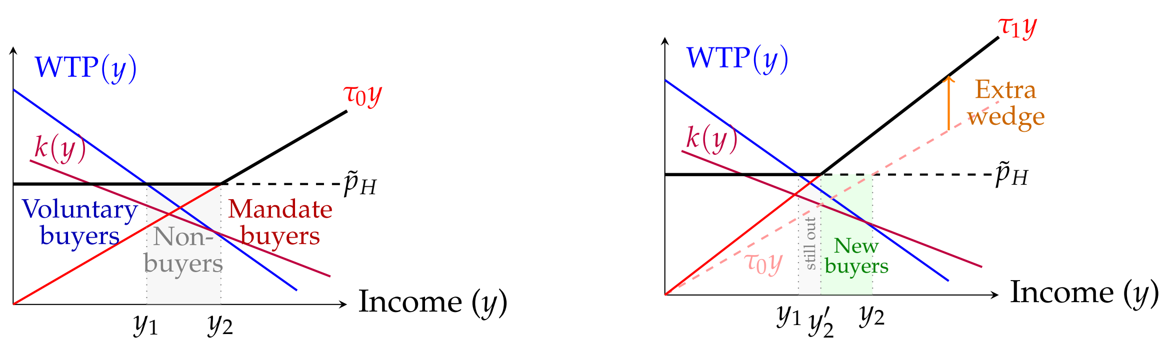

(a) Baseline: three enrollment regions

(b) Mandate increase: $\tau_0 \to \tau_1$

Notes: The figure represents the situation of a market when the income-spending correlation is negative, i.e., medical spending falls with income and richer people are less costly to insure than poorer people (because, for instance, income and health are positively related). Plan H enrollment under income-based mandate. WTP for plan H is decreasing in income due to declining risk aversion and decreasing cost curve; the cost curve $k(y)$ is decreasing under positive income-health correlation. Panel (a): under τ_0, voluntary buyers (low-income, high-WTP, high-cost) and mandate buyers (high-income, bound, low-cost) enroll, while a middle group does not. Panel (b): when $\tau_1 > \tau_0$, some non-buyers in (y_2', y_2) enter (green); already-bound enrollees pay a higher wedge (orange). Whether the new buyers' costs are above or below the pool average determines the sign of the cost composition effect. The base premium $\tilde{p}_H$ may also shift downward in equilibrium due to competitive pressure from extra mandate revenue (not shown).

The graph shows the WTP curve as a function of income, as well as the mandate. The solid black line shows the final premium each enrollee will pay based on income: if the base premium is higher than the mandate curve, people pay the base premium; otherwise, people pay what is prescribed by the mandate, generating the non-linear price in the figure. For people unbound by the mandate, who I call voluntary buyers, they acquire plan H as long as the WTP curve is above the base premium. When the mandate starts to bind people, since there is no price differential for them anymore and they value insurance, they acquire plan H; I call these mandate buyers. The people with incomes between y_1 and y_2 are unbound by the mandate, but their WTP is less than the base premium: these are the non-buyers. Notice that we have two groups of buyers then: one that is sicker and buys insurance because they significant value insurance (the voluntary buyers) and richer, healthier people, who only buy plan H because they have to pay the mandate regardless of their choice and they value insurance to some extent.

Figure 1b shows what happens when we move the mandate to τ_1. First, notice that, for mandate buyers, insurers gain an extra wedge indicated in orange, which can generate a downward pressure on base premiums. For clarity, I do not present the base premium movements in the figure, but notice that this might induce poorer (and more costly, in this case) people to choose plan H, dampening this effect.[15]

Regarding switchers, now the mandate starts to bind people with income above $y_2' < y_2$, which was the income at which the mandated started to bind under τ_0. Now, people with incomes between $y'2$ and y_2, despite having WTP lower than the base premium, have to pay according to the mandate regardless of their choice. Since they value insurance, they switch to plan H, changing the plan composition. This will, in turn, affect the base premium $\tilde{p}_H$, but the direction in which this effect will push is ambiguous: it depends on whether these switchers are more or less costly than the average enrollee of plan H. Notice that this is unclear, since there are two groups of buyers with different sizes and cost profile.

The case of positive correlation between income and health spending, where richer people are more costly to insure than poorer people, is shown in Figure 2. Notice that, despite the costs being increasing in income, I still assume the WTP curve is decreasing in income. This happens because WTP depends not only on health, but also on how the insurance value for individuals varies with income. If we consider, for instance, a utility of consumption with a log form, the same risk (in absolute terms) will be less valued by rich people than by poor people. Hence, the WTP increases with income only if this effect is dominated by the increase in health risk, which is more likely at higher incomes, where the mandate is more likely to bind. In any case, in Appendix C, I show the analysis for the case where WTP is increasing in income.

When income and spending are positively correlated, the analysis is qualitatively unam-

[15]In the graph, this effect would amount to an unambiguous downward shift of the base premium, with the rest of the analysis, particularly concerning switchers, remaining the same.

Figure 2: Mandate increase under positive income-spending correlation

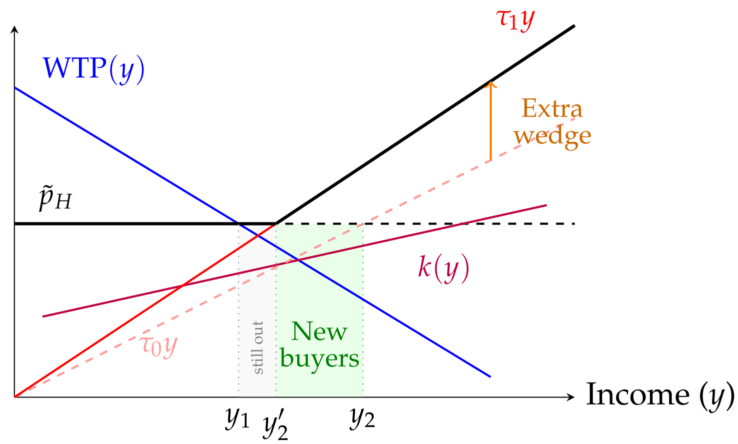

Notes: The figure shows what happens when the mandate increases under positive income-spending correlation (richer people are more costly to insure). WTP is still decreasing (risk aversion dominates), but the cost curve $k(y)$ is now increasing. Voluntary buyers on the left are low-cost (poor but healthy), while mandate buyers on the right are high-cost (rich but sick). New buyers in (y_2', y_2) have intermediate-to-high costs, making it more likely that they worsen the risk pool relative to the incumbent average.

biguous. Now, the voluntary buyers are the inexpensive ones, whereas the mandate buyers are the costly ones to insure. The same effect of the increase in the mandate wedge exists, which will tend to push the base premium downwards in equilibrium. However, now the effect of switchers goes in the same direction: people with lower income, who in this case are definitely cheaper to insure, join plan H, decreasing the average cost per enrollee and, in equilibrium, pushing base premiums down.

Another interesting point concerns the strength of the correlation between income and medical spending and its relationship to the effect generated by the switchers via enrollee composition. This is shown in Figure 3.

In Figure 3, we again look at the situation where health and spending are negatively correlated, but now we include two potential cost curves: one where the income-spending correlation is weak, with an essentially flat curve, and another with a stronger correlation. The figure also indicates the average cost per enrollee, which is assumed to be the same for both cases. With a weak correlation, notice that people who may switch plans when the mandate varies are all very similar in terms of costs, and also very similar to the initial enrollees (and the average cost). Hence, they have a limited scope to affect base premiums. When the correlation is stronger, the switchers may exhibit larger differences with respect to the average cost per enrollee, which generates larger effects on equilibrium premiums.

Figure 3: Cost composition effect of new buyers and magnitude of income-spending correlation

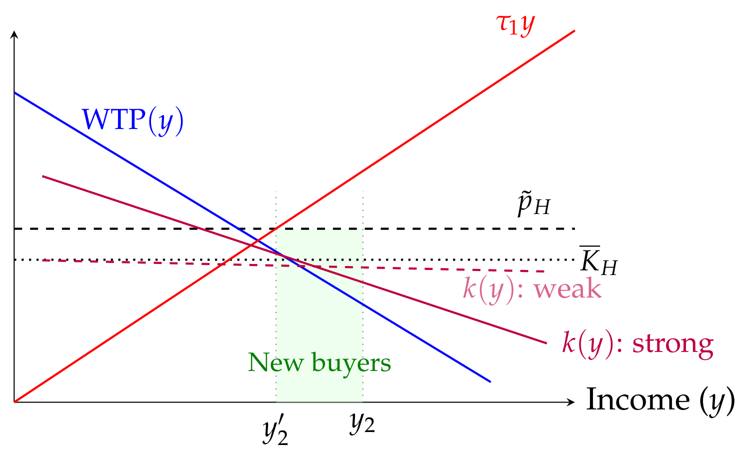

Notes: The graph shows the cost composition effect of new buyers under different strengths of the income-spending correlation. Under a strong negative income-spending correlation (solid red), new buyers in (y_2', y_2) have costs below $\bar{K}_H$, improving the pool. Under a weak correlation (dashed red), new buyers have costs near $\bar{K}_H$, with a muted effect regardless of direction.

In light of this analysis, note that, for the equilibrium channel to be potentially relevant, two conditions must typically be met. First, there must be a non-negligible pool of unbound enrollees whose plan choices could change under the mandate increase. Second, the *residual* correlation between income and costs (i.e., *after* accounting for any risk-rating) must be sufficiently strong. Risk-rating absorbs cost variation driven by observable demographics, flattening the residual cost curve. In markets where risk-rating is pervasive (as in the Chilean ISAPRE system), even a strong unconditional correlation between income and costs does not guarantee a strong equilibrium channel.

4 Data

In this section, I describe my main data sources. Moreover, I make explicit the assumptions I make to be able to combine all the data I have, as well as those I make to allow the data to be fully compatible with the empirical model described in the next section.

Overview of the Main Datasets. The main dataset comes from Chile's private insurance regulatory agency, *Superintendencia de Salud* (Superintendencia de Salud, 2012). This administrative dataset contains claims data for all privately insured people in Chile. Despite being anonymized, these data can be complemented with other data from the same agency on the basic characteristics of each enrollee — age, gender, district of residence, number of dependents, and income[16] — and with data on plan characteristics — base premium, age-gender risk adjustment factors, general outpatient and inpatient coinsurance rates — for plans that are held by at least one person.[17]

I have access to the data for 2011 and 2012. Unlike in markets with a limited enrollment period (as in the U.S.), ISAPRE contracts are signed throughout the year, with each contract running for one year from the month of signing. To capture a consistent measure of annual healthcare utilization, I aggregate each enrollee's medical claims over the 12 months following her contract start month, regardless of whether that window begins in 2011 or January 2012. Each enrollee thus contributes exactly one year of data, with the calendar window varying across enrollees according to their contract start date. I attribute to each enrollee the plan she held as of January 2012; if she changes plans during her contract year, her aggregated claims still include all observed expenditures under both the prior and subsequent plans.[18]

Additionally, for some auxiliary graphs regarding the general characteristics of the Chilean market discussed in section 2, for both the public option and the private market, and considering the full population, I use data from the 2013 Chile National Socioeconomic Characterization Survey (CASEN) (Instituto Nacional de Estadísticas, 2013).

Adjustments to Plans Data. The Chilean private health insurance market has a significant number of formally distinct plans. However, many of these plans are essentially the

[16]Technically, what is reported is the income relevant for the calculation of the mandate, which is capped at a statutory ceiling (*tope imponible*). The unit of observation throughout the analysis is the *cotizante* (policyholder, i.e., the contracting party who pays the premium), not total covered lives. Each *cotizante* represents one contract and one decision unit; a household with dependents (*cargas*) covered under a single contract appears linked to one *cotizante* in the data. Although a non-negligible share of policyholders in my sample seem to have their incomes censored due to this restriction, this is unlikely to affect my results since the cap is relatively large compared to the typical income and because the vast majority if not all of these enrollees are inframarginal in the counterfactual exercises.

[17]These data were obtained through a request for public information submitted to the Superintendencia de Salud under Chile's access-to-information law (Ley 20.285); the agency released the data used here and withheld part of the original request as protected under Ley 19.628, none of which is used in the analysis. The 2011-2012 records are no longer hosted in the open version on the agency's public portal but remain available to any researcher who submits an equivalent request. Access is therefore subject to the agency's response to such a request.

[18]If an enrollee leaves ISAPRE for FONASA (the public option) during the contract year, post-departure claims are not observed in the ISAPRE data, which understates her total annual spending. Given the strong segmentation between the public and private systems, this is rare and unlikely to introduce meaningful measurement error.

same in content. Then, I adopt an approach similar to Starc and Town (2020) and group the plans by their characteristics. More specifically, I group them based on the main characteristics used when they are marketed by insurers: the most common coinsurance rates for outpatient and inpatient procedures, and the identity of the insurer company. This yields 20 plan groups across the five major open insurers. In order to define the other characteristics for each plan group, I compute enrollment-weighted means for continuous variables — the base premium and demographic-specific effective coinsurance rates — and take modal values for discrete characteristics such as the risk adjustment factor table.[19]

The base premium for each plan group is not necessarily the price enrollees pay. Rather, it is a plan-level price that is then adjusted by a demographic-specific risk factor, based on age and gender. Then, the premium for individual i in plan j is

$$p_{ij} = \max\{\tilde{p}_{ij},\ \tau y_i\}$$

where $\tilde{p}_{ij} = \bar{p}_j \cdot r_j(X_i)$ is the risk-adjusted premium, $\bar{p}_j$ is the plan group base premium, $r_j(X_i)$ is the age-gender risk factor, $\tau = 0.07$ is the mandatory contribution rate, and y_i is individual i's income. When $\tau y_i > \tilde{p}_{ij}$, the enrollee is bound by the mandate: the insurer receives the full mandatory contribution, which exceeds the actuarially determined premium.[20]

Another characteristic of private plans in the Chilean market is that, although they have no deductibles, they do have limits on the amount insurers cover. One of these limits is an annual cap on the amount of expenses covered by the given plan, which is generally high and rarely binding. The other, more relevant limit is a cap on what the insurer covers per procedure, each time it is performed. This last cap on the amount of expenses covered

[19] Within each plan group, the modal factor table accounts for approximately 68% of enrollment on average (enrollment-weighted). Since factor tables within the same insurer tend to produce similar risk adjustments, this is unlikely to introduce meaningful measurement error.

[20] In principle, when the mandatory contribution exceeds the plan price, the difference — called *excedente* — is deposited into an individual account that can be used for copayments and uncovered health services. However, ISAPREs routinely adjust plan prices to absorb the contribution through the annual *adecuación* process: when excedentes exceed 10% of the legal contribution, ISAPREs are required to offer an alternative plan whose price better matches the contribution; as of December 2012, while 48.2% of ISAPRE affiliates had an active excedente account from past accumulation, the average per-capita balance was approximately \$160 USD (Superintendencia de Salud, 2013). In the full cotizantes data for 2012, only approximately 1.0% of enrollees have a total premium below their 7% contribution, a share that is even lower for my final estimating sample. Therefore, I abstract from *excedentes* in the model.

by the insurer makes the actuarial value of plans differ from the typical coinsurance rates used to market them. Then, for each (original) plan and each demographic group I use in the estimation, I calculate the effective coinsurance rate as one minus the ratio of total insurer-covered expenses to total medical expenses among enrollees in that cell. When aggregating to plan groups, I take the enrollment-weighted mean of these rates across constituent plans.

Population of Interest. The empirical analysis focuses on a specific subset of the privately insured population, defined by a set of scope restrictions. First, I focus on formal-sector workers and retirees who are required to purchase insurance (*Dependiente* or *Pensionado*), since these are the enrollees subject to the income-based contribution mandate. Second, I restrict the sample to people living in Santiago's metropolitan area, where the healthcare market structure is most uniform; outside the capital, provider networks are narrower and the public sector plays a larger role.[21] Third, I focus on single-person contracts (no dependents, i.e., *cotizantes* with no associated *cargas*), for whom the choice of health plan can be cleanly characterized as a single-agent decision.[22] I also exclude the so-called closed insurers and their enrollees, which serve specific employer groups and constitute a small share of the market, and two other open insurers with negligible enrollment. This leaves the five major open insurers, which together cover approximately 95% of the private market. Finally, I restrict to people who are 18 years old or above.[23]

Estimation Sample. On top of the scope restrictions described above, I apply two further types of restrictions. The first set concerns data quality and ensures that the model's inputs are well-defined: I require that enrollees have valid plan information (matched to the plans dataset), are in plans with non-fixed copayments, and have reported income above the legal minimum wage (which also serves to exclude observations with missing

[21]The Santiago metropolitan area is home to approximately 40% of the Chilean population and a substantially larger share of ISAPRE enrollment. Studying Santiago in isolation is a scope choice, not a sample selection in the welfare-bias sense: I am studying the policy's effect on a well-defined population rather than excluding observations on the basis of unobservables.

[22]Family contracts involve joint decisions across household members and complicate the model in ways that lie outside the methodological scope of this paper. An approach similar to Marone and Sabety (2022) could in principle handle family contracts, but at the cost of additional assumptions about intra-household decision-making. The mechanism underlying my main result — the dominance of the fiscal channel under income-based pricing — generalizes to family contracts, since the mandate floor is income-based and premiums are risk-rated regardless of household composition.

[23]The few people below 18 years old left in the data after the previous sample restrictions are children with substantial income reported, which are likely to be mistakes in the reported data.

or unreliable income data). I also drop plan groups with less than 1% market share, both for data-quality reasons (small plans tend to have insufficient enrollment in some demographic cells to compute reliable cell-specific coinsurance rates) and because the estimation framework requires adequate identification of plan-level parameters; this reduces the number of plan groups from 24 to 15.

The second set of restrictions are required by the structural estimation framework. First, I drop enrollees whose implied consumption under their chosen plan is negative ($z_{ij} = y_i - p_{ij} - (1 - c_{ij})h_i < 0$), which removes 778 individuals. Second, the estimation requires that, on each enrollee's chosen plan, positive health spending can only arise when the health shock is positive ($h^*_{ij} = 0 \iff \lambda_i = 0$). This regularity condition, $y_i - p_{ij} > 1 - c_{ij}$, is violated by only 19 enrollees, who are also dropped.[24]

Appendix Table A.1 reports the number of observations at each stage of sample construction. The initial pool consists of approximately 1.53 million formal-sector contributors in January 2012. Of these, 91.0% are workers required to have insurance (*Dependiente* or *Pensionado*); 45.3% live in the Santiago metropolitan area; and 25.8% are in single-person contracts.[25] The remaining scope restrictions (open ISAPREs, age 18+) and the data-quality restrictions (valid plan info, non-fixed copay, income above minimum wage, plans with $\geq$ 1% share) cumulatively retain 13.3% of the original pool. After the two estimation-specific filters described above, the final sample consists of 202,662 enrollees in 15 plan groups.

