Insurance Plan Heterogeneity and Hospital Price Dispersion: Theory and Evidence

Wei G. Xiong*

March, 2026

Abstract

How does cost sharing affect hospital prices when prices are negotiated? I show that coinsurance can increase negotiated hospital prices by weakening insurers' incentives to bargain. In Chile's managed-competition individual market, negotiated inpatient prices vary widely across insurers within the same hospital and are higher for insurer-hospital pairs with lower expected coverage, especially at high-priced hospitals. A stylized model combining competitive selection with Nash-in-Nash bargaining delivers this prediction when insurer revenues do not fully internalize consumer surplus. Structural demand estimates indicate that inpatient demand is inelastic at the observed coverage, consistent with the insurer exposure channel driving the pattern.

Keywords: *Hospital Prices; Bargaining; Provider Network*
JEL Codes: *D43; I11; L14*

*Universidad de los Andes, School of Business and Economics, Santiago, Chile (email: wxiong@uandes.cl). I'm grateful to Juan-Pablo Montero and Carlos Noton for their invaluable guidance and advice. I benefited from the advice of Ignacio Cuesta and Benjamin Vatter, and from their joint work with Carlos Noton, which uses similar data. I thank Juan Pablo Atal, Jorge Alé, Sebastian Fleitas, Jeanne Lafortune, Claudio Lucarelli, Alexander Olssen, Nicolas Riquelme and Nicholas Tilipman for helpful comments and suggestions. I also thank Benjamin Vergara and Martin Fuentes for technical assistance. ChatGPT recommended alternative phrasing for part of this work. All errors are my own.

1 Introduction

Many health systems provide health insurance through managed competition (Enthoven, 1988) but differ in the extent of discretion regulators grant insurers to differentiate plan attributes.[1] In an insurance market, product differentiation interacts with selection incentives because "consumers vary not only in how much they are willing to pay for a product but also in how costly they are to the seller" (Einav, Finkelstein, & Mahoney, 2021; Rothschild & Stiglitz, 1976). Hence, regulators actively debate how much plan variety to allow.[2] The economics literature has largely focused on how differentiation affects sorting and welfare within the insurance market (e.g., Einav, Finkelstein, and Cullen (2010) and Marone and Sabety (2022)). This paper shows that plan heterogeneity also has supply-side spillovers: it reshapes insurers' incentives in provider-price negotiations and can amplify negotiated hospital price levels and dispersion.

In Chile's individual health insurance market, five private insurers offer thousands of plans that vary not only in generosity but also in provider-network rules, including tiered designs that reimburse a higher share at a subset of preferred hospitals. In this market, negotiated inpatient prices vary widely across insurers within the same hospital, and dispersion is largest at high-priced "star" hospitals (Ho, 2009). A central contrast with Gowrisankaran, Nevo, and Town (2015) is that coinsurance (i.e., the complement of coverage) does not appear to discipline prices in this setting: insurer-hospital pairs with lower expected coverage negotiate *higher* inpatient prices, both across insurer-years and within insurer-hospital pairs over time. For outpatient care, by contrast, price dispersion is smaller, and the coverage-price relationship is muted.

I interpret the negative coverage-price slope through a simple bargaining mechanism implied by cost sharing. When coinsurance is substantial, a marginal reduction in a negotiated hospital price lowers total spending for patients and insurers, but it lowers an insurer's own spending only in proportion to its expected coverage rate. Whether this translates into higher negotiated prices depends on how much insurers internalize the consumer-side gains of price reductions through premiums and enrollment. In competitive individual markets, revenues need not track consumers' willingness-to-pay closely – especially when a large component of revenues comes from mandated income-based contributions – so insurers may under-internalize these consumer-side benefits. In that case, coinsurance weakens insurer price sensitivity in bargaining and shifts negotiated prices toward hospitals' list-price benchmarks (i.e., the prices that would prevail absent effective insurer bargaining).

To formalize the link from plan heterogeneity to negotiated prices, I develop a stylized model with two markets. In the insurance market, I extend Azevedo and Gottlieb (2017)'s perfect-

[1] In the United States, the Affordable Care Act marketplaces standardize financial generosity into four metal tiers (bronze, silver, gold, and platinum), and Medicare Advantage plans must provide benefits at least as generous as traditional Medicare. In both markets, however, provider network design is only lightly regulated, and plans are typically distinguished by network types such as HMO and PPO. By contrast, several European-style managed-competition systems (e.g., the Netherlands and Belgium) and systems such as Israel's restrict plan variety by imposing a standardized benefit package across insurers.

[2] Examples of regulatory proposals include "network adequacy" for the U.S. ACA exchanges and "unique plan" in Chile.

competition equilibrium with adverse selection to a contract space in which plans differ in both generosity and network quality. In the hospital market, transaction prices are determined by bilateral Nash-in-Nash bargaining (Horn & Wolinsky, 1988). The model yields two implications that discipline the empirical analysis. First, when the high-quality hospital is also high-cost, competitive selection makes it difficult to sustain plans that offer strictly better coverage at the expensive hospital than at the cheaper alternative. Second, holding enrollment and hospital choice fixed, more coinsurance dampens insurers' incentives to bargain for lower prices whenever insurers place relatively more weight on their own medical spending than on consumers' out-of-pocket costs.

The analysis also clarifies why the effect of coinsurance on negotiated prices can differ across institutional settings. In employer-sponsored insurance (ESI), employer incentives may place substantial weight on consumer surplus relative to medical spending (Ho & Lee, 2017; Tilipman, 2022), so coinsurance can strengthen incentives to negotiate for lower prices, as estimated by Gowrisankaran, Nevo, and Town (2015). On the contrary, in a competitive individual market, insurer revenues are only imperfectly tied to consumer willingness-to-pay so that coinsurance may elevate negotiated prices.

Empirically, I analyze industry-wide plan enrollment and itemized claims data from Chile's health regulator. The breadth of plan offerings suggests vigorous competition in differentiated products: each insurer simultaneously sells hundreds to thousands of active plans, spanning a wide range of network rules and coverage rates. At the same time, insurers differ sharply in their enrollment allocations across plan types. Insurers with greater enrollment in Preferred Provider plans exhibit systematically lower expected coverage at high-priced hospitals, consistent with tiering that places expensive star hospitals on a less generous tier. Because claims are reported using standardized service codes that map into a common public fee schedule, I can express negotiated prices as insurer-hospital-year multipliers relative to the same public benchmark. A stark example comes from the two largest high-priced star hospitals: the smaller insurer m_5, which relies less on hospital tiering and therefore has higher expected coverage at these hospitals, pays roughly 25-35 percent lower inpatient prices than the larger insurer m_2, even though their outpatient prices at the same hospitals are almost identical.

The claims data further allow me to separate negotiated prices from realized resource intensity. I document that selection also differs by care setting: conditional on hospital fixed effects, outpatient utilization rises with coverage (adverse selection), while inpatient resource intensity falls with coverage (advantageous selection).[3] The contrast suggests that inpatient risk and hospital preferences are harder to forecast at the time of plan choice than outpatient needs.

These facts matter for welfare because they weaken the insurance function precisely where the stakes are highest. For inpatient care, coverage is lower where prices are higher, and resource intensity is higher among those with lower coverage; together, these patterns imply that severely ill patients face higher out-of-pocket spending when they draw high-priced hospitals, and this exposure

[3]See Einav and Finkelstein (2011) for the characterization of adverse or advantageous selection as a positive or negative correlation between expected loss and insurance coverage.

rises with price dispersion. If hospital choice is even modestly price-sensitive for inpatient care, this induces substitution away from high-priced hospitals, creating a mismatch in which the sickest patients are least likely to use the highest-quality hospitals. For those whose inpatient demand is inelastic, network tiering that shifts inpatient spending to out-of-pocket payments can erode risk protection and generate welfare losses.

Finally, I estimate a hospital demand model that uses individual-level variations in coinsurance, resource intensity and hospital choices to recover consumer price elasticities and the components of insurer surplus that map coverage into bargaining incentives, following the guidance in Berry and Haile (2021). Inpatient demand is inelastic at observed coverage levels, while outpatient demand is substantially more elastic, implying a limited role for demand discipline in inpatient price negotiation. An accounting decomposition shows that, among the objects entering the bargaining first-order condition, insurer cost exposure $\partial C/\partial p$ largely explains the inpatient coverage-price slope at high-priced star hospitals: once $\partial C/\partial p$ is included, the within-hospital relationship between expected coverage and negotiated prices becomes statistically indistinguishable from zero. Consistent with limited internalization of consumer-side gains from lower prices, I show that insurer revenues rely heavily on the income-based mandatory contribution, while voluntary contributions are well below the consumer willingness-to-pay implied by demand – especially for larger, more reputable insurers.

These results have implications for market design in managed-competition systems. Allowing plans to freely combine substantial coinsurance with aggressive tiering can reduce insurers' exposure at expensive hospitals while simultaneously weakening the bargaining forces that would otherwise constrain hospital prices. In settings where inpatient demand is inelastic and ex post needs make substitution difficult, these supply-side effects can amplify both price levels and price dispersion, and can compromise the risk protection that insurance is meant to provide.

Related literature. This paper's main contribution is to fill a gap in the literature on the supply-side effects of tiered provider networks, as highlighted by Handel and Ho (2021).[4] Existing evidence shows that tiered or narrow networks can reallocate utilization toward lower-cost providers and can interact with selection when high-quality hospitals are also high-cost (Ericson & Starc, 2015; Frank et al., 2016; Gruber & McKnight, 2016; Liebman & Panhans, 2021; Prager, 2020; Serna, 2025; Shepard, 2022). Much of this work emphasizes demand-side mechanisms – how consumers value networks and how network rules change provider choice.

A growing literature on narrow networks shows that network exclusion – akin to exclusive dealing in the broader industrial organization economics literature (e.g., Bernheim and Whinston (1998)) – can create bargaining leverage that disciplines prices at *in-network* hospitals (Ghili, 2022; Ho & Lee, 2019, 2023); by contrast, care obtained out of network must be paid out-of-pocket at posted charges rather than negotiated rates. In Chile, however, Exclusive Provider plans

[4]Differential, but non-zero, coverage at preferred versus non-preferred hospitals is sometimes referred to as in-network versus out-of-network coverage, such as in Einav, Finkelstein, and Cullen (2010).

are unpopular, so fully excluding major hospitals would forgo substantial enrollment in a market where insurers compete through plan heterogeneity. Insurer-level networks are therefore effectively complete, largely eliminating exclusion-based leverage. The empirical setting in this paper is close to the preferred pharmacy network, which benefits consumers by lowering out-of-pocket costs at preferred pharmacies (Starc & Swanson, 2021). Hospital tiering similarly reduces out-of-pocket costs at preferred hospitals. My results, however, emphasize that inpatient prices are higher at non-preferred hospitals – precisely where demand is least elastic.

This paper contributes to the classic debate on whether insurance coverage elevates medical spending. Cost sharing is typically justified as a tool to discipline utilization and reduce spending on the grounds of moral hazard (Cutler & Zeckhauser, 2000; Einav & Finkelstein, 2018; Manning et al., 1987; Pauly, 1968). This paper highlights an additional supply-side channel: when provider prices are negotiated, higher coinsurance can weaken insurers' incentives to bargain and thereby raise transaction prices, rather than merely shifting payment responsibility to patients.

This paper also contributes to a small but growing literature on hospital price dispersion. Studies using U.S. claims data document substantial variation both across providers and within providers by payer, but the empirical evidence is often based on a limited set of insurers and a handful of relatively homogeneous procedures (Cooper et al., 2019; Craig, Ericson, & Starc, 2021; Gaynor, Ho, & Town, 2015). In contrast, I use market-wide claims and standardized coding to measure inpatient prices at the admission level.

Finally, this paper connects to a broader literature on vertical markets. Crawford and Yurukoglu (2012) show that downstream product competition shapes downstream-upstream bargaining in the cable television market. In health care, coinsurance adds an additional margin: consumers pay not only insurance premiums but also a share of provider bills, creating a "triangle" structure in which bargaining outcomes depend on both insurer cost exposure and consumer out-of-pocket exposure.

Overview. The paper proceeds as follows. Section 2 presents the theoretical model. Section 3 describes the institutional setting, data and documents the empirical patterns. Section 4 estimates hospital demand and links the estimated primitives to negotiated prices, and Section 5 discusses limitations and policy implications.

2 Theoretical framework

In this section, I introduce a stylized model that captures the key aspects of the two health care markets. The insurance market is competitive, while hospitals are vertically and horizontally differentiated. Central to both markets is a continuum of consumers with heterogeneous, non-contractible preferences over hospitals. This heterogeneity generates positive hospital markups through imperfect substitution and creates selection incentives in insurance: by tiering coverage across hospitals, insurers can screen consumers on hospital preferences, as their hospital choice has differential cost implications for the insurer.

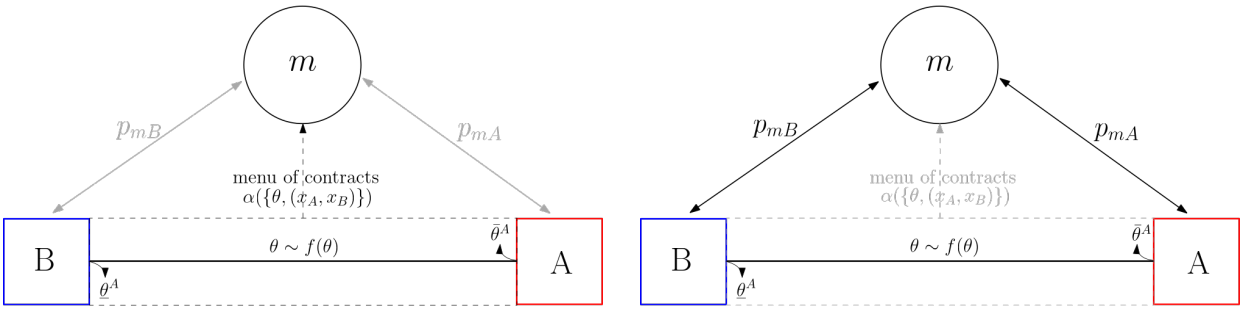
2.1 A model of two markets

Consider an economy with a set of insurers $m \in \mathcal{M}$, two hospitals $A, B \in \mathcal{H}$, and a continuum of consumers. Each m is a multi-product firm offering health insurance plans (i.e., contracts). Hospitals are single-product firms producing substitutes. Hospital A is of higher quality, with common valuations $v_A > v_B > 0$, and a higher marginal cost $mc_A > mc_B$ (to focus on the asymmetric-price case). Consumers differ in (i) θ^A proximity to hospital A (and $-\theta^A$ is interpretable as distance to A in a Hotelling (1929) formulation)[5] and (ii) other dimensions θ^X such as risk and risk attitudes. Let $\theta = (\theta^A, \theta^X)$ be drawn from a joint distribution $f(\theta)$ on a compact support. Insurers only know of $f(\theta)$ because consumer type θ is private information.

Assumption 1 (Taste-risk Independence)**.** *The distribution of θ^A is orthogonal to the distribution of θ^X. Formally, $f(\theta^A) \perp f(\theta^X)$.*

Hospital proximity θ^A affects which hospital is chosen conditional on illness, while θ^X affects the probability and intensity of illness; the two dimensions are independent. This is a simplifying assumption for tractability that will be relaxed in the empirical analysis.[6]

Figure 1 illustrates the two markets with a representative insurer m. A contract specifies hospital-contingent coverage $x \equiv (x_A, x_B) \in [0,1]^2$, where x_h is the insurer's share of the bill if hospital h is chosen. Premiums are denoted $t(x)$. An allocation $\alpha(\{\theta, x\})$ is a measure of θ types purchasing contract x.



(a) Insurance market (b) Hospital market

Figure 1: Theoretical model

Notes: Consumers have common valuations for hospitals A and B such that $v_A > v_B > 0$, while further differ in horizontal hospital preference θ^A. p_{mA} and p_{mB} are negotiated unit prices. A contract of insurance plan is defined as a pair of hospital-specific coverage (x_A, x_B); $\alpha(\cdot)$ denotes plan allocation.

The timing of the model is as follows:

- Stage 1. Each year, insurers post plan menus and negotiate transaction unit prices with hospitals, conditional on the menu and anticipated enrollment allocation. Consumers choose

[5] Empirically, proximity may arise from spatial distance or other idiosyncratic taste, such as personal connection, peer recommendation and so on.

[6] Assumption 1 rules out the correlation that patients who prefer higher-quality hospitals may also be sicker (i.e., riskier), which is reasonable in a context of unregulated contract space. It differs from Veiga and Weyl (2016)'s model's setup, in which a single product attribute must sort consumers of multidimensional types.

plans given premiums and anticipated out-of-pocket costs. Equilibrium is a fixed point at which negotiated prices are consistent with the demand implied by plan and hospital choices.

- Stage 2. Risk realizes. Consumers choose hospitals given negotiated prices and plan-specific coinsurance.

Assumption 2 (Unit Demand)**.** *Both insurance demand and hospital demand are inelastic unit demands.*

Inelastic insurance demand assumes away market unraveling at the extensive margin, in which low-risk consumers do not participate (Akerlof, 1970). Inelastic hospital demand assumes away ex post moral hazard, where patients make extensive or intensive margin choices based on insurance coverage (Pauly, 1968). This simplifying assumption will also be relaxed in the empirical analysis.

Under Assumption 2, I specify consumer preferences in a backward order. In Stage 2, conditional on illness, a consumer of type θ who holds contract $x = (x_A, x_B)$ from insurer m chooses a hospital to maximize ex post utility

$$U_H(h; x, p_{mA}, p_{mB}, \theta) = \begin{cases} v_A + \theta^A - (1 - x_A)p_{mA} & \text{if } h = A, \\ v_B - (1 - x_B)p_{mB} & \text{if } h = B, \end{cases}$$

where v_h is baseline hospital value and θ^A captures proximity or tastes toward hospital A. The consumer chooses hospital A if and only if

$$\theta^A \geq \hat{\theta}^A(x) \equiv (1 - x_A)p_{mA} - (1 - x_B)p_{mB} - (v_A - v_B).$$

To rule out corner solutions with zero or unit market share, I assume an interior-choice condition: for the contracts and price vectors considered, the cutoff lies in the support of tastes, $\underline{\theta}^A < \hat{\theta}^A(x) < \bar{\theta}^A$. Equivalently, the vertical quality gap $v_A - v_B$ is not so large that hospital choice becomes degenerate. Since $p_{mA} > p_{mB}$, any increase in the relative coverage tilt x_A/x_B lowers the out-of-pocket price of A relative to B and therefore (weakly) increases the probability of choosing hospital A when ill.

In Stage 1, the type θ with initial income Y anticipates a probability π of falling ill and chooses contract $x = (x_A, x_B)$ to obtain an ex ante expected utility

$$U_I(x, t; \theta) = \pi u(Y - t - OOP) + (1 - \pi)u(Y - t),$$

where π measure risks and the curvature of the function $u(\cdot)$ measures risk aversion, both captured by θ^X. And

$$OOP = (1 - x_A)p_{mA}\sigma_A(x, \theta^A) + (1 - x_B)p_{mB}\sigma_B(x, \theta^A),$$

where σ_h is the type's second-stage choice probability of hospital $h \in \{A, B\}$ given her contract x and taste θ^A. Under Assumption 2, $\sigma_A + \sigma_B = 1$.

Definition 1 (Generosity)**.** *Given unit hospital demand, a contract's generosity is summarized by* $g(x) \equiv \max\{x_A, x_B\}$*, the maximum coverage rate available across hospitals.*

Definition 2 (Network Quality)**.** *Given the common valuations $v_A > v_B$, network quality is summarized by $q(x) \equiv x_A/x_B$. Specifically, there are three network types: Free Choice (FC) contracts satisfy $q = 1$, Preferred Provider (PP) contracts satisfy $q \neq 1$ (tiering), and Exclusive Provider (EP) contracts are corner cases with $x_A = 0$ or $x_B = 0$.*[7]

Adverse selection operates on both generosity and network quality dimensions: higher θ^X consumers purchase more generous coverage and are costlier to insure (Akerlof, 1970; Rothschild & Stiglitz, 1976); higher θ^A consumers purchase higher coverage for hospital A and are more likely to use hospital A's expensive services (Shepard, 2022). Hence, a plan's expected cost depends on the types that choose it. Under Assumption 1, an insurer can separately screen (Mussa & Rosen, 1978) for consumer heterogeneity over θ^X and θ^A with plan generosity and network quality.

Lemma 1 (Single-crossing in generosity and network quality)**.** *Fix negotiated prices (p_{mA}, p_{mB}) and consider ex ante utility $U_I(x, t; \theta)$ with $\theta = (\theta^X, \theta^A)$.*

(i) ***(Generosity).*** *Fix a network quality q. Among contracts with $x_A/x_B = q$, for any two contracts x, x' with $g(x') > g(x)$, the difference $U_I(x', t; \theta) - U_I(x, t; \theta)$ has single crossing in θ^X. Equivalently, holding network quality fixed, U_I has increasing differences in (θ^X, g).*

(ii) ***(Network quality).*** *Fix a generosity level g. Among contracts with $\max\{x_A, x_B\} = g$, the difference $U_I(x', t; \theta) - U_I(x, t; \theta)$ has single crossing in θ^A whenever $q' = x'_A/x'_B > q = x_A/x_B$. Equivalently, holding generosity fixed, U_I has increasing differences in (θ^A, q).*

Lemma 1 follows because higher coverage reduces illness-state out-of-pocket spending, while a higher x_A/x_B lowers the relative out-of-pocket price of A, which benefits types more likely to choose A.

2.2 Aggregate demand and a two-market equilibrium

The average cost of a contract $x = (x_A, x_B)$ is

$$c(x, \theta) = \left[x_A p_{mA} s_A(x, \hat{\theta}^A) + x_B p_{mB} s_B(x, \hat{\theta}^A)\right] \bar{\pi}_x, \tag{1}$$

where $\bar{\pi}_x$ is the average probability of illness and $s_A(x, \hat{\theta}^A)$, $s_B(x, \hat{\theta}^A)$ the market shares of hospitals A and B among the N_x consumers enrolled in the plan x. Aggregating across all plans, I denote m's total patient volume at hospital h as

$$Q_{mh} \equiv \sum_{x \in [0,1]^2} s_{mh}(x, \hat{\theta}^A) \bar{\pi}_x N_x. \tag{2}$$

[7] Classification of network policy adopts names known in the U.S. as PPOs and EPOs. FC plans do not have a direct equivalent in the U.S.; they can be considered as broad-network PPOs with all available providers on the preferred tier. In Chile, FC plans are called "plan libre elección", PP plans are called "plan preferente", and EP plans are called "plan cerrado".