Scope and External Validity. The sample restrictions described above define a specific subpopulation: single-person Santiago contracts in open ISAPREs. Within this subpopulation, the empirical analysis can cleanly identify the welfare effects of marginal changes in the contribution rate and decompose them into fiscal and equilibrium channels. The substantive welfare numbers reported in this paper characterize the mandate's effect on this subpopulation. Extrapolating to the full ISAPRE universe, including family contracts,

[24]As explained in subsection 5.2, this condition ensures that zero observed spending on the chosen plan corresponds to a zero health shock, which simplifies the inversion of observed spending to recover the latent health shock. In principle, one could accommodate these observations, but this would substantially complicate the estimation. Since the share of observations for which this condition does not hold is negligible, I opt for dropping these observations.

[25]The geographic restriction to Santiago is the single largest reduction in sample size, but it reflects a scope choice rather than sample selection: healthcare market structure differs substantially outside the capital, with narrower provider networks and greater public-sector dominance, making within-Santiago analysis methodologically cleaner.

non-Santiago residents, and below-minimum-wage workers, would require additional assumptions about how the mechanisms identified here scale to those settings. The methodological contribution of the paper, however, generalizes immediately: the decomposition framework applies to any contribution-floor mandate, and the conditions under which the fiscal channel dominates are present whenever income-based mandates interact with risk-rated premiums.

Descriptive Statistics. I present the descriptive statistics of the final sample in Table 1. The final sample consists of 202,662 people across 15 plan groups. When comparing results with other studies from the literature, it is important to notice that in my setting, incomes and expenses tend to be significantly lower than those for developed countries: the average income is \$24.48 thousand per year (in 2012 U.S. Dollars) and average medical expenses amount to \$1.33 thousand per year. In addition, there is a large degree of inequality, which can be surmised by the wide interquartile range of \$18.77 thousand. Finally, the sample is somewhat young, with the group of people between 30 and 44 years old comprising 44% of the sample and the combined share below 45 years old amounting to 71% of the sample.

Approximately 50% of enrollees are bound by the mandate on their chosen plan, meaning their 7% mandatory contribution exceeds the risk-adjusted premium. These are the individuals for whom a mandate increase directly affects the fiscal channel.

5 Empirical Strategy

5.1 Empirical Model

In this section, following Cardon and Hendel (2001), Einav et al. (2013), and Marone and Sabety (2022), I present a parameterization of the two-stage model of health insurance choice and health care utilization from the theoretical framework section, along with some extensions to better reflect the Chilean empirical setting. This model is then used for estimation and as the main tool for counterfactual analysis.

Table 1: Descriptive Statistics

	Mean	SD	Q1	Q3
Enrollees				
Age	39.36	13.82	29	47
Share Age: 18–29	0.27			
Share Age: 30–44	0.44			
Share Age: 45–64	0.21			
Share Age: 65+	0.07			
Share Women	0.45			
Income (USD 1000/year)	24.48	11.16	15.12	33.89
Health Exp (USD 1000/year)	1.33	3.54	0.08	1.03
Share Zero Spending	0.14			
Share Bound by Mandate	0.50			
N	202,662			
Plan Groups				
Effective Coinsurance	0.33	0.07	0.27	0.39
Base Premium (USD 1000/year)	1.07	0.31	0.81	1.38
N	15			

Notes: The main columns report statistics for the estimation sample. Plan groups are formed by aggregating plans within each insurer that share the same basic outpatient and inpatient coinsurance rates. Effective coinsurance is the enrollment-weighted average across demographic cells within each plan group. Base premium is the enrollment-weighted mean across constituent plans within each group. "Bound by mandate" indicates enrollees whose 7% mandatory contribution exceeds their risk-adjusted premium on the chosen plan.

Individual First Stage Utility. In the first stage, before the health shock is realized, agents choose their plan $j \in \{1, 2, ..., J\}$ according to[26]

$$j = \underset{j'}{\text{argmax}}\ V_{ij'}$$

where

$$V_{ij'} = \mathbb{E}_{\lambda_i}[U_{ij'}] + \xi_{j'} + \varepsilon_{ij'}$$

where U_{ij} is the (optimal, conditional on the λ_i realization) second stage utility of agent i when choosing plan j. Compared to the theoretical framework, two elements were added here. First, the ξ_j plan fixed effect, which accounts for characteristics of plans not explicitly included (e.g., marketing). Second, I include an idiosyncratic, i.i.d. preference shock ε_{ij}. This shock is assumed to have an extreme value type I distribution with unitary variance. Hence, the probability that agent i chooses plan j is given by

$$s_{ij} = \frac{e^{\mathbb{E}_{\lambda_i}[U_{ij}] + \xi_j}}{\sum_{j'} e^{\mathbb{E}_{\lambda_i}[U_{ij'}] + \xi_{j'}}}$$

Importantly, recall that, when choosing plans, agents have private information about their ω_i. I assume the variable is distributed among the population as

$$\log(\omega_i) \sim \mathcal{N}(\overline{\omega}, \sigma_\omega^2)$$

where

$$\overline{\omega} = \omega_0 + \alpha_\omega \log(y_i / y_{ref}) + \delta_\omega \mathbb{1}\{y_i \text{ censored}\}$$

where y_i is agent i's income and y_{ref} is a reference income that is fixed throughout the empirical exercises.[27] Since the contribution income is capped and a number of observations is censored at this income,[28] for agents with income at or above the cap, I consider a shift

[26]The choice set includes ISAPRE plans only. As discussed previously, movement between the ISAPRE and FONASA systems is limited and the two markets are largely segmented. Moreover, the counterfactual exercises involve a 1 percentage point mandate change, which is unlikely to be large enough to affect the extensive margin between systems.

[27]The reference income is used solely for numerical reasons when estimating the model. Notice that it does not affect anything at a conceptual level.

[28]I do observe some incomes censored at the cap, while others are reported even above it.

in $\overline{\omega}$ for the observations I observed being at the capped income.[29]

Individual Second Stage Utility. In the second stage, after plan j is chosen and the health shock is realized, agent i chooses consumption z_{ij} and health spending h_{ij}. For $\lambda_i > 0$, the second stage utility is parameterized as

$$U_{ij} = \log(z_{ij}) + (h_{ij} - \lambda_i) - \frac{1}{2\omega_i \lambda_i}(h_{ij} - \lambda_i)^2$$

$$\text{s.t. } z_{ij} = y_i - p_{ij} - (1 - c_{ij})h_{ij}$$
$$z_{ij} \geq z_{min} \, , h_{ij} \geq 0$$

where $z_{min} > 0$ is a minimum consumption level, c_{ij} is the coverage of plan j for individual i, and p_{ij} is the premium of plan j to agent i; assuming the mandate is such that enrollees must pay in premiums at least τy_i,

$$p_{ij} = \max\{\tilde{p}_{ij}, \tau y_i\}$$

where $\tilde{p}_{ij}$ is the base premium of plan j, with the relevant risk factor based on age and gender applied to it. I also assume that, when $\lambda_i = 0$,

$$U_{ij} = \log(z_{ij})$$

under the same constraints.

For the part of utility associated with health spending, my specification follows Einav et al. (2013)'s alternative specification with multiplicative moral hazard. I depart from this specification in the utility of consumption: whereas Einav et al. (2013) and Marone and Sabety (2022) consider the utility function quasi-linear in consumption, I assume it is strictly concave. This naturally embeds the risk aversion that makes people value insur-

[29] In practice, I identify enrollees with censored income using the clear spike in the income distribution I observe and considering a very small window around it. For people with censored income, this will estimate an average for $\overline{\omega}$. Since the income cap is relatively large, people with censored income will most likely be inframarginal in the counterfactual exercise, and hence this will likely not be of any significant consequence for the results. The only point of potential relevance is in the consumption utility. However, given the relatively high value of the cap, this will have a very small effect on the results, since the log function is relatively flat around the cap.

ance.

For $\lambda_i > 0$, we can solve for the agent's optimal health spending assuming the constraints on consumption and health spending are not binding; let us call it the *unconstrained optimal spending*, h_{ij}^{unc}. The first order condition can be written as

$$h_{ij}^{unc} = \lambda_i \left[1 + \omega_i \left(1 - \frac{1 - c_{ij}}{z_{ij}^{unc}}\right)\right]$$

where $z_{ij}^{unc} = y_i - p_{ij} - (1 - c_{ij})h_{ij}^{unc}$. Notice that, when $1 - c_{ij}$, if we have that $y_i - p_{ij} > 1 - c_{ij}$, then $h_{ij}^{unc} = 0 \iff \lambda_i = 0$.[30] Solving for h_{ij}^{unc}, we have:[31]

$$h_{ij}^{unc} = \frac{y_i - p_{ij} + \lambda_i(1 + \omega_i)(1 - c_{ij})}{2(1 - c_{ij})} - \frac{\sqrt{(y_i - p_{ij} - \lambda_i(1 + \omega_i)(1 - c_{ij}))^2 + 4\omega_i\lambda_i(1 - c_{ij})^2}}{2(1 - c_{ij})}$$

Also, for $c_{ij} < 1$, define $h_{ij}^{max} \in \mathbb{R} \cup \{\infty\}$ as

$$h_{ij}^{max} = \frac{y_i - p_{ij} - z_{min}}{1 - c_{ij}}$$

If $c_{ij} = 1$, let $h_{ij}^{max} = \infty$. Then, the optimal spending h_{ij}^* is given by

$$h_{ij}^* = \max\{\min\{h_{ij}^{unc}, h_{ij}^{max}\}, 0\}$$

It is also important to notice that, compared to Einav et al. (2013), the interpretation of ω_i is slightly more complex, but encompasses a richer scenario.[32] Although the first order condition only defines the unconstrained optimal spending implicitly, it helps to understand the role of ω_i. If we consider only the non-linear part of the utility of health spending, the unconstrained optimal choice would be to set $h = \lambda$; hence, deviations from this come from the utility gains of increasing h from the linear part of health spending utility *versus*

[30]This is particularly relevant in the estimation for the observed enrollee-plan pairs.

[31]When solving for h_{ij}^{unc}, there are two potential solutions: the one I present and another in which the second term is positive. However, this second solution would imply a negative consumption.

[32]In Einav et al. (2013) multiplicative moral hazard model, the optimal spending is given by $h_{ij}^* = \lambda_i(1 + \omega_i c_{ij})$, and ω_i can be interpreted as the increase in spending, in percentage terms, from changing the coverage from 0 to 1.

the loss of utility from consumption necessary to finance this. In the Einav et al. (2013) model, since the utility of consumption is linear, this becomes a simple monetary comparison: with coverage c_{ij}, an agent loses $1 - c_{ij}$ in consumption utility to get a unit increase in utility coming from the linear part of health utility, with ω_i controlling the importance of this difference, or price sensitivity. In my model, the cost of financing a unit increase in linear health utility must be weighted by the marginal utility, since consumption utility is no longer linear. However, notice that ω_i still controls the degree to which this difference matters for the spending choice, or the price sensitivity of the agent, hence determining the level of ex-post moral hazard. However, the model allows for richer patterns that depend, including that people (especially poorer people) spend less than λ_i.

Health Shock Parameterization. To accommodate both the typical mass at zero and the possibility of heavy tails, I assume λ_i follows a three-part model, similar in spirit to Karlsson et al. (2024). Specifically, I assume λ_i follows

$$\lambda_i \sim \begin{cases} 0 & \text{with probability } \kappa(X_i) \\ G^+(\cdot \mid X_i, y_i) & \text{with probability } 1 - \kappa(X_i), \end{cases}$$

where X_i are individual characteristics (age and gender groups) and $\kappa(X_i) \in [0, 1]$ is the demographic-specific probability of zero health spending. Conditional on $\lambda_i > 0$, the density is a spliced log-normal-Pareto distribution:

$$g^+(\lambda \mid X_i, y_i) = \begin{cases} \phi_{\mathrm{LN}}(\lambda;\, \mu(X_i, y_i),\, \sigma(X_i)) & \text{if } \lambda \in (0,\, \bar{\lambda}\,] \\ S_{\mathrm{LN}}(\bar{\lambda};\, \mu(X_i, y_i),\, \sigma(X_i))\, \dfrac{\theta\, \bar{\lambda}^{\,\theta}}{\lambda^{\,\theta+1}} & \text{if } \lambda > \bar{\lambda}, \end{cases}$$

where $\phi_{\mathrm{LN}}(\cdot;\, \mu, \sigma)$ denotes the log-normal density and $S_{\mathrm{LN}}(\bar{\lambda};\, \mu, \sigma) = 1 - \Phi\left(\frac{\log \bar{\lambda} - \mu}{\sigma}\right)$ is the log-normal survival function evaluated at the splice point.[33] The splice point $\bar{\lambda}$ and the Pareto shape parameter θ are common across demographic groups, while $\mu(X_i, y_i)$, $\sigma(X_i)$, and $\kappa(X_i)$ are group-specific. μ also includes a linear dependence on log income (with the shift for the maximum contribution income) in the same way as $\overline{\omega}$:

$$\mu(X_i, y_i) = \mu_0(X_i) + \alpha_\mu \log(y_i / y_{ref}) + \delta_\mu \mathbb{1}\{y_i \text{ censored}\}$$

[33]Notice the density integrates to one by construction.

5.2 Estimation and Identification

The vector of parameters to be estimated is denoted by

$$\boldsymbol{\Theta} = [\boldsymbol{\xi}', \omega_0, \alpha_\omega, \delta_\omega, \sigma_\omega, \boldsymbol{\mu}_0', \alpha_\mu, \delta_\mu, \boldsymbol{\sigma}', \boldsymbol{\kappa}', \theta]'$$

where bold letters indicate vectors, which happens when the parameter varies with age and gender or for the plan fixed effects. I consider four age bins — 18-29 years old, 30-44 years old, 45-64 years old, and 65 years old and above — which are interacted with gender, resulting in eight demographic groups. For the estimation, I fix $z_{min} = 0.001$, and I set the splice point $\bar{\lambda}$ at the 85^{th} percentile of the observed positive spending distribution.

Likelihood. The model is estimated via Markov Chain Monte Carlo (MCMC).[34] The key challenge is that the moral hazard parameter ω_i is private information of the agent: it enters both the plan choice (through expected utility) and the spending decision (through the first order condition), but is not directly observed. I handle this via data augmentation, treating ω_i as a latent variable that is drawn jointly with the population parameters.[35]

I adopt a non-centered parameterization for ω_i, writing $\log(\omega_i) = \overline{\omega}(y_i) + \sigma_\omega \varepsilon_{\omega,i}$ where $\varepsilon_{\omega,i} \sim \mathcal{N}(0,1)$. The latent draws are then over $\varepsilon_{\omega,i}$ rather than ω_i directly, which improves mixing when σ_ω is small.[36]

Given ω_i (or, equivalently, $\varepsilon_{\omega,i}$), the individual contribution to the likelihood has two components: plan choice and health spending. The plan choice component follows from the logit structure:

$$\mathcal{L}_i^{\text{choice}}(\boldsymbol{\Theta}, \omega_i) = s_{i,j_i}(\boldsymbol{\Theta}, \omega_i)$$

where j_i is the observed plan choice and s_{i,j_i} is as defined above.

For the spending component, the key observation is that, given ω_i and the chosen plan j_i, the first order condition establishes a one-to-one mapping between observed spending h_i and the health shock λ_i. This allows me to evaluate the spending density via a change of

[34] Details on the estimation procedure are provided in Appendix B.

[35] See Tanner and Wong (1987) and Albert and Chib (1993) for the general data augmentation framework, and Train (2009) for its application to discrete choice models.

[36] This is standard practice in hierarchical Bayesian models; see, e.g., Papaspiliopoulos et al. (2007).

variables. Specifically, for $h_i > 0$, I invert the first order condition to obtain

$$\lambda_i = \Lambda(h_i, \omega_i, j_i) \equiv \frac{h_i}{1 + \omega_i \left(1 - \frac{1-c_{ij_i}}{z_i}\right)}$$

where $z_i = y_i - p_{ij_i} - (1 - c_{ij_i})h_i$ is observed consumption.[37] Then, the spending component of the likelihood is

$$\mathcal{L}_i^{\text{spend}}(\boldsymbol{\Theta}, \omega_i) = \begin{cases} \kappa(X_i) & \text{if } h_i = 0 \\ \big(1 - \kappa(X_i)\big)\ g^+(\Lambda(h_i, \omega_i, j_i) \mid X_i, y_i) \left|\frac{d\Lambda}{dh}\right| & \text{if } h_i > 0 \end{cases}$$

where g^+ is the spliced log-normal-Pareto density defined previously and $|d\Lambda/dh|$ is the Jacobian of the inversion.[38]

The full individual likelihood, conditional on ω_i, is then

$$\mathcal{L}_i(\boldsymbol{\Theta}, \omega_i) = \mathcal{L}_i^{\text{choice}}(\boldsymbol{\Theta}, \omega_i) \times \mathcal{L}_i^{\text{spend}}(\boldsymbol{\Theta}, \omega_i)$$

and the complete data likelihood, augmented with the latent $\varepsilon_{\omega,i}$, is

$$\mathcal{L}(\boldsymbol{\Theta}, \varepsilon_\omega) = \prod_{i=1}^{N} \mathcal{L}_i(\boldsymbol{\Theta}, \omega_i(\varepsilon_{\omega,i})) \times \phi(\varepsilon_{\omega,i})$$

where ϕ denotes the standard normal density. Notice that, given ω_i and the observed plan choice, λ_i is not a separate latent variable. Rather, it is pinned down deterministically by the first-order condition inversion for individuals with positive spending, and equals zero otherwise. Hence, $\varepsilon_{\omega,i}$ is the only latent variable requiring MCMC sampling.

MCMC Implementation. The model is estimated using a Metropolis-Hastings algorithm that alternates between updating the population parameters $\boldsymbol{\Theta}$ and the individual latent variables $\varepsilon_{\omega,i}$. For the population parameters, I use component-wise Metropolis-Hastings updates, cycling through each parameter conditional on the current values of

[37] This inversion is valid when $y_i - p_{ij_i} > 1 - c_{ij_i}$ on the chosen plan, which is the regularity condition discussed above. Under this condition, $h_{ij}^* = 0 \iff \lambda_i = 0$, so the inversion is well-defined for all positive spending levels.