The multi-plan m's total cost on its network is therefore

$$C_{m,\mathcal{H}} = \sum_{h \in \{A,B\}} \bar{x}_{mh} p_{mh} Q_{mh} \tag{3}$$

where $\bar{x}_{mh}$ is the average share of the bill that m pays to h.

Each hospital h's profit from m is[8]

$$\Pi_{h,m} = (p_{mh} - mc_h) Q_{mh} \tag{4}$$

and m's enrollees' total ex post consumer surplus is

$$W_{m,\mathcal{H}} = \sum_{h \in \{A,B\}} \Big(v_{mh} - (1 - \bar{x}_{mh}) p_{mh} \Big) Q_{mh} + \Gamma(Q_{mA}), \tag{5}$$

where $\Gamma(Q_{mA})$ collects the Hotelling proximity component of surplus and depends on hospital choice only through the share choosing hospital A.[9]

In stage 1, consumers purchase health insurance for two reasons. One is to reduce uncertainty (Ψ_m); the other is to maximize ex post surplus conditional on risk realization, given hospital preferences ($W_{m,\mathcal{H}}$). Under assumption 1, consumers' ex ante willingness-to-pay (WTP) for insurance is additively separable as $\Psi_m + W_{m,\mathcal{H}}$. In the spirit of Gowrisankaran, Nevo, and Town (2015), I define an insurer m's surplus as

$$V_{m,\mathcal{H}} = \bar{R}_m + \tau \left(\Psi_m + W_{m,\mathcal{H}} \right) - C_{m,\mathcal{H}} \tag{6}$$

where $\bar{R}_m$ is the insurer's fixed revenue independent of consumer WTP, and $\tau \in \mathbb{R}_+$ indicates the extent to which m's revenue is proportional to consumer WTP.[10] Equivalent to its weight on consumer WTP relative to its costs, τ measures the marginal dollar of insurer revenue generated by a one-dollar increase in consumer WTP, capturing how tightly premium and enrollment responses translate consumer-side gains from lower hospital prices into insurer surplus.

Azevedo and Gottlieb (2017) prove that equilibrium exists in a competitive insurance market with free entry of contracts and endogenous contract characteristics. I define a joint equilibrium of the two markets as (i) perfect competition with adverse selection in the insurance market (Azevedo & Gottlieb, 2017) and (ii) Nash-in-Nash negotiation in the hospital market (Horn & Wolinsky, 1988).

Definition 3 (Competitive Selection-Bargaining Equilibrium)**.** *An equilibrium consists of (i) a menu of contracts and premiums* $\{x, t^*(x)\}$*, (ii) an allocation* $\alpha^*(\{\theta, x\})$ *mapping consumer types to contracts, and (iii) negotiated prices* $\{p^*_{mh}\}_{m \in \mathcal{M}, h \in \mathcal{H}}$ *such that:*

[8] The assumption of a constant marginal cost assumes away capacity constraint concerns, as a static model captures long-run equilibrium.

[9] For example, if assuming uniform distribution $\theta^A \sim U[\underline{\theta}^A, \bar{\theta}^A]$, $\Gamma(Q_{mA}) = \sum_{x \in [0,1]^2} \frac{1}{2}[(\bar{\theta}^A)^2 - (\hat{\theta}^A(x))^2] \bar{x}_x N_x$.

[10] Administrative costs are commonly assumed to be constant in the literature and omitted.

1. ***Consumers select contracts optimally.*** *For almost every* θ*,*

$$\alpha^*(\{\theta, x\}) \in argmax_{(x_A, x_B)} U_I(x; t^*(x), \theta)$$

2. ***Zero profit on traded contracts.*** *For any contract with positive sales that* $x) \in Im(\alpha^*)$*,*

$$t^*(x) = \mathbb{E}[c(x, \theta)|\alpha^*]$$

3. ***Free entry.*** *For any contract* $(x') \notin Im(\alpha^*)$ *that is not traded, there exists* (θ, x) *in the support of* α^* *such that*

$$U_I(x; t^*(x), \theta) = U_I(x'; t^*(x'), \theta) \quad and \quad c(x', \theta) \geq t(x')$$

where consumer θ *makes the cost of* x' *weakly higher than its competitive price.*

4. ***Nash-in-Nash Bargaining.*** *For each pair* $(m, h) \in \mathcal{M} \times \mathcal{H}$*,* p^*_{mh} *maximizes the Nash product of their respective gain-from-trade (GFT):*

$$\max_{p_{mh}} [\underbrace{V_{m,\mathcal{H}} - V_{m,\mathcal{H}\setminus h}}_{GFT_{m,h}}]^{b_m} [\underbrace{\Pi_{h,\mathcal{M}} - \Pi_{h,\mathcal{M}\setminus m}}_{GFT_{h,m}}]^{b_h} \tag{7}$$

holding the other negotiated prices $\boldsymbol{p}_{-mh}$ *fixed.*[11] *Nash bargaining weights* $b_m, b_h \in (0, 1)$ *are normalized so that* $b_m + b_h = 1$*.*

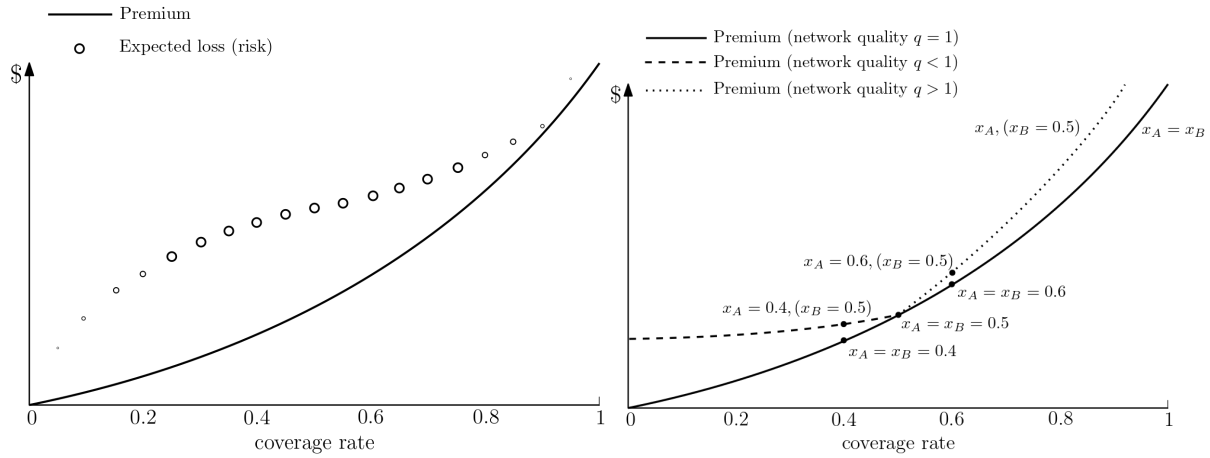
2.3 Equilibrium

In this subsection, I first characterize each market holding the other market's key objects fixed: the insurance equilibrium takes negotiated hospital prices as given, while the hospital bargaining problem takes enrollment and expected coverage $\{\bar{x}_{mh}\}_{h\in\{A,B\}}$ as given. I then describe the feedback between the two markets when these objects are jointly determined.

2.3.1 The insurance market

Figure 2(a) replicates the simulated competitive equilibrium in Azevedo and Gottlieb (2017) for linear contracts that cover between 0 and 100 percent of expenditures. Adverse selection is implied by higher-risk consumers purchasing more generous coverage. Under free entry and zero-profit pricing for traded contracts, expected costs increase with coverage, so the equilibrium premium schedule is convex.

[11] The assumption is valid when there are substitute firms in both the insurance market and the hospital market. See Collard-Wexler, Gowrisankaran, and Lee (2019) for the micro-foundation of Nash-in-Nash bargaining and its application to the health care markets.



(a) Replication of Azevedo and Gottlieb (2017)'s simulated equilibrium

(b) Premium schedule with network quality

Figure 2: Insurance market equilibrium with adverse selection

The contracts in Azevedo and Gottlieb (2017) correspond to FC plans in the present model, with network quality $q \equiv x_A/x_B = 1$. Allowing coverage to vary by hospital introduces an additional contract dimension – network quality – that interacts with selection incentives when hospital A is both higher quality and higher priced ($p_{mA} > p_{mB}$). The next result characterizes which tiering patterns can be sustained under perfect competition.

Proposition 1 (No "A-preferred" tiering under perfect competition)**.** *Suppose $p_{mA} > p_{mB}$. Under free entry and zero-profit pricing for traded contracts, no Preferred Provider (PP) contract with $x_A > x_B$ is sustained: an entrant can offer a Free Choice (FC) contract that attracts a subset of lower-cost types at a lower premium, causing the $x_A > x_B$ contract to unravel.*

The proof is in the Appendix. Figure 2(b) illustrates the intuition. Start from an FC plan with $(x_A, x_B) = (0.5, 0.5)$. Consider perturbations that change network quality by tilting coverage toward one hospital. By Lemma 1 (single-crossing), higher-θ^A consumers select into contracts with higher network quality q.

- **"B-preferred" tiering ($q < 1$).** If $x_A < x_B$, tiering reduces coverage at the expensive hospital A while holding generosity fixed at $g = \max\{x_A, x_B\} = x_B$. For example, the PP plan $(0.4, 0.5)$ has the same generosity as $(0.5, 0.5)$ but makes hospital A more expensive out of pocket. It therefore attracts fewer high-θ^A consumers (who are costlier because they use A more often), lowering expected costs relative to the FC plan. Its premium, therefore, lies between that of the FC plan $(0.5, 0.5)$ and that of the lower-generosity FC plan $(0.4, 0.4)$.

- **"A-preferred" tiering ($q > 1$).** If $x_A > x_B$, the contract both increases generosity (since $g = x_A$) and tilts coverage toward the expensive hospital A. Relative to the FC plan $(0.6, 0.6)$ with the same generosity, the PP plan $(0.6, 0.5)$ attracts a more A-oriented (higher-θ^A) pool of enrollees. Because these enrollees are more likely to use the expensive hospital A when

ill and because $x_A > x_B$ further tilts utilization toward A, the PP contract has a higher expected cost than the corresponding FC contract despite delivering a weakly lower expected utility.

Proposition 1 implies that, in a competitive equilibrium with $p_{mA} > p_{mB}$, traded PP plans must place hospital A on the non-preferred tier: $x_A < x_B$.

Remark 1. *The competitive benchmark in this section isolates the selection incentives created by $p_{mA} > p_{mB}$. With imperfect competition or frictions in plan choice, insurers may sustain "A-preferred" designs through markups or behavioral demand.*

2.3.2 The hospital market

Insurer m's expected coverage rate $\bar{x}_{mh}$ is determined by the allocation of its enrollees across plans and by hospital choices. To isolate how coinsurance affects bargaining outcomes, I assume away the reverse impact by holding fixed expected coverage throughout this subsection:

$$\frac{\partial \bar{x}_{mh'}}{\partial p_{mh}} = 0 \qquad \text{for all } h, h' \in \{A, B\}.$$

Via cost sharing, both the insurer and consumers are exposed to negotiated prices. The effect of a price change can be decomposed into a *direct exposure* channel, which measures how a price change affects payments holding hospital choice fixed, and a *patient-steering* channel, which measures how a price change shifts hospital choice. Differentiating (3) with respect to p_{mh} yields

$$\frac{\partial C_{m,\mathcal{H}}}{\partial p_{mh}} = \underbrace{\bar{x}_{mh} Q_{mh}}_{\text{direct exposure}} + \underbrace{\sum_{h' \in \{A,B\}} \bar{x}_{mh'} p_{mh'} \frac{\partial Q_{mh'}}{\partial p_{mh}}}_{\text{patient steering}}. \tag{8}$$

Differentiating (5) with respect to p_{mh} gives

$$\frac{\partial W_{m,\mathcal{H}}}{\partial p_{mh}} = \underbrace{-(1 - \bar{x}_{mh}) Q_{mh}}_{\text{direct exposure}} + \underbrace{\sum_{h' \in \{A,B\}} \Big(v_{h'} - (1 - \bar{x}_{mh'}) p_{mh'} \Big) \frac{\partial Q_{mh'}}{\partial p_{mh}}}_{\text{patient steering}} + \underbrace{\Gamma'(Q_{mA}) \frac{\partial Q_{mA}}{\partial p_{mh}}}_{\text{taste term}}. \tag{9}$$

Insurer surplus is $V_{m,\mathcal{H}} = \bar{R}_m + \tau\big(W_{m,\mathcal{H}} + \Psi_m\big) - C_{m,\mathcal{H}}$, so

$$\frac{\partial V_{m,\mathcal{H}}}{\partial p_{mh}} = \tau \frac{\partial W_{m,\mathcal{H}}}{\partial p_{mh}} - \frac{\partial C_{m,\mathcal{H}}}{\partial p_{mh}}. \tag{10}$$

Direct exposure. To focus on the direct exposure channel, I hold hospital choice fixed:[12]

$$\frac{\partial Q_{mh'}}{\partial p_{mh}} = 0 \qquad \text{for all } h, h' \in \{A, B\}.$$

A one-unit change in p_{mh} applies to Q_{mh} units of care. The insurer pays a share $\bar{x}_{mh}$, and consumers pay the complement $1 - \bar{x}_{mh}$, so negotiated prices affect insurer spending and consumer surplus through *direct exposure*. A marginal price reduction generates total resource savings Q_{mh}, but only the fraction $\bar{x}_{mh}$ accrues to the insurer; the remaining fraction accrues to consumers as lower out-of-pocket spending. This creates a positive bargaining externality.

Proposition 2 (Coinsurance and price exposure)**.** *Let insurer surplus be* $V_{m,\mathcal{H}} = \bar{R}_m + \tau(W_{m,\mathcal{H}} + \Psi_m) - C_{m,\mathcal{H}}$. *Holding enrollment and hospital choice fixed so that* $\partial \bar{x}_{mh'}/\partial p_{mh} = 0$ *and* $\partial Q_{mh'}/\partial p_{mh} = 0$ *for* $h' \in \{A, B\}$*, the insurer's marginal gain from reducing* p_{mh} *is*

$$-\frac{\partial V_{m,\mathcal{H}}}{\partial p_{mh}} = \big(\tau + (1-\tau)\bar{x}_{mh}\big) Q_{mh}.$$

Moreover,

$$\frac{\partial}{\partial \bar{x}_{mh}} \left(-\frac{\partial V_{m,\mathcal{H}}}{\partial p_{mh}} \right) = (1-\tau) Q_{mh}.$$

Thus, greater coinsurance (lower $\bar{x}_{mh}$*) strengthens price sensitivity if* $\tau > 1$ *and weakens it if* $\tau < 1$.

A brief proof is in the Appendix. Intuitively, τ measures the extent to which the insurer internalizes the consumer portion of bargaining gains through revenue. If $\tau < 1$, the insurer underweights consumer out-of-pocket savings, so higher coinsurance reduces the insurer's perceived benefit from bargaining down prices. The opposite occurs if $\tau > 1$.[13]

Negotiated prices. I now define the disagreement payoffs and characterize negotiated prices. Under managed competition, consumers may switch insurers (or pay out of pocket) to access a given hospital. Let $\tilde{Q}_{mh} \in [0, Q_{mh}]$ denote the quantity of insurer m's enrollees whom hospital h would *lose* in the event of disagreement with insurer m. The remaining volume, $Q_{mh} - \tilde{Q}_{mh}$, are "recapturable" (Ho & Lee, 2017). The volume *at stake* for hospital h in negotiations with insurer m is therefore $\tilde{Q}_{mh}$. I assume $\tilde{Q}_{mh}$ is locally independent of p_{mh} along the equilibrium path.

Hospital h's gains from trade (GFT) from reaching an agreement with insurer m are

$$GFT_{h,m} \equiv \Pi_{h,\mathcal{M}} - \Pi_{h,\mathcal{M}\setminus m} = (p_{mh} - mc_h)\, \tilde{Q}_{mh},$$

[12] Gowrisankaran, Nevo, and Town (2015) emphasize patient steering when $\tau = 1$. However, $\tau \neq 1$, the effect of steering on dispersion is generally ambiguous and depends on substitution patterns, τ, and the relative magnitudes of $\bar{x}_{mA}p_{mA}$ and $\bar{x}_{mB}p_{mB}$. Online Appendix A derives the full terms.

[13] In ESI settings, Gowrisankaran, Nevo, and Town (2015) estimate τ at 2.79 and Tilipman (2022) obtains similar magnitudes at 3.70.

and insurer m's GFT from reaching an agreement with hospital h are

$$GFT_{m,h} \equiv V_{m,\mathcal{H}} - V_{m,\mathcal{H}\setminus h} = \tau\big(W_{m,\mathcal{H}} - W_{m,\mathcal{H}\setminus h}\big) - \big(C_{m,\mathcal{H}} - C_{m,\mathcal{H}\setminus h}\big).$$

The Nash-in-Nash first-order condition (FOC) corresponding to (7) is therefore

$$p_{mh} = mc_h - \left(\underbrace{\frac{\partial Q_{mh}}{\partial p_{mh}}}_{\Omega} + \underbrace{\omega \left(\tau \frac{\partial W_{m,\mathcal{H}}}{\partial p_{mh}} - \frac{\partial C_{m,\mathcal{H}}}{\partial p_{mh}} \right)}_{\Lambda} \right)^{-1} \tilde{Q}_{mh}. \tag{11}$$

where Ω is the standard consumer price sensitivity and Λ captures the additional disciplining force coming from insurer price sensitivity; the relative importance of which is weighted by

$$\omega \equiv \frac{b_m}{b_h} \frac{\tilde{Q}_{mh}}{GFT_{m,h}}.$$

Gowrisankaran, Nevo, and Town (2015) refers to $\Omega + \Lambda$ the *effective price sensitivity*.

When $\Lambda = 0$, negotiated prices are pinned down by consumer substitution alone, yielding the standard Bertrand-Nash markup implied by Ω, providing a useful list-price benchmark. A larger magnitude of $-\Lambda$ (recall $\partial V_{m,\mathcal{H}}/\partial p_{mh} < 0$) increases effective price sensitivity and reduces the negotiated markup, moving p_{mh} toward marginal cost mc_h.

Remark 2. *For any pair that reaches agreement, $\tilde{Q}_{mh} > 0$ and $GFT_{m,h} > 0$, so $\omega \in \mathbb{R}_{++}$. If $\tilde{Q}_{mh} = 0$, the hospital has no incremental volume at stake in contracting with insurer m, so the transacted price reverts to the non-negotiated benchmark. If $GFT_{m,h} \leq 0$, insurer m optimally excludes hospital h from its network.*[14]

The ratio ω can be interpreted as capturing consumers' relative loyalty to the insurer versus the hospital. If consumers readily switch insurers to access hospital h, then $\tilde{Q}_{mh}$ is large and $V_{m,\mathcal{H}\setminus h}$ is small, reducing ω and limiting the role of insurer bargaining in disciplining prices. The opposite occurs when switching is limited, so that ω is large and the hospital risks losing substantial volume in the event of disagreement.

2.4 Interaction between the two markets

Some consumers enroll in PP plans but receive care at the non-preferred hospital A ex post, reducing insurer m's exposure at hospital A: $\downarrow \bar{x}_{mA}$. This tiered exposure creates a two-way feedback loop. First, $\bar{x}_{mA}$ determines insurer m's marginal gain from a price reduction as in Proposition 2. Second, holding generosity fixed, the hospital price wedge determines the premium wedge between FC and PP plans, shifting enrollment across network structures and feeding back into $\bar{x}_{mA}$.[15]

[14] For example, Shepard (2022) studies an episode in which a plan excludes an expensive hospital system when cost savings exceed the associated loss in enrollment revenue.

[15] Holding generosity fixed, an FC contract is more expensive than a PP plan because it both (i) attracts relatively more high-θ^A consumers and (ii) induces a higher share of utilization at the higher-price hospital A. A larger price

Table 1 summarizes the sign of these comparative statics for a tiered insurer ($\bar{x}_{mA} < \bar{x}_{mB}$) across values of τ (how strongly insurer revenue internalizes consumer surplus) and ω (the weight placed on insurer price sensitivity relative to consumer price sensitivity in (11)).

Table 1: Comparative statics under tiered exposure ($\bar{x}_{mA} < \bar{x}_{mB}$)

Shock: $\downarrow \bar{x}_{mA}$	Bargaining regime			
	(1) $\tau < 1,\ \omega > 1$	(2) $\tau < 1,\ \omega < 1$	(3) $\tau > 1,\ \omega > 1$	(4) $\tau > 1,\ \omega < 1$
Hospital market				
Negotiated price p_{mA}	$\uparrow$	$\downarrow$	$\downarrow$	$\downarrow$
Price wedge $\Delta p \equiv p_{mA} - p_{mB}$	$\uparrow$	$\downarrow$	$\downarrow$	$\downarrow$
Insurance market				
Premium wedge $\Delta t \equiv t_{FC} - t_{PP}$	$\uparrow$	$\downarrow$	$\downarrow$	$\downarrow$
Insurer exposure at A: $\bar{x}_{mA}$	$\downarrow$	$\uparrow$	$\uparrow$	$\uparrow$

Notes: The shock $\downarrow \bar{x}_{mA}$ corresponds to greater tiering/coinsurance at hospital A (holding enrollment fixed in the hospital-market subsection). τ governs how strongly insurer revenue internalizes consumer surplus, and ω governs the weight placed on insurer incentives relative to consumer demand in the bargaining FOC. Δt is the premium differential between an FC plan and a PP plan with the same generosity; Δp is the negotiated price differential between hospitals.