[38] The Jacobian can be computed in closed form from the first order condition. Its derivation is provided in Appendix B.2.

the others and the latent variables. Each proposal is evaluated using the relevant components of the likelihood and accepted or rejected according to the Metropolis-Hastings ratio. For the plan fixed effects ξ_j, I adopt a hierarchical prior $\xi_j \sim \mathcal{N}(\mu_\xi, \sigma_\xi^2)$, where the hyperparameters μ_ξ and σ_ξ are estimated jointly with the other parameters. This regularizes the fixed effects and avoids overfitting, particularly for plan groups with smaller market shares.

The model is estimated using a Metropolis-Hastings algorithm that alternates between updating the population parameters $\mathbf{\Theta}$ and the individual latent variables $\varepsilon_{\omega,i}$. For the population parameters, I use component-wise Metropolis-Hastings updates, cycling through each parameter conditional on the current values of the others and the latent variables. Each proposal is evaluated using the relevant components of the likelihood and accepted or rejected according to the Metropolis-Hastings ratio. To make estimation tractable at this sample size, I evaluate the plan choice likelihood on a random subsample of individuals at each Metropolis-Hastings step for the population-parameter blocks, rescaling to recover the full-sample log-likelihood.[39] For the plan fixed effects ξ_j, I adopt a hierarchical prior $\xi_j \sim \mathcal{N}(\mu_\xi, \sigma_\xi^2)$, where the hyperparameters μ_ξ and σ_ξ are estimated jointly with the other parameters. This regularizes the fixed effects and avoids overfitting, particularly for plan groups with smaller market shares.

For the individual latent variables $\varepsilon_{\omega,i}$, proposals are evaluated using both the plan choice likelihood and the spending likelihood for individual i, together with a standard normal prior on $\varepsilon_{\omega,i}$. Given the large sample size, I update a random subsample of individuals at each iteration rather than the full sample, which substantially reduces the computational cost per iteration without affecting the ergodicity of the chain.[40]

I use weakly informative priors for all parameters. The chain is run for 60,000 iterations after a burn-in period of 15,000 iterations, with posterior means used as point estimates in the counterfactual analysis.[41]

[39] This is the noisy-Metropolis-Hastings approach (Korattikara et al., 2014; Bardenet et al., 2017; Maire et al., 2019), which targets a small-bias approximation to the posterior in exchange for substantial computational savings; Quiroz et al. (2019) develop a more complex, bias-corrected variant that I do not implement. Further details on the subsampling strategy, the adaptation of proposal variances, and convergence diagnostics are provided in Appendix B.

[40] Each selected individual is updated using the exact likelihood for that individual; this is random-scan Gibbs sampling (Liu, 2001; Roberts and Rosenthal, 1997), which targets the exact posterior, in contrast to the noisy-Metropolis-Hastings approach used for the population-parameter blocks above.

[41] The posterior means of the individual latent variables $\varepsilon_{\omega,i}$ are computed as running averages over the post-burn-in iterations, yielding individual-specific posterior mean estimates of ω_i.

Identification. The parameters are identified from different sources of variation in the data, which I discuss informally here.

The zero spending probability $\kappa(X_i)$ is identified from the fraction of enrollees with zero medical expenses within each demographic group. The log-normal parameters $\mu(X_i, y_i)$ and $\sigma(X_i)$ are identified from the distribution of positive spending within each demographic group, conditional on ω_i and the chosen plan's coverage rate, which together determine the mapping between observed spending and the latent health shock λ_i. The Pareto shape parameter θ is identified from the tail behavior of high-spending individuals across all demographic groups.

The parameters $(\omega_0, \alpha_\omega, \sigma_\omega)$ are identified from how spending responds to variation in coverage across plans. Two individuals with similar health profiles but enrolled in plans with different coverage rates c_{ij} will exhibit different spending levels, with the gap governed by ω_i. This is the standard identification argument in the literature: coverage variation traces out the healthcare price response, separating the effect of coverage on spending from underlying health differences. The income gradient α_ω is identified from how the spending-coverage relationship varies with income.

The plan fixed effects ξ_j are identified from observed market shares, conditional on the premiums and coverage rates that enter expected utility. Plans that attract more enrollees than predicted by their financial characteristics alone must have positive unobserved attributes captured by ξ_j.

Finally, the income gradient in the health shock distribution, α_μ, is identified from the relationship between income and spending levels, after controlling for the moral hazard channel. Since both ω_i and μ depend on income, the joint estimation separates them: ω_i governs how spending responds to coverage variation at a given health level, while α_μ governs how the health level itself varies with income.

5.3 Parameter Estimates and Model Fit

Table 2 reports posterior means and credible intervals for the structural parameters, while the estimated fixed effects are shown in Figure 4..

Recall that the income-spending relationship is crucial for the potential welfare gains via selection (within the equilibrium channel). This relationship will depend on how the distribution of health shocks (or underlying health) varies with income and on how price

Table 2: Estimates

	Mean	Median	SD	95% CI	$\hat{R}$	ESS
ω Distribution						
ω_0	0.131	0.121	0.102	[-0.051, 0.339]	1.003	306.641
α_ω	0.612	0.612	0.034	[0.551, 0.680]	1.009	307.889
σ_ω	0.786	0.786	0.010	[0.767, 0.804]	1.018	337.400
δ_ω	-4.975	-4.982	0.025	[-4.999, -4.909]	1.000	3966.337
Health Distribution						
α_μ	-0.517	-0.515	0.027	[-0.570, -0.468]	1.005	306.611
δ_μ	2.357	2.356	0.029	[2.299, 2.414]	1.014	336.422
θ	1.139	1.138	0.014	[1.113, 1.165]	1.001	8568.135
Men – Age: 18-29						
μ_0	-1.378	-1.372	0.079	[-1.535, -1.228]	1.001	313.292
σ	0.375	0.375	0.007	[0.362, 0.389]	1.000	1825.215
κ	0.451	0.451	0.003	[0.445, 0.456]	1.000	12178.72
Men – Age: 30-44						
μ_0	-1.292	-1.287	0.079	[-1.446, -1.144]	1.001	308.430
σ	0.358	0.358	0.004	[0.351, 0.365]	1.002	1666.228
κ	0.045	0.045	0.001	[0.043, 0.047]	1.000	13988.40
Men – Age: 45-64						
μ_0	-1.278	-1.272	0.079	[-1.435, -1.130]	1.001	313.141
σ	0.351	0.351	0.006	[0.339, 0.363]	1.001	1983.507
κ	0.030	0.030	0.001	[0.027, 0.032]	1.000	14224.79
Men – Age: 65 and above						
μ_0	-1.236	-1.232	0.081	[-1.398, -1.085]	1.001	328.762
σ	0.343	0.343	0.012	[0.320, 0.365]	1.001	1548.118
κ	0.002	0.002	0.000	[0.001, 0.004]	1.000	13898.77
Women – Age: 18-29						
μ_0	-1.335	-1.329	0.081	[-1.497, -1.186]	1.001	316.300
σ	0.356	0.356	0.008	[0.340, 0.371]	1.001	1388.115
κ	0.395	0.395	0.003	[0.388, 0.401]	1.000	12950.18
Women – Age: 30-44						
μ_0	-1.285	-1.280	0.080	[-1.440, -1.135]	1.001	309.940
σ	0.363	0.363	0.004	[0.354, 0.372]	1.011	1626.725
κ	0.042	0.042	0.001	[0.040, 0.044]	1.000	14112.53
Women – Age: 45-64						
μ_0	-1.285	-1.279	0.079	[-1.441, -1.138]	1.001	311.427
σ	0.348	0.348	0.005	[0.337, 0.358]	1.000	1858.887
κ	0.006	0.006	0.001	[0.006, 0.008]	1.000	14065.31
Women – Age: 65 and above						
μ_0	-1.272	-1.266	0.080	[-1.431, -1.127]	1.002	319.439
σ	0.324	0.325	0.009	[0.306, 0.343]	1.000	1608.434
κ	0.000	0.000	0.000	[0.000, 0.001]	1.000	13751.31
Fixed Effects Hyperparameters						
μ_ξ	0.807	0.808	0.542	[-0.267, 1.890]	1.000	60000
σ_ξ	1.983	1.931	0.365	[1.421, 2.844]	1.000	53684.07

Notes: Results reported from the posterior distributions obtained from the MCMC estimation with the 95% credible intervals. $\hat{R}$ is the Gelman-Rubin statistic and ESS stands for effective sample size. The estimation works with logit(κ); in the table, I report κ directly via the sigmoid transformation $s(x) = 1/(1 + exp(-x))$. For the standard deviations, I use the delta method, so if the standard deviation of the logit(κ) chain is $std_{logit(\kappa)}$ and $\text{mean}_{logit(\kappa)}$ is its mean, the standard deviation reported in the table is $std_\kappa = s(\text{mean}_{logit(\kappa)}) \times (1 - s(\text{mean}_{logit(\kappa)})) \times std_{logit(\kappa)}$.

Figure 4: Estimated Fixed Effects (95% Credible Intervals)

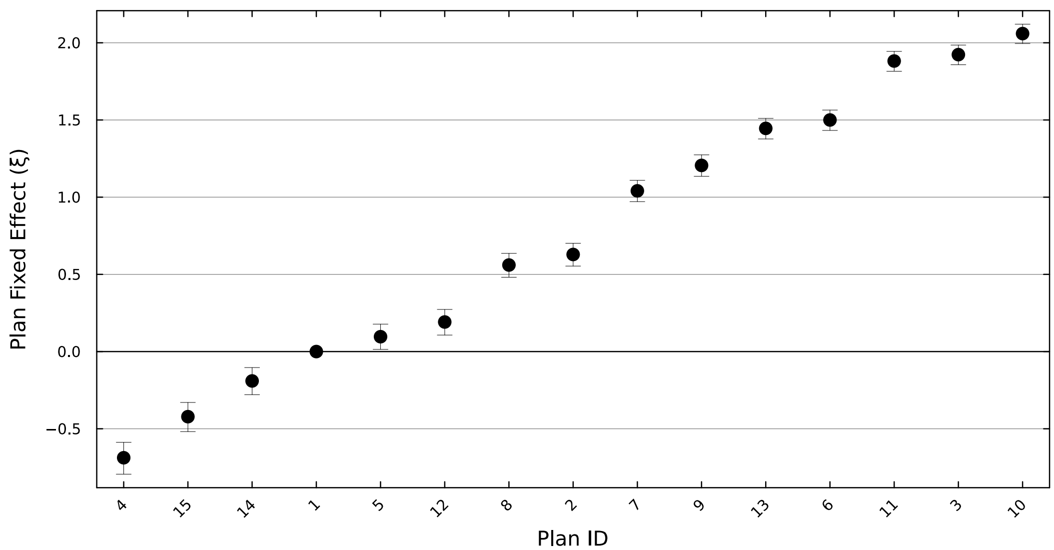

Notes: Estimated fixed effects, sorted by value of the point estimate, with their respective 95% credible intervals. Plan 1 is the reference plan. These results in table format are shown in Appendix Table A.2

sensitivity to healthcare prices relates to income. The income coefficient in the health shock distribution is negative ($\alpha_\mu = -0.517$), so higher-income enrollees draw from a health shock distribution shifted toward lower realizations: richer enrollees are healthier, conditional on demographics. At the same time, the income coefficient in the moral hazard parameter is positive ($\alpha_\omega = 0.612$). It is important to remember, though, that this does not automatically imply that richer people are more price-sensitive (i.e., exhibit larger degrees of moral hazard) than poorer people, since income also has a direct effect due to the concavity of consumption utility that has to be taken into account, as noted in section 3. I explore this relationship in more detail below.

The demographic-specific health parameters are estimated precisely and follow the expected patterns. The probability of zero spending, κ, is large for the youngest enrollees (around 45% and 40% for the 18-29 groups of men and women, respectively) and falls sharply with age (near zero for the 65 and above groups). The dispersion of the health shock distribution, σ, is stable across demographic cells, and the Pareto tail index, $\theta = 1.14$, implies a heavy right tail in spending typical of health spending distributions (Karlsson et al., 2024). All parameters display satisfactory convergence, with most Gelman-Rubin $\hat{R}$ statistics below 1.01, all below 1.02, and effective sample sizes adequate for the quantities of interest.[42]

Aggregate fit. The model reproduces the aggregate features of the market that the counterfactual relies on. Figure 5 compares observed and model-predicted plan shares across the fifteen plan groups, which track closely. Figure 6 compares the observed and predicted distributions of health spending, which are likewise close, including in the right tail that drives insurer costs. Fit at the level of individual demographic cells is reported in Appendix A.[43]

Figure 5: Model Fit – Plan Shares

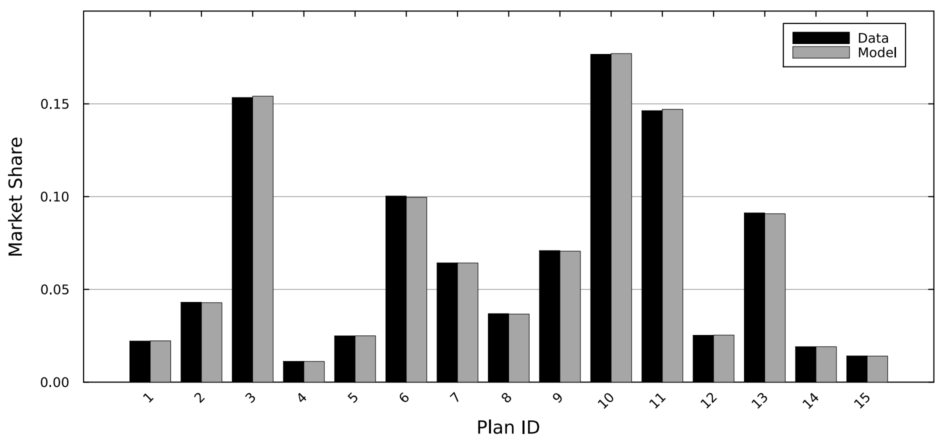

Notes: Figure shows the plan market shares observed in the data and the ones implied by the model for each of the 15 plan groups using the reported posterior means.

The spending-income relationship. Figure 7 plots mean health spending against income (in windows), in the data and as predicted by the model. The model captures the empirical pattern well. Two features stand out. First, the gradient is modest: spending rises with income, but not dramatically, consistent with the offsetting health and moral hazard channels discussed above. Second, Figure 8 shows that holding demographics

[42]The population-level location parameters (ω_0, α_ω, α_μ, μ_0) mix more slowly because they enter the plan-choice integral for every enrollee; their effective sample sizes (around 300) remain sufficient for stable posterior means, which are the inputs to the counterfactual analysis.

[43]Cell-level share fit is less tight than the aggregate fit. This reflects a deliberate modeling choice. Plan preferences enter through a hierarchically-shrunk fixed effect ξ_j rather than a saturated demographic-by-plan specification. A richer interaction would improve cell-level fit mechanically but at the cost of overfitting preference heterogeneity that the counterfactual (a change in the income-proportional floor, operating through income rather than through demographic plan taste) does not require. The aggregate share and spending fits are the relevant targets for the welfare exercise.

Figure 6: Model Fit – Healthcare Spending

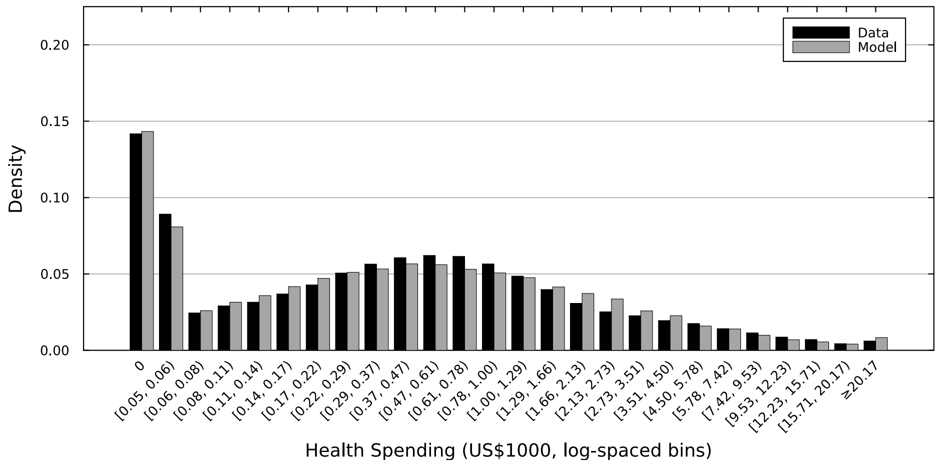

Notes: Figure shows the model fit for spending considering the reported posterior means. For zero spending, I report the share observed in the data and the share implied by the model. For positive spending, I consider spending windows (in log scale) and report the shares observed in the data and implied by the model.

fixed flattens the residual relationship further. Because ISAPRE premiums are risk-rated on precisely these demographics, this pattern is relevant when analyzing the scope for selection.

Figure 7: Income-spending Relationship – Data vs. Model

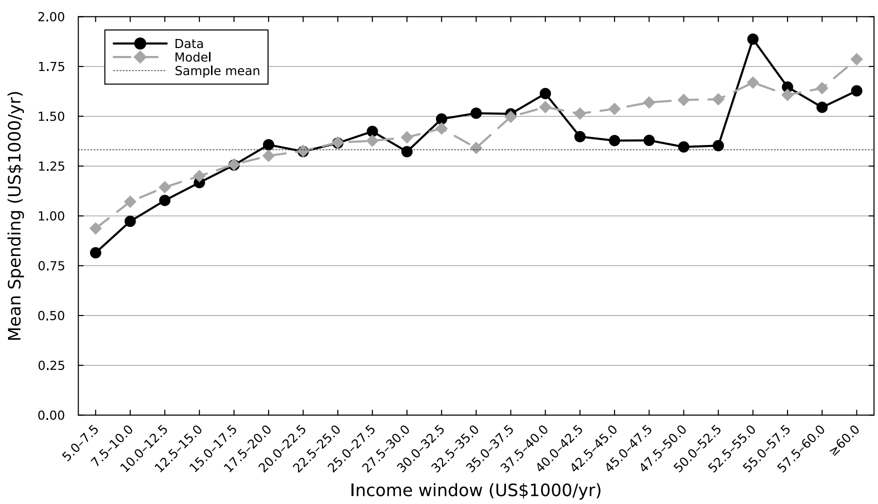

Notes: Relationship between income and healthcare spending in the data *versus* the one implied by the model. In each case, I consider income windows of fixed size (except the last, catch-all window) and calculate the mean spending within that income window.