Case (1): $\tau < 1$ and $\omega > 1$ (most relevant for individual markets). When insurer revenues imperfectly track consumer surplus ($\tau < 1$) and insurer incentives materially affect bargaining outcomes ($\omega > 1$), a reduction in $\bar{x}_{mA}$ weakens insurer m's marginal benefit from lowering p_{mA}, raising p_{mA} and increasing the price wedge $\Delta p \equiv p_{mA} - p_{mB}$. A larger Δp widens the premium differential $\Delta t \equiv t_{FC} - t_{PP}$ between FC and PP tiered plans of the same generosity, shifting enrollment toward PP plans and further lowering $\bar{x}_{mA}$. The resulting feedback loop tends to amplify both tiering and negotiated price dispersion.

The feedback loop admits an interior equilibrium. Prices are bounded between marginal cost and the list-price benchmark, and $\bar{x}_{mA} \in (0, 1)$ as long as PP plans retain positive coverage at A (i.e., do not collapse into EP designs) at any finite premium wedge. Since the mappings from $\bar{x}_{mA}$ to p_{mA} and from p_{mA} to $\bar{x}_{mA}$ are continuous, a fixed point exists by a standard fixed-point argument. Uniqueness is not guaranteed globally but holds locally when the feedback elasticities are sufficiently small.

Case (3): $\tau > 1$ and $\omega > 1$ (employer-sponsored settings). In environments where insurers (or employers) place substantial weight on consumer surplus, coinsurance can strengthen incentives to bargain down prices (Proposition 2). This is the configuration emphasized by Gowrisankaran, Nevo, and Town (2015), who estimate $\tau > 1$ and find that increasing coinsurance reduces negotiated prices in their setting. In this case, reductions in $\bar{x}_{mA}$ tend to lower p_{mA} and compress Δp, which narrows Δt and attenuates the tendency for enrollment to concentrate in tiered plans.

wedge Δp amplifies these cost differences, widening the premium wedge Δt. See Online Appendix A for details.

Cases (2) and (4): $\omega < 1$ (consumer demand dominates). When $\omega < 1$, negotiated prices are governed primarily by consumer price sensitivity Ω. In this regime, negotiated prices are close to the benchmark implied by consumer substitution alone (the $\Lambda \simeq 0$ case in the bargaining FOC). The feedback loop is therefore driven mainly by the standard demand-discipline mechanism: higher out-of-pocket exposure increases consumer responsiveness, thereby constraining markups and compressing negotiated price dispersion.

3 Institutional setting and empirical patterns

To test the relationship between expected coverage ($\bar{x}_{mh}$) and negotiated hospital prices (p_{mh}), I study Chile's private health care markets. Three institutional features make it well-suited for empirical analysis of insurer-hospital transactions in a managed-competition context. First, the data is comprehensive: the health regulator collects the universe of plan enrollment and itemized claims from all private insurers. Claims are reported using standardized service codes that map directly into the public fee schedule.[16] Second, plan contracts are largely linear, with coinsurance as the primary financial attribute. This simplifies consumers' price exposure relative to non-linear structures with deductibles and out-of-pocket maxima (Diaz-Campo, 2023; Prager, 2020). Network design is also typically defined at the hospital level, and primary-care gatekeeping is uncommon (unlike HMO plans in the U.S.). Third, insurance is mandatory and private coverage is subsidized through an income-based contribution that favors the young and healthy (Tebaldi, 2024). More broadly, the value of private coverage in Chile derives primarily from access to private providers rather than from avoiding an uninsured state (Gruber, 2008), since the outside option is a public system that offers comprehensive care.

3.1 The mixed health care system in Chile

Chile's health care system is mixed with both public and private insurance and providers.[17] Health insurance is mandatory. The law requires a minimum contribution of 7% of one's taxable income (up to a cap) in either the state insurance fund (called *Fonasa*) or the purchase of a private plan from a private insurance company (called *Isapre*). Effectively, the state provides an income-based tax credit for purchasing private insurance. Consumers are free to choose an insurance company, but the mandatory premium contribution amount is not adjustable. Consumers can optionally top up with a voluntary contribution to purchase better plans.[18] Public insurance offers no plan choice and provides coverage primarily through the public hospital system. By contrast, private insurance offers access to both public and private hospitals. As a result, private coverage is most attractive to

[16] The public price list is set administratively by the public insurer, listing its willingness-to-pay for private providers for medical services that it reckons *basic*. The list is updated annually.

[17] Variations of mixed health care systems are common around the globe, especially among Latin American countries (McPake & Noramnd, 2008).

[18] Consumers can also purchase complementary insurance from a different set of companies (e.g., banks) that reimburses their out-of-pocket expenses but do *not* engage in setting hospital prices. In 2015, 45% of the Isapre enrollees held complementary insurance. The contraction of complementary insurance is not observable in my data.

higher-income individuals for whom the income-based contribution is sufficient to finance a private plan that extends coverage to the private hospital network. Most insurance buyers are enrolled in individual plans traded in a regulated market. A few companies, mostly in the mining sector, hire their own "closed" insurers to provide ESI.

Vertical integration between private insurers and hospitals in Chile is achieved through a third-party holding company that simultaneously owns majority shares in an insurer and one or more hospitals. Unlike the integrated models in the U.S., such as Kaiser Permanente, insurers and hospitals under a common holding maintain separate balance sheets and conduct transactions at established fee-for-service prices. Every open insurer has hospitals under a common holding, but the high-priced star hospitals are independent (Cuesta, Noton, & Vatter, 2024).[19]

Hospital prices are partially transparent. Chile's price transparency law requires private hospitals to publicly disclose a fee-for-service price list for the uninsured and the publicly insured, as well as a reference price for privately insured patients. Most private hospitals offer a personalized budgetary quote conditional on one's insurance. Therefore, cross-hospital price variation is salient, but cross-insurer variation is not.[20]

3.2 Data

I use administrative data from Chile's health care regulator covering the universe of private insurers, including the five open insurers in the individual market and six closed (ESI) insurers.

The plan data is at the product level and covers the years 2006 to 2019. It lists all plans with at least one enrollee and provides each plan's characteristics and enrollee count. I complement the plan data with beneficiary data that provides enrollee characteristics and residential region at the individual level. Notably, plans in Chile do not include deductibles, fixed-amount copays, or out-of-pocket maximums; they specify coinsurance rates for inpatient care versus outpatient care, and for preferred providers versus non-preferred providers.[21] I observe whether a plan is classified as FC, PP, or EP, but I do not observe the plan's hospital-level tier schedule (i.e., which hospitals are on the preferred tier). Accordingly, I recover insurer-hospital "expected coverage" from claims-based realized coverage.

The claims data is at the service level and covers the years 2012 to 2019. It lists service-level itemized transactions between three parties: a private insurer, the patient's coinsurance payment and the medical service provider. I merge it with each year's public price list, then aggregate service-level claims into outpatient visits and inpatient admissions, yielding 89 million visits and 4.4 million admissions. A visit is defined as a patient's annual spending at a provider. An admission is defined as a patient's hospitalized treatment for a diagnosis within a quarter; when treatment

[19] Direct ownership between insurer and hospital is outlawed. See Cuesta, Noton, and Vatter (2024) for a welfare discussion of vertical integration in this context.

[20] Interested readers may compare the price lists at Clínica Alemana and at Clínica Santa María (accessed January 2026).

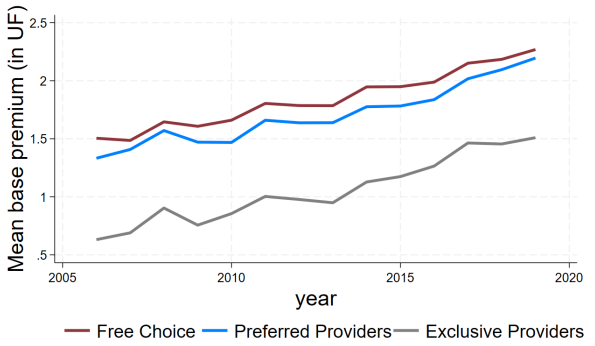
[21] Plans also differ in exclusion (e.g., childbirth) and annual coverage cap. Once coverage reaches the cap, the insurer only covers the legal minimum of 25%. The cap is set quite high and rarely binds.

spans multiple providers, the main provider is the hospital with the longest stay.

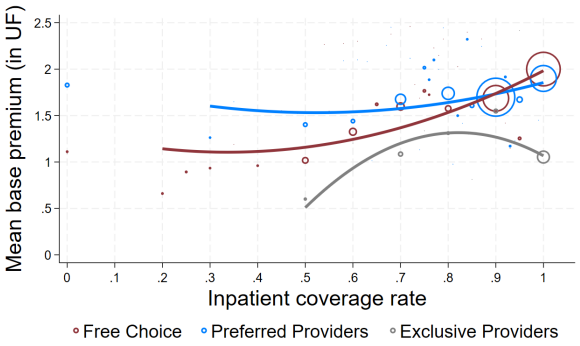
Additionally, accounting data was collected from the financial reports of open insurers and major hospitals. I provide additional background descriptions, along with other data sources, in the Online Appendix B.

3.3 Plan characteristics and enrollment

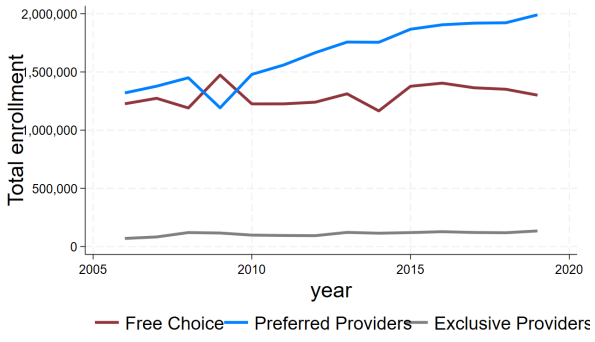
Figure 3 provides an overview of the insurance market. On average, PP and EP plans are cheaper than FC plans. Consistent with Proposition 1, they obtain lower costs by placing lower-cost hospitals on the preferential coverage tier – i.e., most PP and EP plans have lower network quality than FC plans. Moreover, FC plans' premium schedules are convex in coverage. At the same time, its convexity is reduced or even reversed for PP and EP plans – consistent with the lower-quality PP plans' flatter premium schedule depicted in Figure 2. Enrollment trends show that PP plans are the most popular network type, accounting for nearly all enrollment growth. Most chosen plans have inpatient and outpatient coverage rates between 0.7 and 1, because the risk protection of lower-coverage plans would be inferior to that of public insurance.



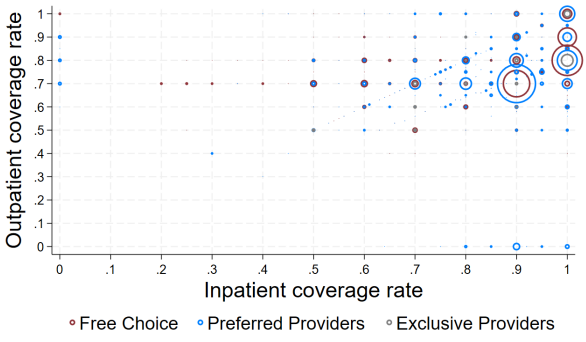
(a) Premium by network



(b) Premium by coverage and network



(c) Enrollment by network



(d) Enrollment by coverage and network

Notes: The size of each circle reflects the number of enrollees. Premiums are base prices for a risk factor of 1. Premiums are denominated in UF (Unidad de Fomento), a non-circulating unit indexed to inflation that maintains constant real value. 1UF≃ 30-56 USD during the study period.

Figure 3: Overview of the insurance market (based on plan and enrollee data)

The five open insurers appear to compete fiercely in plan offering. Table 2 shows that each insurer enrolls consumers across thousands of plans, and the number of plans that each insurer actively sells each year ranges from 328 to 2,870. The large number of plan offerings suggests that a similar combination of plan characteristics is likely offered by multiple companies.[22] The median number of enrollees per plan ranges from 2.4 to 15.9, suggesting the possibility of individual risk-based pricing.[23] Inactive plans persist because guaranteed renewability leads to contract lock-in: consumers with deteriorated health status stay enrolled in legacy contracts, which leads to unprofitable plans becoming inactive (Atal, 2019).

The five open insurers serve enrollees with broadly similar incomes and therefore receive similar per capita premium contributions. Relative to beneficiaries in ESI, their enrollee population is substantially younger, less likely to be female, has fewer dependents, and is more likely to reside in the capital city, Santiago. These patterns suggest that the insurance mandate and generous subsidies already bring many "young invincibles" into the market (Tebaldi, 2024), strengthening the risk pool and limiting the scope for unraveling through nonparticipation. On the other hand, consumers of open insurers exercise choice more frequently, as evidenced by shorter tenure with their insurer and more frequent insurer switching.

Despite overlap in the menu of contract types, insurers differ sharply in how enrollment is allocated across plans. Larger insurers place more enrollment in plans with lower generosity and lower network quality. For example, the largest open insurer (m_1) enrolls almost all members in PP plans. It also has a much smaller share of enrollees with full inpatient coverage than the smallest insurer (m_5), which enrolls predominantly in FC plans with full inpatient coverage. Larger insurers may enjoy a stronger brand reputation, as an average enrollee of m_1 pays a higher premium, remains affiliated longer, and enrollment has grown over the years as more switch in than switch out. These enrollment differences are central to the hospital-market analysis, as they translate into systematic cross-insurer differences in realized coverage at high-priced versus low-priced hospitals, even though insurer-level networks are complete.

3.4 Constructing insurer-hospital transacted price index

Following Cuesta, Noton, and Vatter (2024), I decompose the total paid amount for each visit or admission into (i) an insurer-hospital-year price index p_{mhy} and (ii) an individual resource-intensity weight w_i, using the public fee schedule as a common benchmark. I first separate each patient's total spending P_{imh} into "basic" services that appear on the public price list and "off-list" services. As shown in Table 3, basic services account for the bulk of outpatient revenue and a large share of inpatient revenue. I therefore estimate negotiated prices using only basic services to ensure comparability across insurers and hospitals.[24]

[22] Cuesta, Noton, and Vatter (2024) consolidated the plans by attributes and applied additional filters to reduce the consumer choice set. They analyze 4,110 plan-years, including 1,431 active plans in the spot market.

[23] A reason for the myriad of plan offering is to circumvent regulations that limit risk-rating and cap premium increases.

[24] Drugs and consumables are not separately reimbursed under the public fee schedule.

Table 2: Characteristics of the insurance companies (based on plan and enrollee data)

yearly average 2012-2019		m_1	m_2	m_3	m_4	m_5	ESI
Market share (%)		24.8	20.7	20.7	15.9	14.9	2.9
Plan offering	number of plans with enrollees	15,891	8,421	7,516	18,151	3,249	719
	number of plans active for sale	486	1,658	1,696	2,870	328	173
	median number of enrollees per plan	5.4	11.6	3.0	2.4	15.9	5.6
Policy holder characteristics	mean taxable income	1,157	1,105	975	1,218	1,105	1,300
	mean premium contribution	130	125	103	129	106	137
	mean number of dependents	0.8	0.7	0.7	0.7	0.7	1.1
	mean tenure of affiliation	8.6	7.1	7.5	7.3	6.6	19.2
Enrollee characteristics	mean age	34.1	32.6	33.7	33.0	30.7	41.1
	female (%)	44.7	46.4	41.7	49.7	45.3	51.0
	Santiago resident (%)	71.1	60.9	57.3	62.3	37.5	20.2
Enrollment (%) by type of network	Free Choice (FC)	2.3	39.3	43.9	72.7	69.4	11.7
	Preferred Provider (PP)	97.5	56.5	45.9	27.3	30.6	67.7
	Exclusive Provider (EP)	0.1	4.2	10.2	—	—	20.6
Enrollment (%) by outpatient coverage	1	6.9	3.8	0.0	8.8	0.0	51.9
	[0.9,1)	4.3	13.9	8.0	11.7	15.9	3.9
	[0.8,0.9)	33.0	10.5	44.6	28.7	75.3	9.5
	[0.7,0.8)	45.9	68.8	45.2	43.0	8.8	31.9
	<0.7	9.9	2.9	2.2	7.8	0.0	2.8
Enrollment (%) by inpatient coverage	1	34.8	26.4	43.2	49.4	90.8	59.8
	[0.9,1)	45.7	58.5	53.7	47.8	8.9	28.9
	[0.8,0.9)	7.2	10.0	1.7	2.4	0.2	9.0
	[0.7,0.8)	2.5	5.1	1.2	0.3	0.1	2.3
	<0.7	10.0	0.1	0.2	0.1	—	0.0
Change in enrollment (%)	new to private insurance	6.1	8.0	8.1	7.6	7.2	1.4
	switched in from other insurers	6.4	7.2	6.1	6.8	6.0	0.9
	switched out to other insurers	5.5	7.4	5.9	5.9	9.8	1.0
	left private insurance	5.6	6.2	7.4	4.6	5.8	2.1

Notes: I restrict the sample to 2012-2019 to correspond to the available claims data. ESI, employer-sponsored insurance. Santiago is Chile's capital, where 40% of the population resides.
Policy holder characteristics are only available since 2016; income and premium contribution are in thousand pesos.

I then estimate insurer-hospital-year multipliers relative to the public benchmark by regressing log spending on basic services relative to the log public-priced counterfactual on insurer-hospital-year fixed effects:

$$\ln\left(P_{imh}^{\text{basic}}\right) - \ln\left(P_{imh}^{\text{public}}\right) = \rho_{mhy} + \varepsilon_{imh}, \tag{12}$$

where P_{imh}^{public} is the public-priced counterfactual for the same basket of basic services, ρ_{mhy} is an insurer-hospital-year fixed effect. The residual ε_{imh} captures idiosyncratic variation in treatment intensity (e.g., length of stay and quantity of tests and exams) and service bundles.[25]

[25]Outcomes are winsorized at the 0.1st and 99.9th percentiles.

I define the negotiated price index as

$$p_{mhy} \equiv \exp(\rho_{mhy}),$$

which can be interpreted as an insurer-hospital-year multiplier relative to the public prices for basic services. This multiplier maps directly to the unit price p_{mh} in the theoretical model. I then recover a resource-intensity weight

$$w_i \equiv \frac{P_{imh}^{\text{basic}}}{p_{mhy}},$$

which only counts basic services and is therefore less contaminated by non-medical or non-standard billing components. It is thus a transferable measure of clinical severity that is independent of hospital-specific technology.[26]

As a reference point, I use hospital accounting reports to construct a hospital-year marginal-cost index per unit of resource intensity that is comparable to the price index.[27] This construction motivates treating marginal cost as approximately hospital-specific rather than insurer-specific, conditional on case mix.

3.5 Does negotiated price vary with coverage within hospital?

The Chilean private hospital market is geographically concentrated in Santiago. I focus on the ten largest private hospitals by market share and classify them as high-, medium-, or low-priced using negotiated price indices.[28] The high-priced, high-market share hospitals h_1^{high} and h_6^{high} fit the notion of star hospitals: during the sample period, they were widely regarded as the most prestigious private providers in Chile.[29]

Importantly, insurer-level networks are effectively complete: all insurers transact with all major hospitals, including insurers that operate large vertically integrated facilities. The relevant variation is therefore not whether a hospital is "in network", but how generosity and negotiated prices differ across insurer-hospital pairs within a common network. Table 3 documents substantial within-hospital price dispersion across insurers, concentrated in inpatient care and especially at high-priced hospitals. Vertical integration is associated with lower negotiated prices, consistent with classic internal pricing arguments (Cuesta, Noton, & Vatter, 2024; Spengler, 1950). But insurer size does not mechanically translate into lower negotiated prices: for example, the smaller m_5 pays roughly 25-35 percent lower inpatient prices at the two high-priced hospitals than the larger insurer

[26] The interpretation is analogous to a Diagnosis-Related Group (DRG) weight. I normalize one unit of w_i to correspond to one million Chilean pesos. Because it is constructed from fee-for-service claims, it varies at the patient level rather than only by diagnosis.

[27] I scale a hospital's reported service costs by the share of revenue attributable to private insurance and assume the same outpatient/inpatient split for revenue and cost.

[28] This classification broadly aligns with the regulator's characterization of major private hospitals. See https://www.supersalud.gob.cl/documentacion/666/w3-article-5904.html.

[29] In particular, h_1^{high} Clínica Alemana often features among the top 3 best hospitals in Latin America (see América Economía).

m_2, even though their outpatient prices at those hospitals are nearly identical.[30]

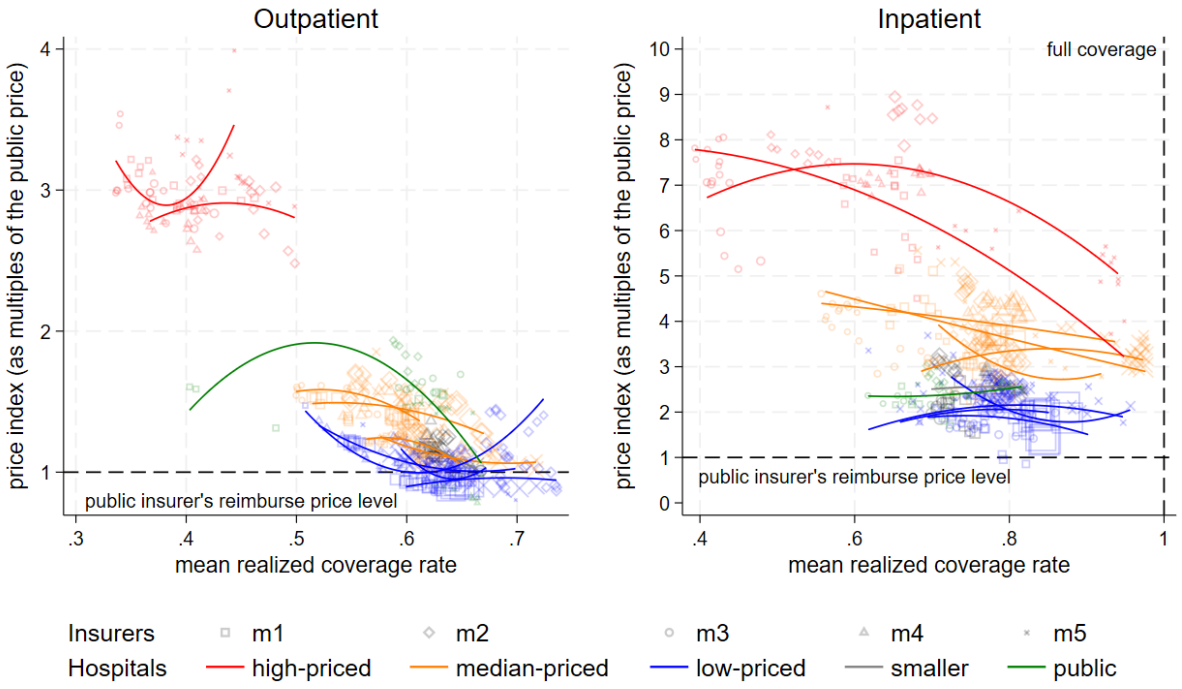
Table 3: Insurer-hospital transacted prices (based on claims data)