Figure 8: Income-spending Relationship – Unconditional vs. Fixed Demographics

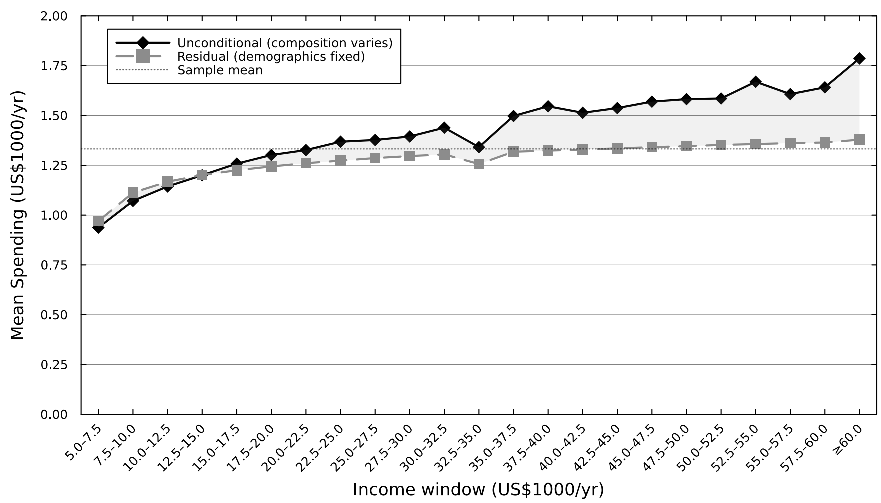

Notes: Relationship between income and healthcare spending implied by the model, unconditional and holding demographic composition constant. In each case, I consider income windows of fixed size (except the last, catch-all window) and calculate the mean spending within that income window.

Finally, one of the key reasons it is important to separate underlying health from moral hazard is that when people switch to plans with better coverage, they may adjust their health spending accordingly, which is not directly observed in the data. If this responsiveness is different across the income distribution, this can affect the income-spending relationship. Recall that this is not simply observed from the estimates, since this is a combination of how the ω_i distribution varies with income and how income interacts with this due to decreasing marginal utility of consumption. To have a sense of how this responsiveness varies across the income distribution, Figure 9 shows the model-implied income-spending relationship for two scenarios, both ignoring premiums. First, I consider the relationship when all agents have a coverage of 67%, the approximate unconditional mean of coverage they face. Then, I consider how the income-spending relationship changes when all enrollees are fully covered. As the figure makes clear, poorer people tend to be *more* responsive than richer people. Since richer people are more likely to acquire better plans (due to the mandate and the fact that better plans are more expensive), this indicates that the income-spending relationship for high-coverage plans is likely even flatter than what we observe in the data, and can even become negative. This tends to make selection even less likely to generate net welfare benefits in the equilibrium channel.

Figure 9: Income-spending Relationship – 60% Coverage vs. Full Coverage

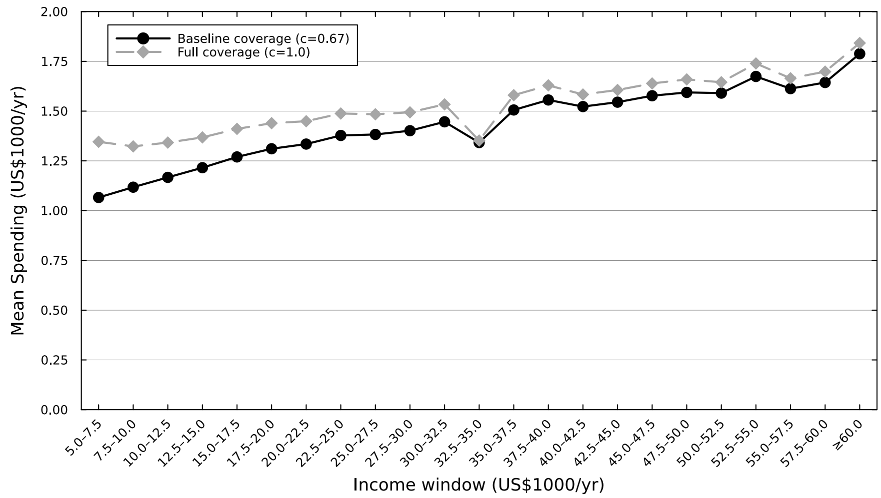

Notes: Relationship between income and healthcare spending implied by the model when all enrollees face a 67% coverage rate *versus* when they all have full coverage. In both cases, premiums are set to zero. In each case, I consider income windows of fixed size (except the last, catch-all window) and calculate the mean spending within that income window.

6 Counterfactual Analysis

In this section, I use the estimated model to evaluate the welfare effects of changing the mandatory contribution rate. The analysis centers on two counterfactual experiments. The first raises the mandate from 7% to 8% of income and traces out the new market equilibrium. Following the discussion in section 3, I decompose its welfare effect into a fiscal channel and an equilibrium channel. The second is a revenue-equivalent direct transfer that raises the same revenue as the mandate increase but distributes it as a cash rebate outside the insurance market, serving as a benchmark for the mandate's value as a redistributive instrument. The remainder of this section specifies the supply side that closes the model (subsection 6.1), the welfare measure I use (subsection 6.2), the formal decomposition (subsection 6.3), and the precise construction of each experiment (subsection 6.4), before turning to results in subsection 6.5.

6.1 Supply Side

The demand model estimated in the previous section takes premiums and contracts as given. To conduct counterfactual exercises, I need to specify how insurers react to changes in the mandate. First, I work in the fixed-contracts tradition (Einav and Finkelstein, 2011), taking the set of offered contracts and their characteristics (except premiums) as given. Second, I have to specify how insurers adjust premiums in response to changes in the mandate. I follow previous papers in the literature (Bundorf et al., 2012; Handel, 2013; Einav et al., 2013; Polyakova, 2016) and summarize pricing behavior using a reduced-form relationship between expected costs and premiums, rather than solving an explicit profitmaximization problem. Specifically, I adopt a constant-margin assumption: each plan's margin, defined as the ratio of base premium to expected insurer costs per enrollee, is held fixed at its estimated value under the baseline 7% mandate.[44] When the mandate changes and the composition of the risk pools of the plans consequently shifts, base premiums adjust proportionally to the change in expected costs, maintaining the same percentage margin. This constant-margin assumption is a natural extension of the perfect competition (where price equals average cost) assumption widely used in the literature (Einav et al., 2010; Hackmann et al., 2015; Einav et al., 2013) and very similar to the approach of assuming that price equals average cost plus a fee used by Handel (2013). The standard assumption that competitive pressure drives prices to expected costs yields zero margins in equilibrium. Since ISAPREs are for-profit insurers operating under risk-adjusted pricing, observed margins are non-zero. Holding these margins fixed at their baseline values under the counterfactual scenarios is the natural analog of the zero-margin assumption when margins are positive in the data. As Einav and Finkelstein (2011) note, prices tracking average costs with fixed margins is also observed in practice. The assumption is also particularly defensible in this setting because the counterfactual exercises I consider involve a 1 percentage-point change in the mandate rate (from 7% to 8% of income), a relatively small perturbation. For larger changes, insurers might respond strategically to substantial shifts in market structure.

More broadly, the small perturbation nature of the exercise is a feature of the analysis in other dimensions. A 1 percentage-point change is small enough that the boundary between FONASA and ISAPREs is unlikely to shift materially, justifying the focus on the within-ISAPRE intensive margin. Larger reforms, such as eliminating the mandate or substantially raising it, could induce movement between the public and private systems and trigger more substantial within-ISAPRE plan re-sorting, thereby activating selection margins that are essentially inactive in the marginal exercises considered here. The decompo-

[44] I am abstracting here from administrative and commercialization costs, focusing on the difference between premiums and expected medical costs.

sition framework applies to such reforms in principle, but quantifying their welfare effects would require, at a minimum, modeling the FONASA-ISAPRE boundary and relaxing the constant-margin assumption. Therefore, the empirical results in this paper should be read as characterizing the welfare consequences of marginal changes around the current 7% contribution rate, not as a global statement about the welfare effects of contribution-floor mandates at arbitrary levels, which, in any case, would depend on other characteristics of the particular market considered as well.

Under the aforementioned assumptions, the counterfactual equilibrium for a given mandate rate τ is computed iteratively. Starting from the initial base premiums (from the baseline scenario with a 7% mandate), I compute choice probabilities and expected insurer costs for each plan, then update base premiums to maintain the baseline margins given the new costs. This process is repeated until premiums converge.[45] At each iteration, the premium paid by each individual is $p_{ij} = \max\{\tilde{p}_{ij}, \tau \min\{y_i, y_{cap}\}\}$, where $\tilde{p}_{ij}$ is the risk-adjusted premium at the current iteration's base premiums.

6.2 Welfare Measurement

I measure consumer welfare using the compensating variation (CV). For each individual i, the CV of a policy change from the baseline (7% mandate, equilibrium premiums $\mathbf{p}^*_{7\%}$) to a counterfactual scenario is the income adjustment Δy_i that makes the individual indifferent between the two scenarios:

$$\mathbb{E}\left[V_i(y_i + \Delta y_i, \mathbf{p}^*_{7\%})\right] = \mathbb{E}\left[V_i(y_i, \mathbf{p}^*_{\text{CF}})\right]$$

where V_i denotes the expected indirect utility of individual i (integrating over plan choice and health shocks), and $\mathbf{p}^*_{\text{CF}}$ denotes the counterfactual premiums. A negative CV indicates the individual is worse off under the counterfactual, since she would have to receive extra money to be indifferent between the two situations.[46]

[45]Convergence is achieved in all exercises within a small number of iterations, reflecting the fact that the perturbation from the baseline is small. I use a damping parameter to ensure stability of the iterations.

[46]I compute the CV evaluating expected utility at the modified income levels given the new mandate, while holding ω_i and the health shock distribution at their baseline values, i.e., holding the income effect on these two objects at the same level as in the baseline. In other words, the income adjustment enters only through the budget constraint, not through the moral hazard parameter or health.

The insurer welfare change is measured by the change in total insurer profits:

$$\Delta\Pi = \sum_j \left[R_j^{\mathrm{CF}} - C_j^{\mathrm{CF}}\right] - \sum_j \left[R_j^{\mathrm{base}} - C_j^{\mathrm{base}}\right]$$

where R_j and C_j denote expected revenues and costs for plan j, computed using the choice probabilities and expected insurer costs from the model.

Total welfare is then

$$\Delta W = \sum_i \mathrm{CV}_i + \Delta\Pi$$

which accounts for both the consumer and insurer sides of the market.

Distributional weights. When evaluating welfare, I follow the standard approach of computing welfare both unweighted and with distributional weights $w_i = (y_i/\bar{y})^{-\varepsilon}$, where $\bar{y}$ is the mean income in the sample and $\varepsilon \geq 0$ is the inequality aversion parameter from Atkinson (1970). At $\varepsilon = 0$, all individuals receive equal weight. As ε increases, the welfare changes of lower-income individuals receive proportionally larger weight. A value of $\varepsilon = 1$ corresponds to a planner with log utility over income (the focal case in Atkinson (1970)), and $\varepsilon = 2$ to a planner with stronger preferences for redistribution. I report results for $\varepsilon \in \{0, 0.5, 1, 2\}$ throughout. The choice of ε matters when comparing the mandate increase against the direct transfer alternative (both of which will be described in detail below), since the policies differ in their distributional incidence.

6.3 Welfare Decomposition

One of the goals of the empirical exercise is to decompose the total welfare effect of the mandate increase into a *fiscal channel* and a *equilibrium channel*. As discussed in the theoretical framework, income-based mandates combine two distinct economic forces: they mechanically transfer resources from bound enrollees to insurers, and they may alter plan choices and the risk pool composition of the health plans. To separate these forces, I perform the following decomposition.

Fiscal channel. The fiscal counterfactual raises the mandate from 7% to 8% but holds plan choices fixed at their baseline (7% mandate) probabilities. Under this counterfactual, each individual's welfare changes because they pay a higher premium floor on their existing plan allocation, but no reoptimization of plan choice occurs. Formally, the fiscal CV for individual i is the income adjustment that satisfies

$$\sum_j s_{ij}^{\text{base}} \cdot U_{ij}(y_i + \Delta y_i^{\text{fiscal}}, \mathbf{p}_{7\%}^*) = \sum_j s_{ij}^{\text{base}} \cdot U_{ij}(y_i, \mathbf{p}_{8\%}^{\text{floor}})$$

where s_{ij}^{base} are the baseline choice probabilities, and $\mathbf{p}_{8\%}^{\text{floor}}$ denotes premiums with the 8% floor applied but base premiums held at their 7% equilibrium values. This is, in essence, the welfare effect of the mandate if it operated purely as a tax, without any behavioral response or supply-side adjustment.

Equilibrium channel. The equilibrium channel is defined as the residual:

$$\Delta W^{\text{equilibrium}} = \Delta W^{\text{total}} - \Delta W^{\text{fiscal}}$$

This channel captures all behavioral responses to the mandate increase: plan switching by consumers in response to the compression of premium differentials across plans, the resulting changes in risk pool composition, and the supply-side premium adjustment that follows. The equilibrium channel is the margin through which the mandate could, in principle, improve welfare by correcting adverse selection. It is also the margin on which the existing literature, focused on extensive margin mandates, has found the largest welfare effects.

6.4 Counterfactual Exercises

I now specify the construction of each experiment.

Mandate increase from 7% to 8%. Raising τ from 0.07 to 0.08, I compute the new equilibrium under the constant-margin assumption of subsection 6.1: starting from the baseline base premiums, I iterate between choice probabilities, expected insurer costs, and the premium update that holds each plan's margin fixed, until premiums converge.[47] The total welfare change is then decomposed into the fiscal and equilibrium channels of subsection 6.3.

[47] Convergence obtains within a small number of iterations in all exercises, as discussed above.

Distributional incidence. I report the welfare effect of the mandate increase across income quartiles, age-gender groups, and quintiles of expected health risk ($\mathbb{E}[\lambda_i]$), and across the joint distribution of income and $\mathbb{E}[\lambda_i]$. For each cell I compute the mean compensating variation and the share of enrollees made better off, both unweighted and under the inequality-aversion weights $w_i = (y_i/\bar{y})^{-\varepsilon}$ for $\varepsilon \in \{0, 0.5, 1, 2\}$.

Revenue-equivalent direct transfer. The transfer levies an additional 1% of income on the enrollees who would be bound under the 8% mandate, thus raising the same revenue as the mandate wedge, and rebates it as a lump-sum cash transfer to unbound enrollees. More specifically, each bound individual's income is reduced by $0.01 \times y_i$, and each unbound individual's income is raised by a per-capita transfer equal to total revenue divided by the number of unbound individuals. The mandate remains at 7%, and insurers face the 7% equilibrium premiums, so the insurer side is unaffected and the transfer is a pure government transaction.

6.5 Counterfactual: Raising the Mandate from 7% to 8%

This section reports the results of the counterfactuals exercises mentioned above, computed under the constant-margin supply side described in subsection 6.1 and decomposed into the fiscal and equilibrium channels defined in subsection 6.3. I report the aggregate decomposition first, then its distributional incidence, and finally the comparison with a revenue-equivalent direct transfer.

6.5.1 Aggregate Effects and Decomposition

Results are shown in Table 3. The main takeaway is that, in the basic scenario ($\varepsilon = 0$), the mandate increase operates as a near-pure fiscal transfer. Moving from 7% to 8% leaves consumers worse off by a compensating variation of approximately 38.9 million 2012 USD in aggregate, while insurer surplus also rises by approximately 38.9 million 2012 USD. The two magnitudes nearly offset each other: the additional contribution paid by bound enrollees is captured by insurers as the mandate wedge, and little else happens. The net welfare effect is 12 thousand USD, three orders of magnitude smaller than the gross transfer between the two sides of the market. With 202,662 enrollees in the sample, this means the total welfare gain per enrollee is low, around 0.06 USD (6 cents of a dollar) per enrollee per year.

The last panel in Table 3 shows that the reason so little else happens is two-fold. First, plan choice probabilities move very little (around 0.01 percentage points), so the mass of unbound enrollees affected by the change in mandate is low. But second, and more importantly, even a larger change would likely generate no gains via equilibrium effects (and selection, in particular), since the equilibrium channel has almost nothing to operate on, exactly as the spending-income evidence in subsection 5.3 anticipated. A strong equilibrium channel requires a relatively strong income-spending correlation among enrollees whose plan choices could change beyond what is already risk-rated based on demographics by insurers. In this market, the residual cost curve is nearly flat, especially once demographics are held fixed (Figure 8). Moreover, many high-income enrollees in this market are already acquiring a plan for 7% of their income, and thus newly-bound people tend to have lower income. Since healthcare price sensitivity decreases with income, newly bound individuals tend to be more sensitive than current enrollees in high-coverage plans. Hence, if these agents switch to better plans, they also increase their spending more substantially, further flattening or even flipping the income-spending relationship for high-coverage plans, which is reflected in the fact that, as Table 3 shows, potential switchers are, if anything, slightly more costly on average than incumbents.

This stands in contrast to the penalty-based mandate literature, where bringing healthier enrollees into the pool is the primary source of welfare gains (Hackmann et al., 2015). In my case, this selection correction is mechanically inactive because the cost variation it would exploit, beyond what is already internalized in premiums, is almost non-existent.

6.5.2 Distributional Incidence

Although the mandate increase is close to welfare-neutral in aggregate, it is not neutral across enrollees. Table A.3 reports the mean compensating variation by income quartile. Two patterns emerge. First, the incidence is concentrated entirely among bound enrollees: those whose risk-adjusted premium already exceeds the 7% floor are slightly better off due to the small decrease in premiums caused by the mandate, while bound enrollees bear a loss equal to the additional one percentage point of income. Because binding status rises with income, the burden falls on higher-income enrollees in absolute terms. Second, within the bound population, the loss is purely a function of income and not of health: conditional on being bound, an enrollee's compensating variation is $-(\tau_1 - \tau_0)y_i$ regardless of λ_i, because the additional contribution buys no additional coverage.