yearly average 2012-2019	smaller providers	**Top Ten Largest Private Hospitals**										public hospitals
		h_1^{high}	h_2^{medium}	h_3^{low}	h_4^{medium}	h_5^{medium}	h_6^{high}	h_7^{low}	h_8^{low}	h_9^{low}	h_{10}^{medium}	
Panel A: Outpatient care												
Market share (%)	78.8	3.8	3.2	1.7	2.6	2.2	2.6	0.7	0.8	0.9	0.7	2.0
Revenue breakdown (in billion pesos)												
Total	806	110	58	18	45	29	61	9	8	11	8	38
basic services	607	106	56	15	40	28	58	8	6	11	8	23
drug/consumable	67	0.7	0.3	0.6	0.3	0.2	0.5	0.08	0.1	0.04	0.05	4
other	131	3	2	2	5	1	2	1	1	0.7	0.7	11
Marginal cost index	—	—	1.1	0.8	—	0.8	2.3	—	0.9	0.7	0.9	—
Transacted price index by insurer												
m_1	1.1	2.9	1.4	0.9	1.5	1.1	3.1	1.2	1.0	0.9	1.2	1.5
m_2	1.2	2.9	1.6	1.1	1.4	1.3	3.0	1.0	1.4	1.0	1.3	1.7
m_3	1.0	2.9	1.6	0.9	1.3	1.1	3.0	1.0	1.0	0.9	1.1	1.6
m_4	1.2	2.8	1.5	1.0	1.3	1.1	2.8	1.2	1.1	1.0	1.2	0.9
m_5	1.2	2.8	1.5	1.0	1.3	1.1	2.8	1.1	1.1	0.9	1.3	1.5
ESI	1.1	2.9	1.4	1.1	1.9	1.1	2.3	1.3	1.1	1.0	1.5	1.1
Panel B: Inpatient care												
Market share (%)	45.3	8.0	8.1	7.1	6.9	5.5	5.3	3.0	2.5	2.4	2.0	3.8
Revenue breakdown (in billion pesos)												
Total	392	170	125	80	106	78	131	38	26	24	32	29
basic services	230	122	78	41	61	51	92	18	15	13	17	21
drug/consumable	114	39	39	29	30	21	31	14	7	7	9	3
other	48	9	8	10	15	6	8	6	5	3	7	5
Marginal cost index	—	—	2.7	1.5	—	2.3	5.3	—	1.7	1.4	2.5	—
Transacted price index by insurer												
m_1	2.1	6.6	3.2	1.8	4.5	2.8	6.4	1.9	1.9	1.5	3.3	2.1
m_2	3.1	8.5	4.8	2.7	3.5	3.6	7.7	2.0	2.6	2.3	3.7	2.7
m_3	2.1	6.7	3.9	2.2	4.0	3.0	6.9	2.0	2.0	1.6	2.8	2.2
m_4	2.8	7.2	4.4	2.5	4.3	3.3	7.1	2.6	2.2	2.3	3.8	2.6
m_5	2.5	5.4	3.4	2.2	4.3	3.2	5.8	2.2	2.6	2.9	3.5	2.5
ESI	1.5	7.8	4.2	2.8	4.7	3.2	5.3	2.7	2.4	2.5	4.5	2.1

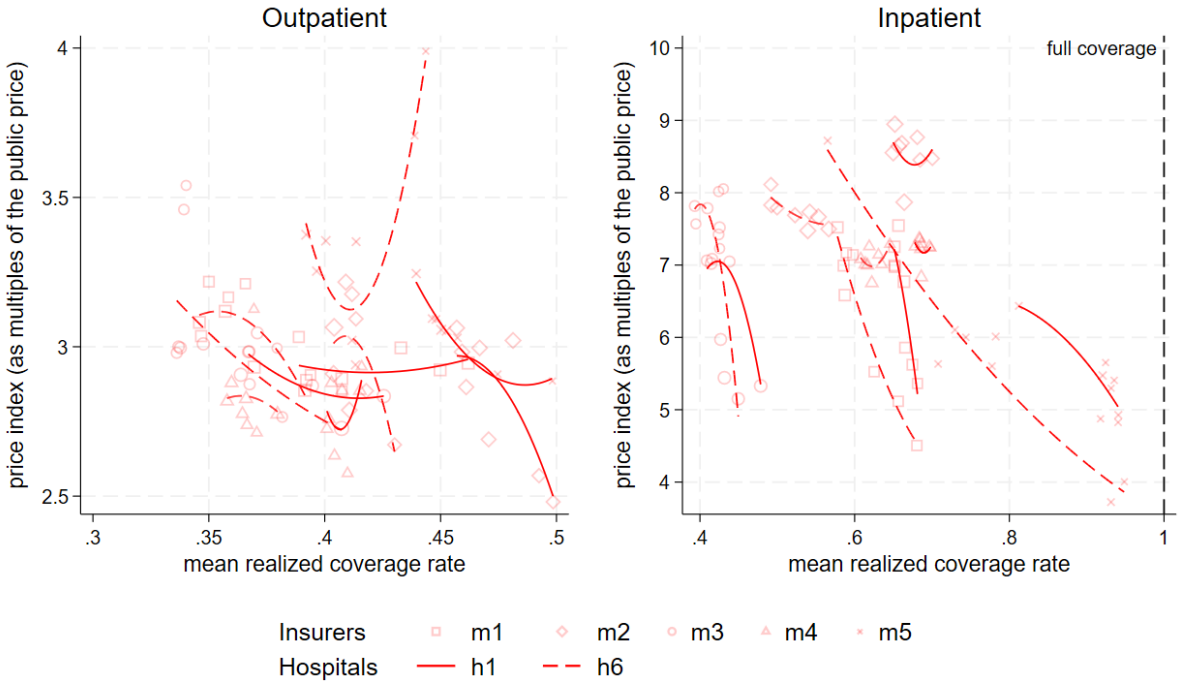
Notes: All marginal cost and price indices are for 1 unit of resource intensity weight. Not all hospitals publish cost reports.
h^{high}, h^{median} and h^{low} are high-priced, median-priced and low-priced hospitals, respectively. 'Smaller providers' include all the rest of the regional and specialized providers. h_4^{median} and h_7^{low} are academic hospitals.
ESI, employer-sponsored insurance. Underlined cells are vertically integrated through a common parent holding. Beyond these underlined large hospitals, all insurers have smaller hospitals under common holding that are included in "smaller providers".

I next relate this dispersion to mean realized coverage. Because insurers sell many contracts with different tiering and coinsurance, an insurer's realized coverage at a given hospital reflects how its enrollees are distributed across plan types. Figure 4 shows a clear, nonlinear yet predominantly negative relationship between an insurer's realized coverage rate at a hospital and the negotiated inpatient price at that hospital. The pattern is strongest at the two high-priced hospitals. Panel (a) plots the within-hospital relationship across insurer-year observations, while panel (b) plots within insurer-hospital changes over time for the high-priced hospitals. For outpatient care, by contrast, the coverage-price relationship is weak and shows no systematic pattern.

[30] Online Appendix B presents price dispersion for relatively homogeneous services. Similarly, cross-insurer variation is near null for outpatient services but substantial for inpatient services.



(a) Correlation across insurer-year, by hospital



(b) Correlation within insurer across year, for high-priced hospitals

Figure 4: Relationship between insurance coverage rates and negotiated hospital price indices

Notes: Each marker represents a insurer-hospital-year combination; size of the marker represents patient volume. In (a), each line is a quadratic fit for a hospital; in (b), each line is a quadratic fit for a hospital-insurer pair.

Consistent with Proposition 1 that PP plans economize on expensive providers, insurers with a higher share of enrollment in PP plans (i.e., m_1, m_2 and m_3) exhibit markedly lower coverage at high-priced hospitals and a steeper coverage gradient across hospital types. Insurers with greater enrollment in FC plans (i.e., m_4 and m_5) display a flatter coverage profile. This mapping from enrollment allocations to hospital-specific coverage provides the empirical bridge from plan heterogeneity in the insurance market to negotiated price dispersion in the hospital market.

3.6 Is selection adverse or advantageous?

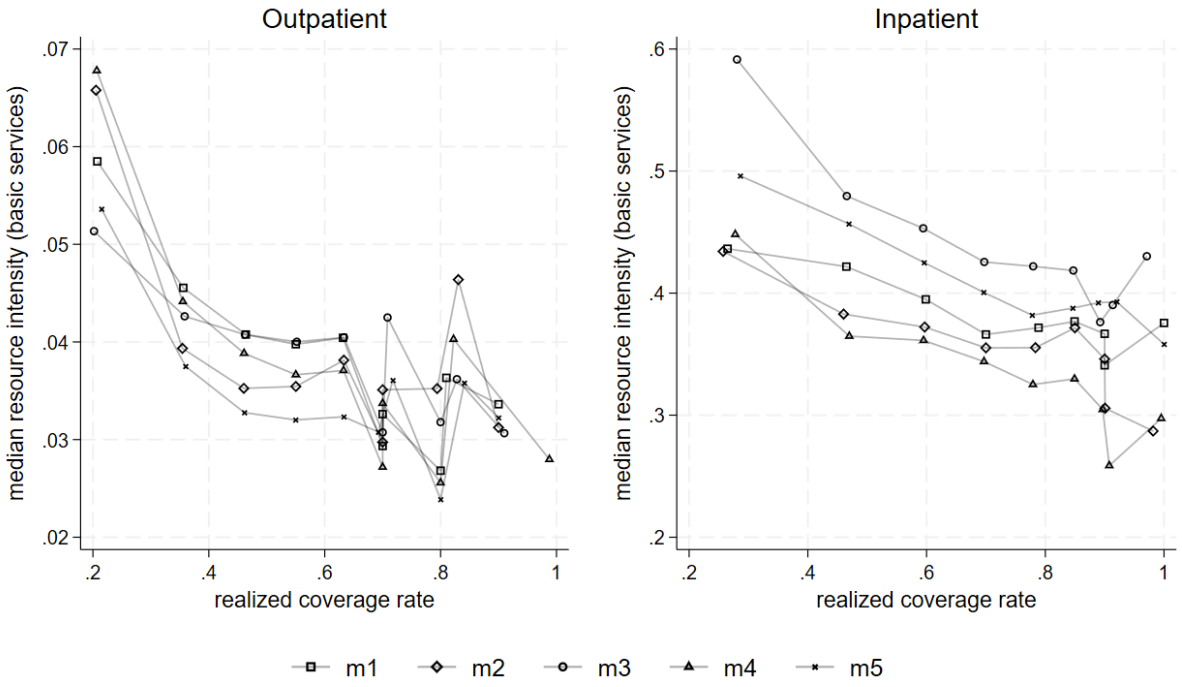
The negative coverage-price slope does not, by itself, imply that tiering and coinsurance improve or reduce welfare. With adverse selection, healthier consumers anticipate low utilization and may rationally prefer low-premium plans that offer less financial protection and weaker access to high-priced hospitals. If selection is strong, insurers with low expected coverage at expensive hospitals may also enroll systematically healthier patients, so part of the coverage-price relationship could be driven by case mix.