Table 3: Welfare Effects of Raising the Mandate from 7% to 8%

	Inequality aversion (ε)			
	0.0	0.5	1.0	2.0
$\sum_i \text{CV}_i$ (weighted)	-38,948.30	-32,666.20	-26,106.06	-14,526.45
$\Delta\Pi$ (insurer profit)	38,960.24	38,960.24	38,960.24	38,960.24
ΔW total	11.93	6,294.04	12,854.18	24,433.78
Decomposition				
Fiscal channel	7.99	6,288.38	12,846.70	24,422.87
Equilibrium channel	3.94	5.66	7.49	10.91
% from fiscal channel	67.0	99.9	99.9	100.0
Mandate wedge				
Volume transferred	40,760.28			
$\Delta\Pi$ except wedge	-1,800.05			
Mechanism Analysis				
Mean Prob Mass Shift (p.p.)	0.01			
Average Spending: Already Bound	0.258			
Average Spending: Newly Bound	0.270			
Average Spending: Never Bound	0.276			

Notes: All welfare quantities are in 1,000 US$ in 2012. Weighted CV uses Atkinson-style social welfare weights with inequality aversion parameter ε, where $\varepsilon = 0$ is utilitarian and $\varepsilon \to \infty$ is Rawlsian. The fiscal channel holds plan choices at their 7% baseline probabilities. The equilibrium channel is the residual. The mandate wedge is the gross transfer from bound enrollees to insurers. $\Delta\Pi$ except wedge captures the small base-premium adjustment. Mean probability mass shift (in percentage points) is calculated as $0.5 \times \sum_{ij} |Pr_{ij}(8\%) - Pr_{ij}(7\%)|$, where $Pr_{ij}(x\%)$ is the probability that individual i chooses plan j when the mandate is x%. The average spending measures (in thousands of 2012 USD) are weighted by choice probabilities for every possible choice; the status of newly or never bound is based on each enrollee-plan match separately in the counterfactual scenario.

This incidence pattern is the key to the policy comparison that follows. Being a near-transfer in this setting, the mandate is essentially taking money from consumers, and in particular from richer consumers, who are more likely to be bound by the mandate, and transferring it to insurers, with a very small gain for some poorer consumers due to a small decrease in premiums for unbound individuals. It becomes then natural to consider how a policy focused on a direct transfer fares against the mandate.

Distributional weighting. The aggregate near-neutrality also masks redistribution once welfare is weighted for inequality aversion. Applying weights $w_i = (y_i/\bar{y})^{-\varepsilon}$, the distributionally weighted welfare effect of the mandate increases, reflecting that the burden falls on higher-income enrollees whose marginal welfare receives lower weight. I report the full schedule for $\varepsilon \in \{0, 0.5, 1, 2\}$ alongside the transfer comparison in Table 4, since the two policies are most naturally compared on the same weighted basis.

6.5.3 Revenue-Equivalent Direct Transfer

The preceding results imply that routing redistribution through the mandate is costly: the instrument transfers resources from bound enrollees to insurers, and the incidence falls without regard to who values the implicit coverage. A natural benchmark is a fiscal policy that raises the same revenue but distributes it directly. I consider a transfer that levies an additional 1% of income on the enrollees who would be bound under the 8% mandate, raising the same revenue as the mandate wedge, and rebates it as a lump-sum cash transfer to unbound enrollees. The transfer enters only through disposable income; the insurance market is left at its 7% equilibrium, so insurers neither gain nor lose.[48]

Table 4 compares the two policies across degrees of inequality aversion. At $\varepsilon = 0$, where all enrollees receive equal weight, the transfer performs worse than the mandate with a negative net welfare change, though the magnitude is still small relative to the volume of transferred resources and the value per enrolee per year is also small, less than 0.20 USD. As ε rises, the policies diverge sharply. With equity concerns (i.e., $\varepsilon > 0.0$), the direct transfer delivers approximately three times the total welfare of the mandate because it avoids the insurer capturing the wedge and can direct the rebate toward lower-income unbound enrollees. The mandate, in contrast, hands the wedge to insurers and burdens

[48] Concretely, each bound enrollee's income is reduced by $0.01 \times y_i$ and each unbound enrollee's income is raised by a per-capita amount equal to total revenue divided by the number of unbound enrollees. Because the mandate stays at 7%, this is a pure government transaction with no insurer incidence.

middle- and higher-income bound enrollees who receive little coverage value in return. On a per-enrollee basis, this translates to a difference on the range between 57 USD per enrollee per year (when $\varepsilon = 0.5$) and 254 USD per enrollee per year (when $\varepsilon = 2$)

Table 4: 8% Mandate vs. Direct Transfer Counterfactual

	Inequality aversion (ε)			
	0.0	0.5	1.0	2.0
ΔW – 8% mandate	11.93	6,294.04	12,854.18	24,433.78
ΔW – direct transfer	-35.87	17,927.78	37,546.94	75,899.38
Gap (transfer − mandate)	-47.80	11,633.74	24,692.76	51,465.60
Ratio (transfer / mandate)	-3.01	2.85	2.92	3.11

Notes: The direct transfer counterfactual collects the same revenue as the 8% mandate increase (from bound enrollees, scaled by 1% of capped income) and redistributes it as a lump-sum transfer to unbound enrollees. Both policies are budget-balanced. The ratio measures how much better an equivalent-revenue direct transfer performs in terms of welfare.

This comparison underscores that, whatever redistributive purpose an increase in the contribution floor might serve, a revenue-equivalent cash transfer in this case serves it more effectively, because the mandate's redistribution is intermediated by an insurer that captures the transfer rather than passing it through. In a risk-rated market where the equilibrium channel is inactive, the contribution floor is dominated, as a redistributive instrument, by a direct fiscal transfer that does not pass through the insurance market. From a social planner's perspective, the burden of addressing the problem of intermediation is justified only if the gains from selection are substantial.

7 Conclusion

This paper develops a framework for analyzing the welfare effects of contribution-floor mandates, a class of regulatory instruments used in health insurance systems in Germany, the Netherlands, Switzerland, Colombia, Chile, and elsewhere, and implements it empirically using administrative data from Chile's private health insurance market. The framework decomposes the welfare effect of a mandate change into a fiscal channel, capturing the mechanical transfer from bound enrollees to insurers, and an equilibrium channel, capturing behavioral responses through plan switching and risk pool recomposition. This decomposition applies to any setting in which the effective price of insurance depends on income.

The main empirical finding is that raising Chile's mandatory contribution from 7% to 8% of income operates as a near-pure fiscal transfer. The mandate increase redistributes approximately \$38.9 million annually from consumers to insurers, with a net total welfare effect of approximately \$12 thousand, or about \$0.06 per person per year without equity concerns. Plan switching is negligible and the scope for gains via selection is virtually non-existent at this margin, which makes the equilibrium channel, the primary source of welfare gains in the penalty-based mandate literature, economically negligible in this setting. A revenue-equivalent direct cash transfer, which avoids routing redistribution through insurers, generates approximately three times the distributionally-weighted welfare of the mandate increase at moderate levels of inequality aversion. The mandate is dominated as a redistributive instrument because the insurer captures the wedge between the premium paid and the plan's actuarial price.

These results are specific to Chile's risk-rated ISAPRE market, where insurers price on age and gender and the mandate operates on the intensive margin of plan choice among individuals who are already insured. In community-rated markets, or in settings where contribution floors affect the extensive margin of insurance take-up, the equilibrium channel could be larger. The decomposition framework is designed to quantify these channels in any specific setting; the relative magnitudes of the fiscal and equilibrium channels are empirical questions that depend on market structure, the degree of risk rating, and the price sensitivity of plan choices. The finding is also local in another sense: it characterizes the welfare effect of marginal changes around the current 7% contribution rate. Larger reforms (e.g., eliminating the mandate or substantially raising it) would likely activate aspects that are inactive at the margin, including movement between the public and private systems and more substantial within-ISAPRE plan switching. The framework applies to such reforms, but the empirical magnitudes could differ substantially.

Several avenues for further work remain. First, the analysis holds the set of plans fixed; allowing insurers to redesign contracts in response to mandate changes could reveal additional welfare effects, particularly if the mandate interacts with plan design incentives. Second, the framework could be applied to other income-based mandate systems—Germany's *Gesetzliche Krankenversicherung* and Colombia's contributory regime are natural candidates—where the institutional intermediation of the wedge differs from the Chilean setting. Third, the role of the contribution floor in shaping the boundary between public and private systems, which I abstract from in this paper, deserves separate analysis: if mandate changes induce switching between FONASA and ISAPREs, the extensive margin effects could be quantitatively important even if the intensive margin effects are not. Finally, the comparison with direct transfers suggests a broader agenda of evaluating whether redistribution through insurance markets is efficient relative to fiscal alternatives, a question that applies beyond health insurance to any setting where regulatory instruments bundle fiscal transfers with market design objectives.

References

Acquatella, A. (2024). Health Care Demand Among Low-Income Individuals. Working Paper.

Albert, J. H. and Chib, S. (1993). Bayesian Analysis of Binary and Polychotomous Response Data. *Journal of the American Statistical Association*, 88(422):669–979.

Andrieu, C. and Roberts, G. O. (2009). The pseudo-marginal approach for efficient Monte Carlo computations. *Annals of Statistics*, 37(2):697–725.

Atal, J. P. (2019). Lock-in in Dynamic Health Insurance Contracts: Evidence from Chile. Working Paper.

Atkinson, A. B. (1970). On the Measurement of Inequality. *Journal of Economic Theory*, 2(3):244–263.

Azevedo, E. M. and Gottlieb, D. (2017). Perfect Competition in Markets With Adverse Selection. *Econometrica*, 85(1):67–105.

Bardenet, R., Doucet, A., and Holmes, C. (2017). On Markov chain Monte Carlo methods for tall data. *Journal of Machine Learning Research*, 18(47):1–43.

Blumel, M. and Busse, R. (2020). Germany: Health System Review. In *International Health Care System Profiles*. Commonwealth Fund. Available at https://www.commonwealthfund.org/international-health-policy-center/countries/germany.

Brown, J., Duggan, M., Kuziemko, I., and Woolston, W. (2014). How Does Risk Selection Respond to Risk Adjustment? New Evidence from the Medicare Advantage Program. *American Economic Review*, 104(10):3335–3364.

Bundorf, M. K., Levin, J., and Mahoney, N. (2012). Pricing and Welfare in Health Plan Choice. *American Economic Review*, 102(7):3214–3248.

Cardon, J. H. and Hendel, I. (2001). Asymmetric Information in Health Insurance: Evidence from the National Medical Expenditure Survey. *The RAND Journal of Economics*, 99:408–427.

Cuesta, J. I., Noton, C., and Vatter, B. (2025). Vertical Integration between Hospitals and Insurers. Working Paper.

Curto, V., Einav, L., Levin, J., and Bhattacharya, J. (2021). Can Health Insurance Competition Work? Evidence from Medicare Advantage. *Journal of Political Economy*, 129(2):570–606.

De Pietro, C., Camenzind, P., Sturny, I., Crivelli, L., Edwards-Garavoglia, S., Spranger, A., Wittenbecher, F., and Quentin, W. (2015). Switzerland: Health system review. *Health Systems in Transition*, 17(4):1–288.

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

Einav, L. and Finkelstein, A. (2011). Selection in Insurance Markets: Theory and Empirics in Pictures. *Journal of Economic Perspectives*, 25(1):115–138.

Einav, L., Finkelstein, A., and Cullen, M. R. (2010). Estimating Welfare in Insurance Markets Using Variation in Prices. *Quarterly Journal of Economics*, 125(2):877–921.

Einav, L., Finkelstein, A., Ryan, S. P., Schrimpf, P., and Cullen, M. R. (2013). Selection on Moral Hazard in Health Insurance. *American Economic Review*, 103(1):178–219.

Einav, L., Finkelstein, A., and Tebaldi, P. (2025). Market Design in Regulated Health Insurance Markets: Risk Adjustment vs. Subsidies. *Journal of Political Economy: Microeconomics*, page Forthcoming.

Finkelstein, A., Hendren, N., and Shepard, M. (2019). Subsidizing Health Insurance for Low-Income Adults: Evidence from Massachusetts. *American Economic Review*, 109(4):1530–1567.

Glazer, J. and McGuire, T. G. (2000). Optimal Risk Adjustment in Markets with Adverse Selection: An Application to Managed Care. *American Economic Review*, 90(4):1055–1071.

Guerrero, R., Gallego, A. I., Becerril-Montekio, V., and Vásquez, J. (2011). Sistema de salud de Colombia. *Salud Pública de México*, 53(s2):144–155.

Hackmann, M. B., Kolstad, J. T., and Kowalski, A. E. (2015). Adverse Selection and an Individual Mandate: When Theory Meets Practice. *American Economic Review*, 105(3):1030–1066.

Handel, B. R. (2013). Adverse Selection and Inertia in Health Insurance Markets: When Nudging Hurts. *American Economic Review*, 103(7):2643–2682.

Instituto Nacional de Estadísticas (2013). Encuesta de caracterización socioeconómica nacional (CASEN) 2013. Santiago: Ministerio de Desarrollo Social, Gobierno de Chile. https://observatorio.ministeriodesarrollosocial.gob.cl/encuesta-casen-2013 (accessed June 26, 2022).

Karlsson, M., Wang, Y., and Ziebarth, N. R. (2024). Getting the right tail right: Modeling tails of health expenditure distributions. *Journal of Health Economics*, 97:102912.

Kauer, L., McGuire, T. G., Schillo, S., and van Kleef, R. C. (2026). Regulated Competition in Health Insurance Markets on Two Sides of the Atlantic. *Journal of Economic Perspectives*, 40(2):4368.

Korattikara, A., Chen, Y., and Welling, M. (2014). Austerity in MCMC Land: Cutting the Metropolis-Hastings Budget. In *Proceedings of the 31st International Conference on Machine Learning*, pages 181–189.

Kroneman, M., Boerma, W., van den Berg, M., Groenewegen, P., de Jong, J., and van Ginneken, E. (2016). Netherlands: Health System Review. *Health Systems in Transition*, 18(2):1–239.

Liu, J. S. (2001). *Monte Carlo Strategies in Scientific Computing*. Springer, New York.

Maire, F., Friel, N., and Alquier, P. (2019). Informed sub-sampling MCMC: approximate Bayesian inference for large datasets. *Statistics and Computing*, 29(3):449–482.

Marone, V. and Sabety, A. (2022). When Should There be Vertical Choice in Health Insurance Markets? *American Economic Review*, 112(1):304–342.

Palmucci, G. and Dague, L. (2014). The Welfare Effects of Banning Risk-Rated Pricing in Health Insurance Markets: Evidence from Chile. Working Paper.

Papaspiliopoulos, O., Roberts, G. O., and Skold, M. (2007). A General Framework for the Parametrization of Hierarchical Models. *Statistical Science*, 22(1):59–73.

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

Polyakova, M. (2016). Regulation of Insurance with Adverse Selection and Switching Costs: Evidence from Medicare Part D. *American Economic Journal: Applied Economics*, 8(3):165–195.

Quiroz, M., Kohn, R., Villani, M., and Tran, M.-N. (2019). Speeding Up MCMC by Efficient Data Subsampling. *Journal of the American Statistical Association*, 114(526):831–843.

Roberts, G. O. and Rosenthal, J. S. (1997). Updating Schemes, Correlation Structure, Blocking and Parameterization for the Gibbs Sampler. *Journal of the Royal Statistical Society: Series B*, 59(2):291–317.

Roberts, G. O. and Rosenthal, J. S. (2001). Optimal scaling for various Metropolis-Hastings algorithms. *Statistical Science*, 16(4):351–367.

Saez, E. and Stantcheva, S. (2016). Generalized Social Marginal Welfare Weights for Optimal Tax Theory. *American Economic Review*, 106(1):24–45.

Sapelli, C. (2004). Risk segmentation and equity in the Chilean mandatory health insurance system. *Social Science & Medicine*, 58(2):259–265.

Starc, A. and Town, R. J. (2020). Externalities and Benefit Design in Health Insurance. *Review of Economic Studies*, 87(6):2827–2858.

Superintendencia de Salud (2012). Bases de datos de cotizantes, planes de salud, tablas de factores y prestaciones bonificadas de isapres. Santiago: Gobierno de Chile.

Superintendencia de Salud (2013). Más de 58 mil millones de pesos acumulan las cuentas de excedentes de los usuarios de Isapre. Comunicado de prensa, Santiago de Chile. Available at `https://web.supersalud.gob.cl/prensa/672/w3-article-8236.html`.

Tanner, M. A. and Wong, W. H. (1987). The Calculation of Posterior Distributions by Data Augmentation. *Journal of the American Statistical Association*, 82(398):528–540.

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. (2009). *Discrete Choice Methods with Simulation*. Cambridge University Press, Cambridge, 2nd edition.

Appendices

A Extra Tables and Figures

Figure A.1: 2013 Income Distribution by Insurance Type (Public vs. Private)

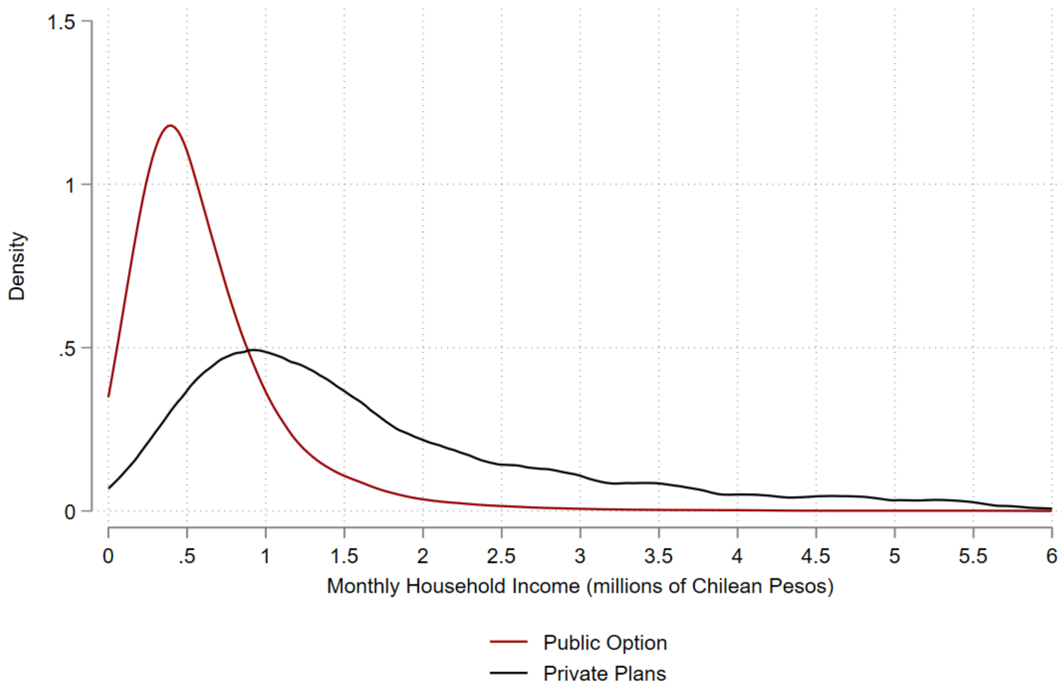

Notes: Kernel estimates of the 2013 income distributions of enrollees by type of insurance: public option vs. private plan. Source: 2013 Chile National Socioeconomic Characterization Survey (CASEN).

Figure A.2: 2013 Age Distribution by Insurance Type (Public vs. Private) and by Gender

(a) Men

(b) Women

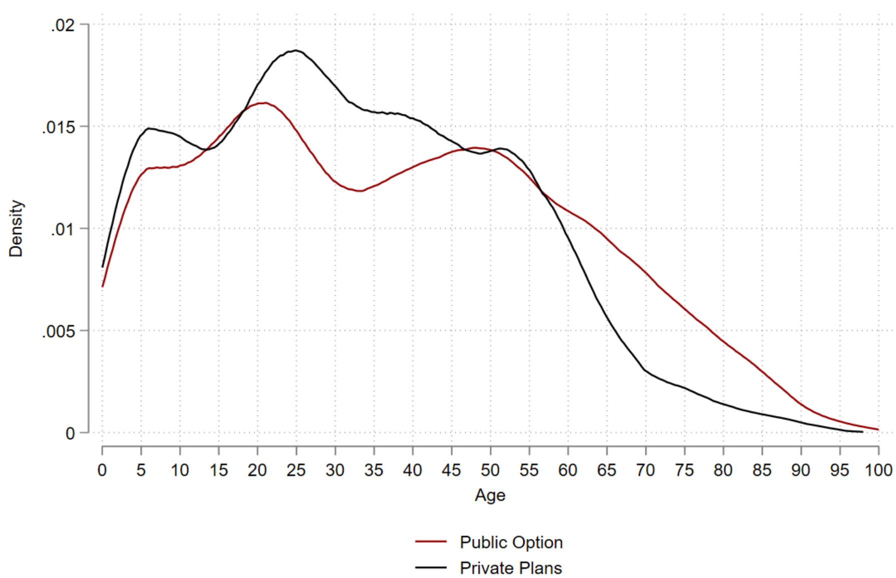

Notes: Kernel estimates of the 2013 age distributions of enrollees by type of insurance — public option vs. private plan — for each gender. Panel (a) shows the age distribution for men and Panel (b) shows the age distribution for women. Source: 2013 Chile National Socioeconomic Characterization Survey (CASEN).

Figure A.3: 2013 Share of Enrollees in Private Market by Self-Assessed Health Status

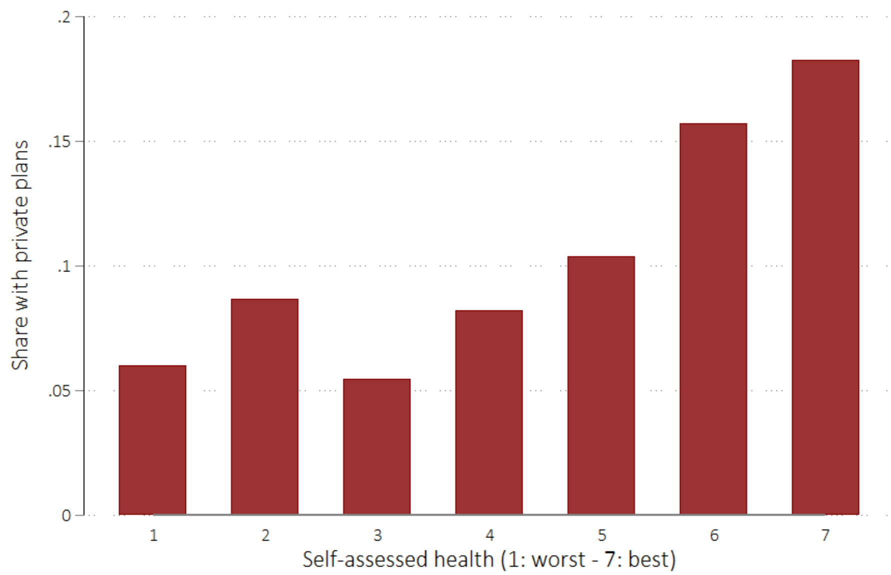

Notes: Share of population with private plans per self-assessed health group. Self-assessed health ranges from 1 (worst health) to 7 (best health). Source: 2013 Chile National Socioeconomic Characterization Survey (CASEN).

Figure A.4: Share Fit by Demographic Group – Men

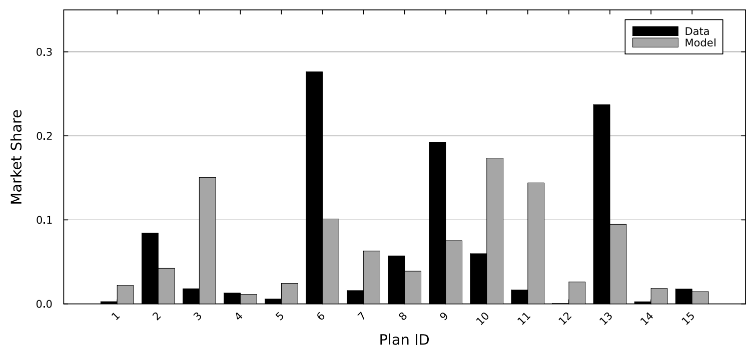

(a) Age: 18-29

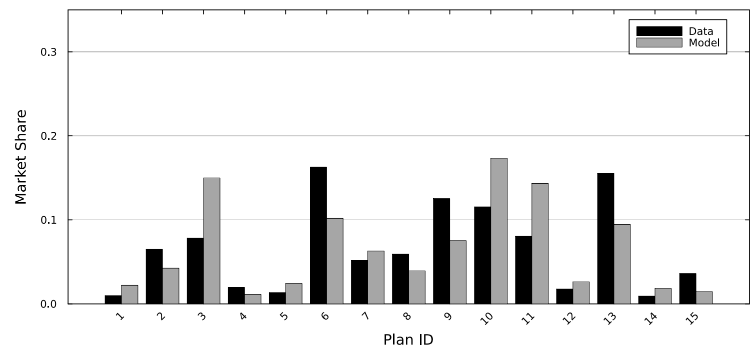

(b) Age: 30-44

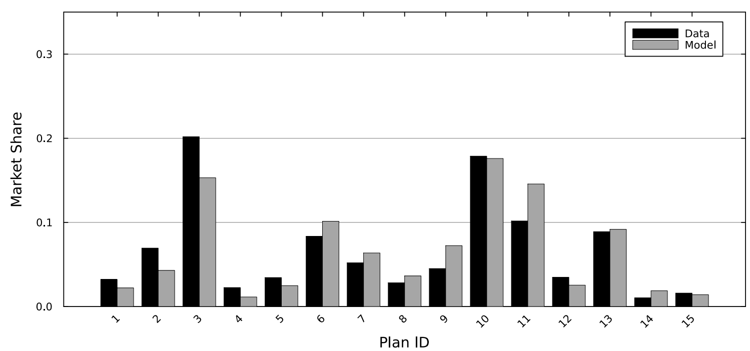

(c) Age: 45-64

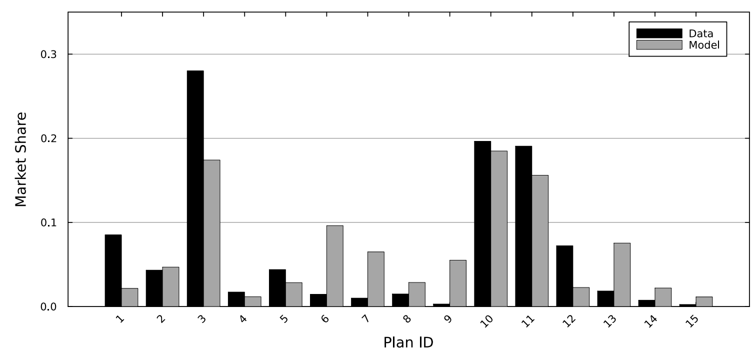

(d) Age: 65 and above

Notes: Observed and model-predicted plan market shares across the 15 plan groups, separately by age group, for men. Bars labeled "Data" are observed shares; "Model" are the shares implied by the estimated model at the posterior means. Each panel corresponds to one age group, as indicated.

Figure A.5: Share Fit by Demographic Group – Women

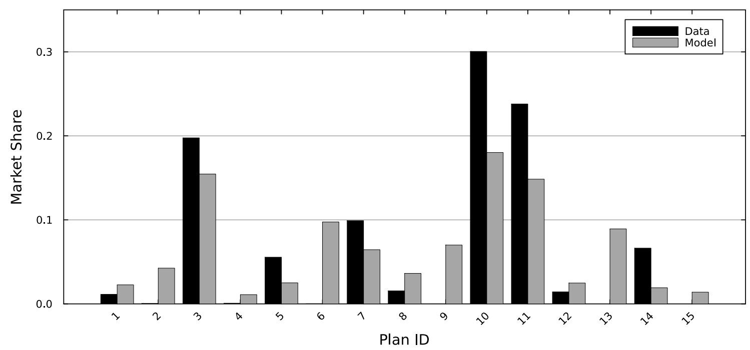

(a) Age: 18-29

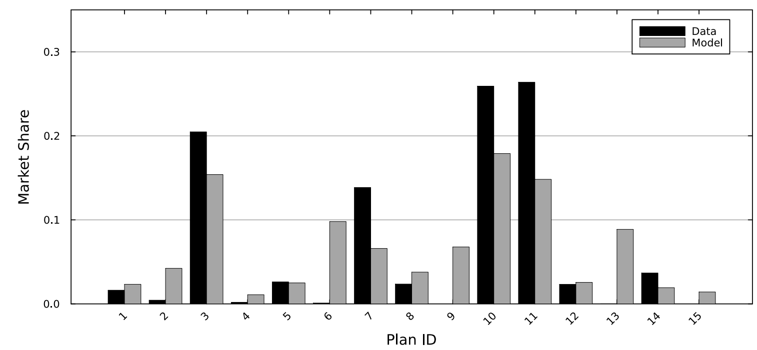

(b) Age: 30-44

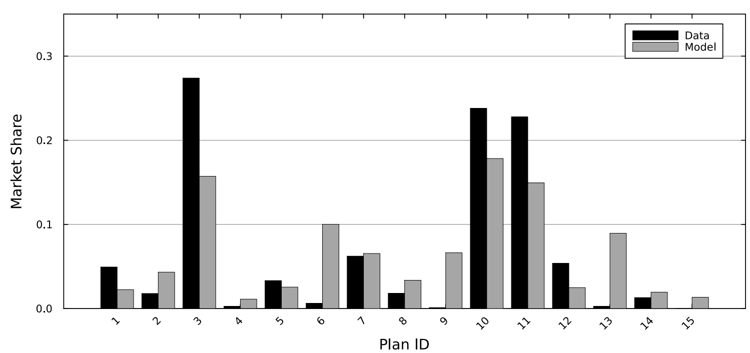

(c) Age: 45-64

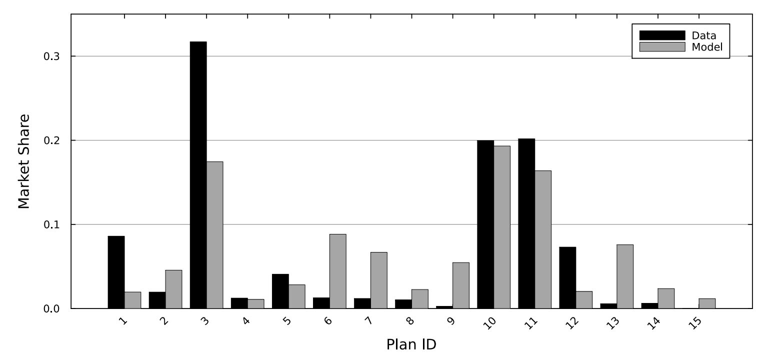

(d) Age: 65 and above

Notes: Observed and model-predicted plan market shares across the 15 plan groups, separately by age group, for women. Bars labeled "Data" are observed shares; "Model" are the shares implied by the estimated model at the posterior means. Each panel corresponds to one age group, as indicated.

Figure A.6: Spending Fit by Demographic Group – Men

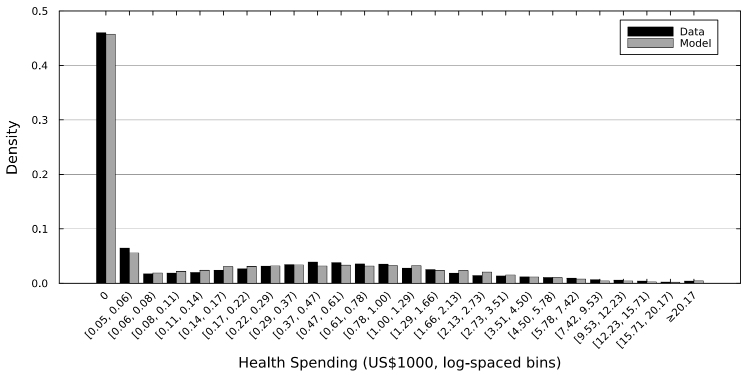

(a) Age: 18-29

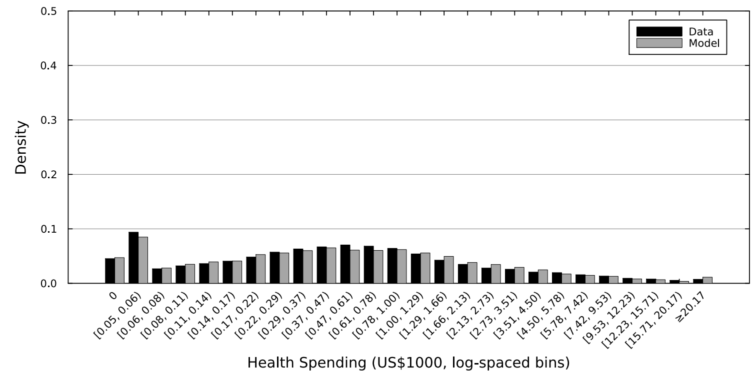

(b) Age: 30-44

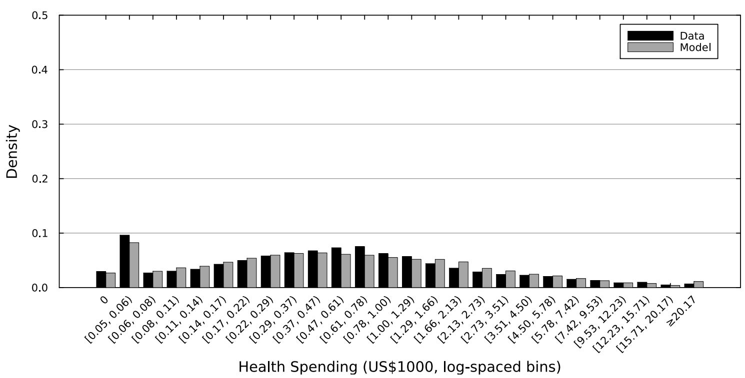

(c) Age: 45-64

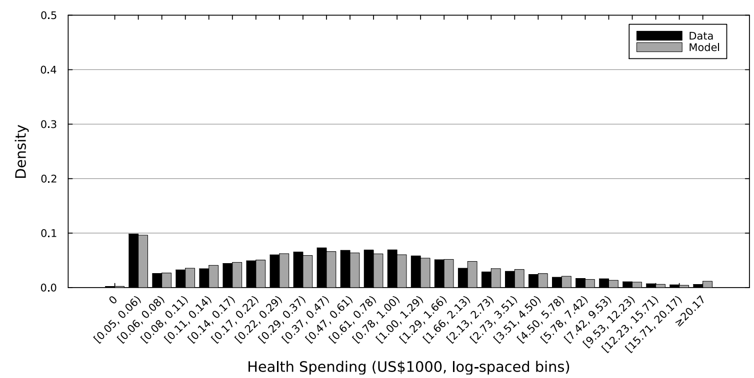

(d) Age: 65 and above

Notes: Observed and model-predicted health spending distributions, separately by age group, for men. For zero spending, the bar reports the observed and model-implied shares; for positive spending, bars report shares within log-spaced spending windows (1000 USD, 2012). "Data" are observed; "Model" are implied by the estimated model at the posterior means. Each panel corresponds to one age group, as indicated..

Figure A.7: Spending Fit by Demographic Group – Women

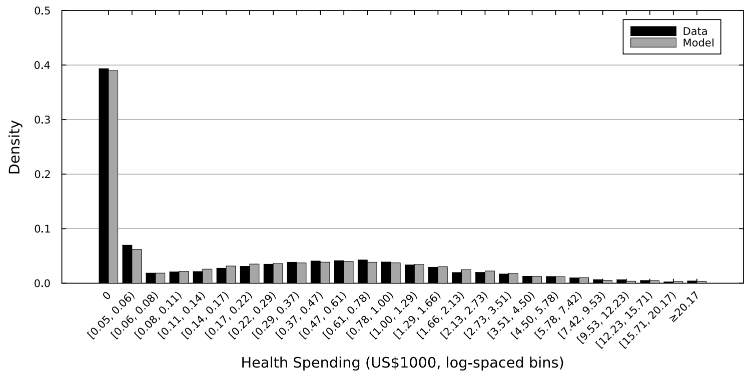

(a) Age: 18-29

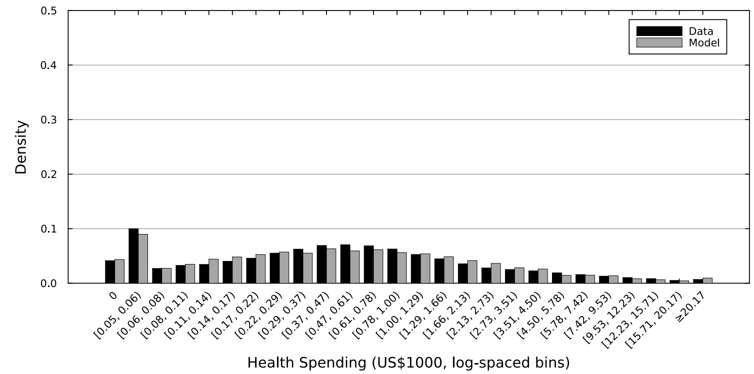

(b) Age: 30-44

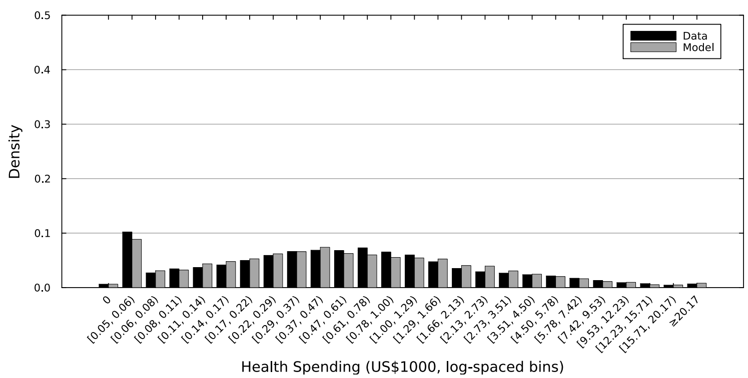

(c) Age: 45-64

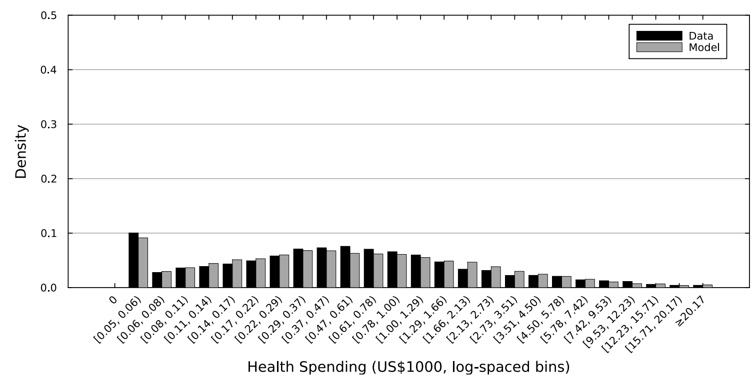

(d) Age: 65 and above

Notes: Observed and model-predicted health spending distributions, separately by age group, for women. For zero spending, the bar reports the observed and model-implied shares; for positive spending, bars report shares within log-spaced spending windows (1000 USD, 2012). "Data" are observed; "Model" are implied by the estimated model at the posterior means. Each panel corresponds to one age group, as indicated..

Table A.1: Sample Construction

Step	Restriction	Category	N	% of Initial
0	Initial *cotizante* file (policyholders, Jan 2012)	—	1,528,236	100.0
Scope restrictions				
1	Workers required to have insurance	Scope	1,391,096	91.0
2	Santiago metropolitan area	Scope	692,578	45.3
3	Single-person contracts (no dependents)	Scope	394,398	25.8
4	Open ISAPREs (five major carriers)	Scope	359,779	23.5
5	Age 18 or above	Scope	213,573	14.0
Data quality restrictions				
6	Has plan ID information	Data Quality	1,527,467	99.9
7	Valid plan info (matched to plans dataset, non-closed)	Data Quality	283,019	18.5
8	Non-fixed-copay plans (model assumes coinsurance)	Data Quality	234,091	15.3
9	Income above minimum wage	Data Quality	213,607	14.0
10	Plans with $\geq$ 1% market share	Data Quality	203,459	13.3
Estimation-specific restrictions				
11	Non-negative implied consumption	Estimation	202,681	13.3
12	Regularity condition satisfied	Estimation	202,662	13.3

Notes: Cumulative number of *cotizantes* (policyholders, not total covered lives) after each restriction; restrictions are applied in the order shown in the underlying code, which groups conceptually-related steps together. "Scope" restrictions define the population of interest — single-person formal-sector workers aged 18 or above in Santiago, enrolled in one of the five major open ISAPREs — rather than constituting sample selection. "Data Quality" restrictions ensure that the model's inputs are well-defined for the retained enrollees; the plan-share filter is implemented in two stages (a 0.5% data-quality threshold in the cleaning pipeline and a stricter 1% threshold in the estimation, but the table reports the combined effect). "Estimation" restrictions are required by the structural framework: non-negative consumption is required by the budget constraint; the regularity condition ensures clean inversion of observed spending into latent health shocks. The largest absolute drops occur at the scope restrictions, particularly the Santiago and single-person restrictions, which together determine the population being studied.

Table A.2: Plan Fixed Effects

Plan ID	Posterior mean (ξ)	Posterior SD	2.5%	97.5%
2	0.6288	0.0381	0.5538	0.7015
3	1.9233	0.0331	1.8578	1.9852
4	-0.6877	0.0521	-0.7942	-0.5879
5	0.0963	0.0419	0.0141	0.1777
6	1.4999	0.0341	1.4323	1.5643
7	1.0409	0.0357	0.9710	1.1093
8	0.5606	0.0395	0.4812	0.6360
9	1.2057	0.0354	1.1351	1.2741
10	2.0596	0.0327	1.9954	2.1206
11	1.8819	0.0332	1.8155	1.9442
12	0.1915	0.0424	0.1069	0.2728
13	1.4456	0.03429	1.3771	1.5104
14	-0.1903	0.0447	-0.2794	-0.1037
15	-0.4216	0.0479	-0.5184	-0.3296

Notes: Posterior summaries of plan-specific fixed effects ξ_j. The reference plan (ID 1) is normalized to $\xi = 0$. Hyperpriors are $\xi_j \sim \mathcal{N}(\mu_\xi, \sigma_\xi^2)$ with hyperparameters jointly estimated.

Table A.3: Welfare Incidence by Income Quartile – 8% Mandate vs. Direct Transfer

		8% Mandate		7% Mandate + 1% tax	
Income quartile	Avg. income	Mean CV	% Better off	Mean CV	% Better off
Q1	10.71	-0.04	54.10	0.41	62.50
Q2	19.68	-0.11	23.30	0.18	43.30
Q3	30.58	-0.28	1.10	-0.24	8.30
Q4	36.96	-0.33	0.60	-0.35	2.30

Notes: Compensating variation by income quartile. "% Better off" is the share of enrollees with positive CV (those for whom the mandate is non-binding or only weakly binding).

B Estimation Details

This appendix provides details on the MCMC estimation procedure: the Jacobian derivation for the spending likelihood, the conditional likelihood components used for each parameter block, the updating scheme, the minibatching strategy, proposal adaptation, prior specifications, and convergence diagnostics.

B.1 Likelihood Components

Before describing the updating scheme, I write out the likelihood components used in each Metropolis-Hastings step. The complete-data likelihood for individual i, conditional on the latent moral hazard parameter ω_i and on the parameter vector Θ, is the product of a plan choice component and a spending component:

$$\mathcal{L}_i(\mathbf{\Theta}, \omega_i) = \mathcal{L}_i^{\text{choice}}(\mathbf{\Theta}, \omega_i) \times \mathcal{L}_i^{\text{spending}}(\mathbf{\Theta}, \omega_i).$$

For each parameter block, only the components that depend on the parameter enter the Metropolis-Hastings ratio.

Plan choice likelihood. Given ω_i, the probability that individual i chooses plan j_i is the multinomial logit

$$s_{i,j_i}(\mathbf{\Theta}, \omega_i) = \frac{\exp\left(\mathbb{E}_{\lambda_i}[U_{ij_i}] + \xi_{j_i}\right)}{\sum_{j'=1}^{J} \exp\left(\mathbb{E}_{\lambda_i}[U_{ij'}] + \xi_{j'}\right)},$$

where $\mathbb{E}_{\lambda_i}[U_{ij}]$ is computed via discretization of the spliced Log-Normal-Pareto distribution (see Appendix B.8). The plan choice likelihood is $\mathcal{L}_i^{\text{choice}}(\mathbf{\Theta}, \omega_i) = s_{i,j_i}(\mathbf{\Theta}, \omega_i)$.

Spending likelihood: two equivalent forms. The spending likelihood appears in two forms depending on which parameters are being updated.

Form 1 (latent-λ form), used when $\{\lambda_i\}$ are held fixed. For the demographic-specific health distribution parameters ($\mu_0(X)$, $\sigma(X)$, $\kappa(X)$, α_μ, δ_μ), the latent health shocks $\{\lambda_i\}$, inferred from the FOC inversion at the start of the iteration, are treated as data. The contribution of individual i to the spending likelihood is then

$$\mathcal{L}_i^{\text{spending,1}}(\boldsymbol{\Theta}) = \begin{cases} \kappa(X_i) & \text{if } \lambda_i = 0, \\ \big(1-\kappa(X_i)\big) \cdot g^{+}\big(\lambda_i \mid X_i, y_i\big) & \text{if } \lambda_i > 0. \end{cases}$$

Here, $g^{+}(\cdot \mid X_i, y_i)$ is the spliced Log-Normal-Pareto density of λ conditional on $\lambda > 0$:

$$g^{+}(\lambda \mid X_i, y_i) = \begin{cases} \phi_{\text{LN}}\big(\lambda; \mu(X_i, y_i), \sigma(X_i)\big) & \text{if } \lambda \in (0, \bar{\lambda}], \\ S_{\text{LN}}\big(\bar{\lambda}; \mu(X_i, y_i), \sigma(X_i)\big) \cdot \dfrac{\theta\, \bar{\lambda}^{\theta}}{\lambda^{\theta+1}} & \text{if } \lambda > \bar{\lambda}, \end{cases}$$

where ϕ_{LN} is the Log-Normal density, S_{LN} is the Log-Normal survival function, and the splice point $\bar{\lambda}$ is fixed at the 85^{th} percentile of observed positive spending. The location parameter is

$$\mu(X_i, y_i) = \mu_0(X_i) + \alpha_\mu \log(y_i / y_{\text{ref}}) + \delta_\mu \mathbb{1}\{y_i \text{ censored}\}.$$

Form 2 (inverted-spending form), used when ω_i varies. For the parameters that control moral hazard (ω_0, α_ω, σ_ω, δ_ω, $\varepsilon_{\omega,i}$), changing the parameter changes ω_i, which changes the inverted health shock $\Lambda(h_i, \omega_i, j_i) = h_i / D_i$ (see subsection B.2). The spending likelihood for individuals with positive spending must therefore be evaluated as

$$\mathcal{L}_i^{\text{spending,2}}(\boldsymbol{\Theta}, \omega_i) = \big(1-\kappa(X_i)\big) \cdot g^{+}\big(\Lambda(h_i, \omega_i, j_i) \mid X_i, y_i\big) \cdot \left| \frac{d\Lambda}{dh} \right|.$$

Form 1 and Form 2 are related by the change-of-variables formula: when λ_i is set equal to $\Lambda(h_i, \omega_i, j_i)$ at the current ω_i, the two forms differ only by the Jacobian $|d\Lambda/dh|$, which does not depend on the parameters being updated when ω_i is fixed (Form 1) but does when ω_i varies (Form 2).

Latent variable. I adopt a non-centered parameterization for the moral hazard parameter, writing $\log \omega_i = \overline{\omega}(y_i) + \sigma_\omega\, \varepsilon_{\omega,i}$ with $\varepsilon_{\omega,i} \sim \mathcal{N}(0,1)$ a priori, where

$$\overline{\omega}(y_i) = \omega_0 + \alpha_\omega \log(y_i / y_{\text{ref}}) + \delta_\omega \mathbb{1}\{y_i \text{ censored}\}.$$

The latent draws are over $\varepsilon_{\omega,i}$ rather than ω_i directly, which improves mixing when σ_ω is small (Papaspiliopoulos et al., 2007).

Identity of the latent shock. For individuals with $h_i > 0$, the latent health shock is pinned down deterministically by the FOC inversion: given the current ω_i, $\lambda_i = \Lambda(h_i, \omega_i, j_i)$. For individuals with $h_i = 0$, the regularity condition ensures $\lambda_i = 0$. Hence the only latent variable that requires MCMC sampling is $\varepsilon_{\omega,i}$. The values $\{\lambda_i\}$ are recomputed from $\{\varepsilon_{\omega,i}\}$ inside the $\varepsilon_{\omega,i}$ update step (subsection B.3), so the demographic blocks see fresh $\{\lambda_i\}$ each iteration.

B.2 Jacobian of the Spending Inversion

Differentiating $\Lambda(h, \omega, j) = h/D$ with respect to h, where $D = 1 + \omega(1 - (1 - c_j)/z)$ and $z = y - p_j - (1 - c_j)h$,

$$\frac{d\Lambda}{dh} = \frac{1}{D} + \frac{h\,\omega\,(1 - c_j)^2}{z^2\,D^2}.$$

Both terms are positive under the regularity condition $y - p_j > 1 - c_j$, which guarantees $z > 0$ and $D > 0$ for all feasible spending levels. Hence the mapping from h to λ is monotone, and the change-of-variables formula applies. The 20 individuals in the sample for whom this regularity condition is violated on the chosen plan are dropped, as discussed in section 4.

B.3 MCMC Updating Scheme

The algorithm cycles through the following blocks at each iteration. For each block, I list the parameters being updated and the relevant likelihood components that enter the Metropolis-Hastings ratio.

Health shock distribution: $\mu_0(X)$, $\log\sigma(X)$, logit$(\kappa(X))$, α_μ. For each demographic group g, I update $\mu_0(g)$, $\log\sigma(g)$, and $\text{logit}(\kappa(g))$ in sequence; the income gradient α_μ is updated using the full sample.

For $\mu_0(g)$, $\log\sigma(g)$, and α_μ, the Metropolis-Hastings ratio uses both the spending likelihood (Form 1) and the plan choice likelihood. Let $\mathcal{I}_g = \{i : X_i = g\}$ and $\mathcal{S}_g^{\text{ch}} \subset \mathcal{I}_g$ denote a random subsample drawn for the plan choice evaluation. The Metropolis-Hastings ratio for $\mu_0(g)$ is

$$r = \min\left\{1,\ \frac{\left[\prod_{i\in\mathcal{I}_g,\,\lambda_i>0} g^+(\lambda_i \mid \mu_0'(g),\cdots)\right]\cdot\left[\prod_{i\in\mathcal{S}_g^{\mathrm{ch}}} s_{i,j_i}(\boldsymbol{\Theta}',\omega_i)\right]^{n_g/n_g^{\mathrm{sub}}}\cdot\pi(\mu_0'(g))}{\left[\prod_{i\in\mathcal{I}_g,\,\lambda_i>0} g^+(\lambda_i \mid \mu_0(g),\cdots)\right]\cdot\left[\prod_{i\in\mathcal{S}_g^{\mathrm{ch}}} s_{i,j_i}(\boldsymbol{\Theta},\omega_i)\right]^{n_g/n_g^{\mathrm{sub}}}\cdot\pi(\mu_0(g))}\right\},$$

where $n_g = |\mathcal{I}_g|$, $n_g^{\mathrm{sub}} = |\mathcal{S}_g^{\mathrm{ch}}|$, π denotes the prior, and Θ' denotes the vector of parameters substituting for the relevant proposed value — in this case, substituting $\mu_0(g)$ by $\mu_0'(g)$. The exponent n_g/n_g^{sub} rescales the subsampled log-likelihood to the full population scale, an unbiased stochastic estimate of the full-sample log-likelihood. The Jacobian and the $\kappa(X_i)$ factor cancel in the spending-likelihood ratio because they do not depend on μ_0. The structure is analogous for $\log\sigma(g)$ and α_μ, with the appropriate likelihood factors and, for $\log\sigma$, a log-Jacobian for the random walk on the log scale. For α_μ, since it does not depend on demographic characteristics, the full sample is used.

For $\mathrm{logit}(\kappa(g))$, both the spending likelihood and the plan choice likelihood enter. The spending likelihood factor is binomial: with $N_g = |\mathcal{I}_g|$ and $N_g^0 = |\{i \in \mathcal{I}_g : \lambda_i = 0\}|$,

$$\mathcal{L}^{\mathrm{spending},1}(\kappa(g)) = \kappa(g)^{N_g^0}\cdot\big(1-\kappa(g)\big)^{N_g-N_g^0}.$$

The plan choice likelihood is included via a numerical shortcut. When κ changes, only the linear blending weights between the zero-shock and positive-shock parts of expected utility change; the λ grid points and their relative weights are unaffected. Concretely, expected utility on plan j for individual i can be written as

$$\mathbb{E}_{\lambda_i}[U_{ij}] = \kappa(X_i)\cdot u_{ij}^{\mathrm{zero}} + \big(1-\kappa(X_i)\big)\cdot \overline{U}_{ij}^{\mathrm{pos}},$$

where $u_{ij}^{\mathrm{zero}} = \log(y_i - p_{ij})$ is the utility under $\lambda = 0$ and $\overline{U}_{ij}^{\mathrm{pos}}$ is the conditional expected utility given $\lambda > 0$. Both objects are independent of κ. I compute them once at the current ω_i and then evaluate the choice likelihood under any κ proposal as a linear blend, avoiding the need to re-discretize the health distribution. With this representation, the Metropolis-Hastings ratio for $\mathrm{logit}(\kappa(g))$ uses the spending likelihood on $\mathcal{I}_g$ and the plan choice likelihood on a subsample $\mathcal{S}_g^{\mathrm{ch}}$ at no additional asymptotic cost beyond the other demographic blocks.

Pareto shape: θ. Only observations with $\lambda_i > \bar{\lambda}$ contribute information to θ. The relevant likelihood factor is

$$\mathcal{L}^{\text{spending}}(\theta) \propto \prod_{i:\lambda_i > \bar{\lambda}} \frac{\theta\, \bar{\lambda}^{\theta}}{\lambda_i^{\theta+1}}.$$

This is combined with a Gamma$(5, 0.5)$ prior, which has mean 2.5 and concentrates the posterior around values that ensure a finite Pareto mean ($\theta > 1$).

Moral hazard: $\omega_0, \alpha_\omega, \sigma_\omega$. I update these three parameters via three sequential component-wise scalar Metropolis-Hastings steps, each conditioning on the current values of the other two. A single subsample $\mathcal{S}^{\text{ch}}$ is drawn at the start of the block and reused across all three steps to amortize the cost of computing expected utilities. For each component $\theta_k \in \{\omega_0, \alpha_\omega, \sigma_\omega\}$, the Metropolis-Hastings ratio is

$$r = \min\left\{1,\ \frac{\left[\prod_{i:h_i>0} \mathcal{L}_i^{\text{spending},2}(\boldsymbol{\Theta}', \omega_i')\right] \cdot \left[\prod_{i\in\mathcal{S}^{\text{ch}}} s_{i,j_i}(\boldsymbol{\Theta}', \omega_i')\right]^{n/n^{\text{sub}}} \cdot \pi(\theta_k')}{\left[\prod_{i:h_i>0} \mathcal{L}_i^{\text{spending},2}(\boldsymbol{\Theta}, \omega_i)\right] \cdot \left[\prod_{i\in\mathcal{S}^{\text{ch}}} s_{i,j_i}(\boldsymbol{\Theta}, \omega_i)\right]^{n/n^{\text{sub}}} \cdot \pi(\theta_k)}\right\},$$

where $\omega_i' = \exp(\overline{\omega}'(y_i) + \sigma_\omega' \varepsilon_{\omega,i})$ uses the proposed parameters. For σ_ω, the random walk is implemented on the log scale, and the Metropolis-Hastings ratio includes the corresponding log-Jacobian. The spending-likelihood factor uses Form 2 because ω_i enters the inverted λ_i and the Jacobian directly. Because these parameters seem to mix more slowly, I sample each of them four times per iteration of the MCMC sampling procedure.

Income-cap shifts: δ_μ, δ_ω. The Chilean tope imponible cap (which is relevant for the mandate computation) censors the income data for part of the observations above it. To capture potential differences in the health and moral hazard distributions at the cap, I include shifts δ_μ and δ_ω in $\mu(X_i, y_i)$ and $\overline{\omega}(y_i)$ respectively for observations that I identify as censored. These are updated using only censored individuals, with structures analogous to the α_μ and ω_0 updates: spending likelihood (Form 1 for δ_μ, Form 2 for δ_ω) plus a subsampled, scaled plan choice likelihood.

Plan fixed effects: ξ_j. Each plan fixed effect ξ_k enters the choice probabilities for all individuals, both directly (numerator of $s_{i,k}$ for those who chose plan k) and through the denominator (for everyone). I update each ξ_k via Metropolis-Hastings using the plan choice likelihood, evaluated on the full sample, plus the hierarchical prior:

$$r = \min\left\{1,\ \frac{\left[\prod_{i=1}^{n} s_{i,j_i}(\mathbf{\Theta}', \omega_i)\right] \cdot \pi(\xi_k' \mid \mu_\xi, \sigma_\xi^2)}{\left[\prod_{i=1}^{n} s_{i,j_i}(\mathbf{\Theta}, \omega_i)\right] \cdot \pi(\xi_k \mid \mu_\xi, \sigma_\xi^2)}\right\}.$$

I normalize one of the ξ_k to zero; thus, effectively, the fixed effects are all reported relative to this normalized fixed effect. The hyperparameters μ_ξ and σ_ξ^2 have closed-form conditional posteriors and are sampled directly via Gibbs steps:

$$\mu_\xi \mid \{\xi_k\}, \sigma_\xi^2 \sim \mathcal{N}\left(\frac{\sigma_\xi^{-2} J \bar{\xi}}{\sigma_{\mu_\xi,0}^{-2} + \sigma_\xi^{-2} J},\ \left(\sigma_{\mu_\xi,0}^{-2} + \sigma_\xi^{-2} J\right)^{-1}\right),$$

$$\sigma_\xi^2 \mid \{\xi_k\}, \mu_\xi \sim \text{InvGamma}\Big(\alpha_0 + J/2,\ \beta_0 + \tfrac{1}{2}\textstyle\sum_{k=1}^{J}(\xi_k - \mu_\xi)^2\Big),$$

where J is the number of plans, $\bar{\xi} = J^{-1}\sum_k \xi_k$, and $(\sigma_{\mu_\xi,0}^2, \alpha_0, \beta_0)$ are the prior hyperparameters reported in Table B.4.

Individual latent variables: $\varepsilon_{\omega,i}$ and λ_i. For each individual selected in the current iteration's minibatch (subsection B.4), I propose $\varepsilon_{\omega,i}' \sim \mathcal{N}(\varepsilon_{\omega,i}, \sigma_{\text{prop},i}^2)$ and accept or reject based on

$$r = \min\left\{1,\ \frac{\mathcal{L}_i^{\text{choice}}(\mathbf{\Theta}, \omega_i') \cdot \mathcal{L}_i^{\text{spending},2}(\mathbf{\Theta}, \omega_i') \cdot \phi(\varepsilon_{\omega,i}')}{\mathcal{L}_i^{\text{choice}}(\mathbf{\Theta}, \omega_i) \cdot \mathcal{L}_i^{\text{spending},2}(\mathbf{\Theta}, \omega_i) \cdot \phi(\varepsilon_{\omega,i})}\right\},$$

where $\omega_i' = \exp(\overline{\omega}(y_i) + \sigma_\omega \varepsilon_{\omega,i}')$ and ϕ is the standard normal density. Whenever a proposal is accepted, λ_i is updated deterministically via $\lambda_i = \Lambda(h_i, \omega_i', j_i)$. For individuals with $h_i = 0$, only the plan choice component enters the Metropolis-Hastings ratio (the spending likelihood reduces to $\kappa(X_i)$, which does not depend on ω_i).

B.4 Subsampling

The algorithm uses two distinct subsampling devices.

Minibatching of individual latent variables. At each iteration, I update the latent variables $\varepsilon_{\omega,i}$ for a random subsample of 15,000 individuals, rather than the full sample of approximately 175,000. This is exact MCMC: the subsampling determines *which* individuals are updated at a given iteration, not *how* their updates are computed. Each selected individual's Metropolis-Hastings ratio uses the exact likelihood, and each individual is updated with positive probability at every iteration, so the chain is ergodic and targets the same stationary distribution as a full-scan Gibbs sampler (Liu, 2001; Roberts and Rosenthal, 1997).

Scaled subsampling for the plan choice likelihood. When updating population parameter blocks (ω block, demographic blocks, income-cap shifts), evaluating the plan choice likelihood for all $n \approx 175{,}000$ individuals at every Metropolis step would dominate the computational cost: each evaluation requires a numerical integration over the λ distribution for every plan in the choice set. I therefore evaluate the plan choice log-likelihood on a random subsample $\mathcal{S}^{\text{ch}}$ of $n^{\text{sub}} = 2{,}000$ individuals and rescale by n/n^{sub} to recover the full-sample log-likelihood:

$$\widehat{\log \mathcal{L}^{\text{choice}}}(\mathbf{\Theta}) = \frac{n}{n^{\text{sub}}} \sum_{i \in \mathcal{S}^{\text{ch}}} \log s_{i,j_i}(\mathbf{\Theta}, \omega_i).$$

This is the simplest form of approximate-MCMC subsampling. Because the rescaled log-likelihood is unbiased on the log scale but not on the likelihood scale (by Jensen's inequality, $\mathbb{E}[\exp(\widehat{\log \mathcal{L}})] \neq \exp(\log \mathcal{L})$), the chain does not satisfy the unbiasedness condition required for pseudo-marginal exactness (Andrieu and Roberts, 2009), and its stationary distribution is a perturbed version of the target posterior, with the perturbation vanishing as $n^{\text{sub}} \to n$. The plain rescaled approximation is widely used in practice when bias-correction methods are computationally infeasible, and is known to perform well under regularity conditions on the per-observation log-likelihood; Maire et al. (2019) analyze it directly and provide bias bounds, and Bardenet et al. (2017) survey the broader literature. Korattikara et al. (2014) develop a related variant with explicit bias bounds, while Quiroz et al. (2019) propose a bias-corrected version that recovers an exact target via control variates and a pseudo-marginal correction. I do not implement these corrections, relying instead on the model's posterior predictive fit (reported in section 5.3) as evidence that the chain has reached a region of parameter space that captures observed market shares and aggregate spending, which is the substantive question of interest. The fit checks demonstrate that the posterior is concentrated in a region of parameter space that reproduces observed data, providing direct evidence that point estimates are reliable. They do not directly bound the variance bias, but the substantive conclusions of the paper rest on point estimates — the welfare decomposition into fiscal and equilibrium channels, the magnitude of the consumer-to-insurer transfer, and the comparison with direct transfers — rather than on the precise width of credible intervals. The spending likelihood, which is cheaper to evaluate because it does not require integrating over λ, is computed on the full sample.

B.5 Proposal Adaptation

Proposal variances for all Metropolis-Hastings steps are adapted during the burn-in period to target acceptance rates between 25% and 45%, in line with standard guidelines (Roberts and Rosenthal, 2001). Specifically, every 100 iterations during the burn-in, I compute the rolling acceptance rate over the most recent 1,000 iterations. If the rate is below 25%, the proposal variance is multiplied by 0.8; if above 45%, by 1.2. After the burn-in period, proposal variances are held fixed.

For the individual latent variables $\varepsilon_{\omega,i}$, proposal variances are adapted individually, since the posterior curvature can vary substantially across individuals — for instance, individuals with high spending have more informative likelihoods and require smaller proposals. For σ_ω and $\sigma(X)$, the random walks are implemented on the log scale, which keeps proposals strictly positive and stabilizes adaptation.

B.6 Prior Specifications

I use weakly informative priors throughout. Table B.4 summarizes the prior distributions.

Table B.4: Prior Specifications

Parameter	Prior	Notes
ω_0	$\mathcal{N}(0, 5^2)$	
α_ω	$\mathcal{N}(0, 1)$	
σ_ω	Half-$\mathcal{N}(0, 2^2) \cap [0.05, 2]$	Truncated; log-scale random walk
δ_ω	$\mathcal{N}(0, 1)$	Censored observations shift
$\mu_0(X)$	$\mathcal{N}(0, 1)$	Per demographic group
α_μ	$\mathcal{N}(0, 1)$	
δ_μ	$\mathcal{N}(0, 1)$	Censored observations shift
$\log \sigma(X)$	$\mathcal{N}(0, 1)$	Log-scale; per group
$\text{logit}(\kappa(X))$	$\mathcal{N}(0, 2^2)$	Logit-scale; per group
θ	$\text{Gamma}(5, 0.5)$	Pareto shape; mean 2.5
ξ_j	$\mathcal{N}(\mu_\xi, \sigma_\xi^2)$	Hierarchical
μ_ξ	$\mathcal{N}(0, 100^2)$	Hyperprior
σ_ξ^2	$\text{InvGamma}(2, 25)$	Hyperprior; $\mathbb{E}[\sigma_\xi] \approx 5$
$\varepsilon_{\omega,i}$	$\mathcal{N}(0, 1)$	Non-centered parameterization

These priors are diffuse relative to the scale of the data, and the posterior is dominated by the likelihood. The truncation is used purely as a numerical device to prevent the sampler from exploring regions where the value of the parameter is unreasonably large or small and that could generate numerical problems. The estimate of the relevant parameter in this case, σ_ω, is not anywhere close to the boundaries, so the estimates are not biased by this procedure.

B.7 Convergence Diagnostics

I run the chain for 75,000 iterations, discarding the first 15,000 as burn-in. Convergence is assessed by visual inspection of trace plots for the population parameters, which show stable mixing after the burn-in period, and by computing the Gelman-Rubin $\hat{R}$ statistic and effective sample size for each parameter. Table 2 reports these diagnostics: all $\hat{R}$ values are below 1.011, and effective sample sizes range from approximately 300 (for the population means ω_0, α_ω, $\mu_0(X)$, and α_μ, which mix more slowly because they affect all individuals through the plan choice integral) to over 14,000 (for the demographic-specific $\kappa(X)$, whose conditional posteriors have very small autocorrelation given the binomial structure of the relevant spending likelihood, which is the one that affects $\kappa(X)$ the most).

For the individual latent variables $\varepsilon_{\omega,i}$, I store the full chain for a random subsample of 500 individuals and verify that their trace plots exhibit adequate mixing. Posterior means of $\varepsilon_{\omega,i}$ are computed as running averages over all post-burn-in iterations in which individual i was selected for updating.[49]

B.8 Discretization of the Health Shock Distribution

Computing the expected utility $\mathbb{E}_{\lambda_i}[U_{ij}]$ requires integrating the second-stage utility $U_{ij}(\lambda)$ against the prior distribution of λ, which is a mixture of a point mass at zero and a continuous spliced Log-Normal-Pareto distribution on $(0, \infty)$. I evaluate this integral by deterministic quadrature, replacing the continuous distribution with a discrete approximation that consists of a zero atom and $n_{\text{points}} = 15$ positive points. The same grid is used in both the estimation and the counterfactual analysis.

The grid is constructed individual-by-individual, since $\mu(X_i, y_i)$ and $\sigma(X_i)$ vary with demographics and income. Given the parameters $(\mu_i, \sigma_i, \kappa_i)$ for individual i, the procedure is as follows.

Zero atom. If $\kappa_i > 0$, the first grid point is $\lambda = 0$ with weight κ_i. The remaining weight $1 - \kappa_i$ is distributed across the positive grid points according to the Log-Normal-Pareto density.

Positive grid: bulk. Roughly 70% of the positive points are placed in the LogNormal bulk on $(0, \bar{\lambda}]$ at quantiles of the LogNormal. Specifically, with $n_{\text{bulk}} = \lfloor 0.7 \cdot n_{\text{points}} \rfloor = 10$ bulk points, the points are

$$\lambda_k^{\text{bulk}} = F^{-1}_{\text{LN}(\mu_i, \sigma_i)}\big(p_k \cdot F_{\text{LN}}(\bar{\lambda})\big), \qquad p_k \in [0.01, 0.99],\ k = 1, \ldots, n_{\text{bulk}},$$

spaced uniformly in the rescaled probability p_k, where F_{LN} is the Log-Normal cumulative density function (CDF) and $\bar{\lambda}$ is the splice point (fixed at the 85th percentile of observed positive spending). The p_k are rescaled by $F_{\text{LN}}(\bar{\lambda})$ so that the bulk grid spans the range $(0, \bar{\lambda}]$ rather than $(0, \infty)$. Each bulk point receives a weight equal to the Log-Normal probability mass between the midpoints of adjacent quantiles, capped above by $F_{\text{LN}}(\bar{\lambda})$.

[49] Because of the minibatching, not every individual is updated at every iteration. Running averages over all post-burn-in iterations remain valid posterior mean estimators; individuals retain their current $\varepsilon_{\omega,i}$ between updates, so the chain visits each coordinate's stationary distribution at the random-scan rate. With a minibatch size of 15,000 and 60,000 post-burn-in iterations, each individual is updated approximately 5,140 times.

Positive grid: Pareto tail. The remaining $n_{\text{tail}} = n_{\text{points}} - n_{\text{bulk}} = 5$ points are placed in the Pareto tail $(\bar{\lambda}, \infty)$. Let $Pr_{\text{tail}} = 1 - F_{\text{LN}}(\bar{\lambda})$ denote the unconditional probability that λ exceeds the splice point. I divide the tail probability mass into n_{tail} equal bins and place a point at the midpoint of each bin under the Pareto CDF:

$$\lambda_k^{\text{tail}} = \bar{\lambda} \cdot (1 - q_k)^{-1/\theta}, \qquad q_k = (k - 0.5)/n_{\text{tail}},\ k = 1, \ldots, n_{\text{tail}},$$

with each tail point receiving equal weight $Pr_{\text{tail}}/n_{\text{tail}}$. This equal-mass scheme places more grid points in the upper tail than a quantile-based scheme would, which improves coverage of catastrophic spending realizations without inflating the total number of points. The weights on the positive grid are normalized to sum to $1 - \kappa_i$ before being combined with the zero atom.

Discussion. The discretization is exact in two limits. As $n_{\text{points}} \to \infty$ with the tail share fixed, the quadrature converges to the continuous integral. As $\sigma_i \to 0$, the Log-Normal collapses to a point mass, and any reasonable quantile-based scheme places all bulk points there. The 70/30 bulk-tail split reflects a trade-off: the bulk dominates the mass and matters most for typical individuals, but the tail dominates expected costs (and therefore insurer pricing in the counterfactual). The results reported in the paper use $n_{\text{points}} = 15$ for computational efficiency.

When $\kappa_i \approx 1$ (degenerate case where $\lambda_i = 0$ with probability one), the grid reduces to a single point at zero. When $\kappa_i \approx 0$, the zero atom is omitted and the grid consists of n_{points} positive points only.

Use in expected utility. Given the grid $\{(\lambda_k, w_k)\}_{k=1}^{n_{\text{points}} + \mathbb{1}\{\kappa_i > 0\}}$, expected utility on plan j is

$$\mathbb{E}_{\lambda_i}[U_{ij}] \approx \sum_k w_k \cdot U_{ij}(\lambda_k),$$

where $U_{ij}(\lambda)$ is computed by solving the second-stage problem at λ using the closed-form FOC inversion described in section 3. The grid (λ_k, w_k) is precomputed and stored for each individual at the start of each MCMC iteration; updates that change μ_i or σ_i trigger a re-discretization, while updates that change only κ_i exploit the linear blend described in the κ block of subsection B.3.

C Theoretical Framework – WTP Increasing in Income

In this appendix, I briefly discuss the effect of an increase in the income mandate when WTP increases with income. Recall that here we are assuming a positive correlation between income and health spending — where richer people are more costly to insure than poorer people — and that health risk increases with income and dominates the decrease in risk aversion, generating an upward-sloping WTP curve. This case is shown in Figure C.8.

Figure C.8: Increasing WTP with positive income-spending correlation

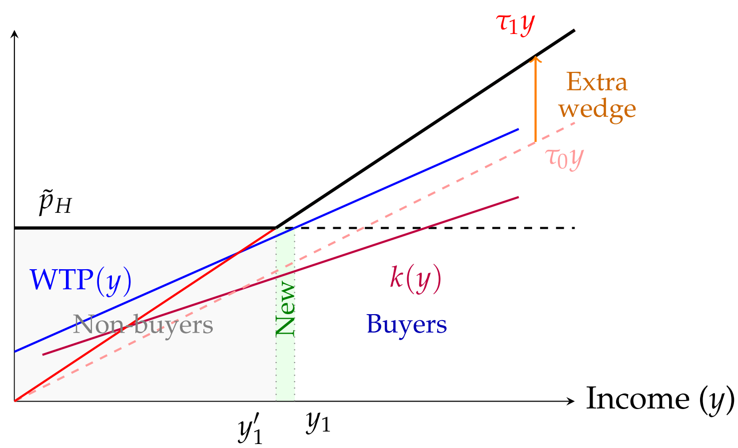

Notes: When the health gradient is strong enough to dominate declining risk aversion, WTP increases in income and non-buyers are concentrated at the low-income (healthy, less costly) end. The cost curve $k(y)$ is also increasing. The mandate increase brings in marginal buyers near y_1 who are at the low-cost end of the buyer distribution, tending to improve the risk pool.

Initially, the first person bound by the mandate had income y_1. As in the other cases, a mandate increase causes this threshold decrease, here to $y_1' < y_1$. These agents with income $y \in [y_1', y_1)$ will now see no premium differential and will prefer to buy the better plan. Since the spending curve is increasing in income, this will improve the risk pool, i.e., the average cost per enrollee will *decrease*. This is coupled with the extra wedge the insurer captures from inframarginal buyers, and both will push the base premium downwards (not shown in the graph for clarity), reinforcing the selection effect.

Notice that the WTP curve is below the base premium in the figure. If the WTP crosses the base premium line at an income level before the mandate binds, and if the change in mandate is not large enough to change that, no extra agents will change their choice due to becoming bound by the mandate, and this part of the effect will be absent. Still, the effect of the extra wedge for inframarginal buyers will exist and the analysis is essentially the same. Also, as in the other cases, the strength of the selection effect depends on the strength of the relationship between spending and income and with the relative sizes of the initial pool of enrollees and the pool of new enrollees. The stronger the spending-income relationship, the higher the difference between the average spending of new buyers and initial enrollees, which strengthens the selection effect; and the larger the pool of new buyers is relative to the initial set of enrollees, the more the average cost for the insurer will change and the stronger the selection effect (conditional on a given difference in average spending).