I test for the presence of adverse selection by plotting the relationship between insurance coverage and patient resource intensity. In Figure 5(a), conditional on insurer fixed effects, coverage is negatively correlated with inpatient resource intensity (advantageous selection). Importantly, cross-insurer differences in patient risk do not line up with cross-insurer dispersion in negotiated prices. For example, m_3 and m_5 exhibit similar resource-intensity profiles yet lie at opposite ends of the coverage-price gradient. This weak association argues against a pure case-mix explanation for the negative coverage-price slope.

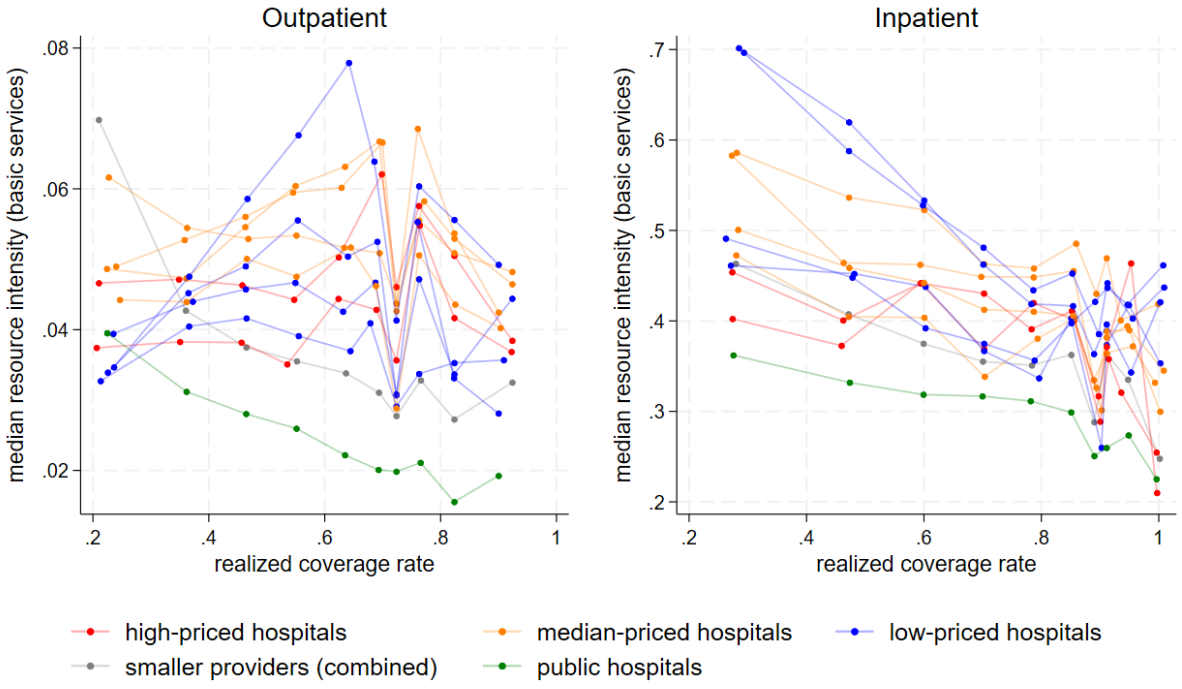
In Figure 5(b), adding hospital fixed effects reverses the pattern for outpatient care at low coverage levels (adverse selection) but not for inpatient care. Moreover, the overall resource insensitivity at higher-priced hospitals is lower, indicating a quality-sickness mismatch in the hospital market. Two factors may have contributed to the emergence of advantageous selection. One is specific to Chile because its generous tax credit *increases* with income, leading to higher-income households – who tend to be healthier and more risk-averse (Cohen & Einav, 2007) – to choose more generous plans. However, it does not explain the distinct pattern between outpatient and inpatient settings. The difference may be due to the unpredictability of inpatient spending. Patients with chronic conditions can estimate outpatient spending and doctor preference at plan choice, but the cost of a hospitalization is less foreseeable. The contrast motivates separate demand estimates by inpatient and outpatient settings.

4 Structural estimates

This section quantifies the demand and surplus components that map insurance coverage into negotiated hospital prices for the five open insurers. The bargaining FOC depends on consumer price sensitivity, $\partial Q_{mh}/\partial p_{mh}$, and on the insurer's marginal incentive to reduce prices, $\partial V_m/\partial p_{mh}$, which in turn depends on $\partial W_m/\partial p_{mh}$ and $\partial C_m/\partial p_{mh}$. I estimate a hospital choice model using



(a) By insurer



(b) By hospital, controlling for insurer fixed effect

Figure 5: Relationship between patients' resource intensity and insurance coverage

Notes: All graphs are binned scatter plots from claims data. Each resource intensity weight is the total of a patient in a hospital within the same year. One resource intensity weight corresponds to 1 million Chilean pesos (1 million Chilean pesos $\simeq$ \$1.4-2.0 thousand U.S. dollars).

microdata to recover these objects and to assess which component covaries with coverage in the cross-section.

I do not model plan choice or disagreement payoffs. Doing so would require specifying consumers' insurance choice sets in a market with thousands of contracts, accounting for dynamic frictions arising from guaranteed renewability and inertia, and making additional assumptions about unobservables, such as income shocks and complementary insurance. Instead, I discipline bargaining with coinsurance by estimating second-stage hospital demand and directly computing the derivatives that enter the bargaining FOC. It is not this paper's objective to pin down τ or ω.

4.1 Predicting hospital-specific coverage

A challenge to estimate hospital demand is that realized coinsurance is endogenous: it reflects plan choice and utilization decisions and may therefore mechanically correlate with unobserved hospital preferences. I therefore construct a predicted coverage rate, $\hat{x}_{imh}$, following Gowrisankaran, Nevo, and Town (2015). The objective is not to rely on realized plan choice, which is itself jointly determined with hospital preferences, but to obtain a coverage proxy that is less tightly linked to individual selection and can be used to form out-of-pocket prices.

I estimate a Tobit model separately for each insurer-hospital pair, with realized coverage as the dependent variable. Explanatory variables include age (5-year brackets), sex, age $\times$ sex, resource intensity weight and its interactions with age and sex, diagnostic specialty indicators, year and region of residence fixed effects, special coverage regulation, indicators for policyholder status, and payment modality (direct payment versus refund). Predicted values $\hat{x}_{imh}$ capture systematic variation in coverage by observable risk and contract terms. At the insurer level, this procedure preserves differences in coverage profiles across hospitals that arise from enrollment allocations across plans (see Online Appendix C).

4.2 Logit hospital demand

To be consistent with the theoretical Hotelling framework, my empirical hospital choice model exploits the richness of the micro data to identify market-specific hospital valuations, following the recommendations by Berry and Haile (2021). Since all ten major private hospitals are located in Santiago within 30 minutes of driving distance, and residents of other regions commonly travel to Santiago for care, I define a common choice set $\mathcal{H}$ for all patients.[31] I estimate hospital demand separately for each insurer m and separately for inpatient and outpatient settings. Markets t are defined by year $\times$ age group (25-year brackets) $\times$ gender for outpatient care, and by year $\times$ age group (25-year brackets) $\times$ diagnosis specialties for inpatient care.[32]

[31] $\mathcal{H}$ includes the ten largest private hospitals, a public option, and an outside alternative that aggregates smaller providers. Smaller providers tend to offer a subset of specialties; treating them as a bundled outside option ensures availability for any diagnosis.

[32] The market definition balances flexibility and cell size. Raval, Rosenbaum, and Tenn (2017) show that consumers' own demographics and diagnoses play a more significant role in hospital choice than hospital characteristics.

Patient i in market t who holds contracts with insurer m obtains the following indirect utility from hospital h

$$u_{imht} = v_{mht} - \alpha_m \, OOP_{imh} - \beta_m D_{r(i)h} + \varepsilon_{imht}, \tag{13}$$

where $OOP_{imh} = (1 - \hat{x}_{imh}) w_i p_{mhy}$ is the predicted out-of-pocket spending from previous sections, and $D_{r(i)h}$ is a region-based distance.[33] ε_{imht} is assumed to follow type-1 extreme value distributions (McFadden, 1974). The parameters of interest are the price coefficient α_m and the distance coefficient β_m, which are common across markets within insurer m; the insurer- and market-specific valuation v_{mht} absorbs all hospital unobservables.[34]

Moment conditions. Let y_{imht} denote the observed individual choice and $\hat{\sigma}_{imht}(\alpha_m, \beta_m, v_{mht})$ the logit choice probability. I estimate all parameters jointly by stacking three sets of moment conditions. First, the likelihood score conditions

$$\mathbb{E}\left[y_{imht} - \hat{\sigma}_{imht}\right] = 0 \tag{14}$$

that pin down v_{mht} for each market. Second, the instrumental variable orthogonality condition

$$\mathbb{E}[v_{mht} \cdot z_{ht}] = 0 \tag{15}$$

that identifies α_m by instrumenting negotiated prices with the ESI price index (Cuesta, Noton, & Vatter, 2024). Since consumers in my analysis are not eligible for ESI, it can be viewed as a Hausman (1996) instrument. ESI prices vary with hospital cost shocks, but are negotiated by a separate payer group, making them plausibly orthogonal to open insurance enrollees' market-specific hospital valuation.[35]

Third, the moment condition

$$\mathbb{E}\left[(y_{imht} - \hat{\sigma}_{imht}) D_{r(i)h}\right] = 0 \tag{16}$$

that identifies β_m under the the common assumption that $D_{r(i)h}$ is exogenous to hospital preferences.

Normalization. Hospital valuations are normalized by setting $v_{\text{public},mt} = 0$ for all consumers. This normalization is economically natural: the outside option for private insurance enrollees is the public insurance system (Fonasa), which covers only public hospitals. Accordingly, all reported

[33] The claims data does not reveal the patient's address. The distance is thus measured from a region's capital city to Santiago. I also assume the distance to the public or outside options is zero, reflecting that public hospitals and small private hospitals are geographically available nationwide.

[34] For example, v_{mht} allows for preferences for academic hospitals and vertically integrated hospitals.

[35] For insurer 5, high coverage rates compress out-of-pocket variation, leaving the IV with insufficient identifying power. I impute α_5 with the average α of other insurers.

valuations and welfare measures are interpreted relative to access to the public hospital network.[36]

Further details on the demand model, estimation results, and validity checks are provided in Online Appendix C.

4.3 How does hospital demand elasticity vary with insurance coverage?

Estimated elasticities differ sharply across care settings. Inpatient demand is relatively inelastic, with own-price elasticities ranging from about -0.04 to -1.8, whereas outpatient demand is substantially more elastic, ranging from roughly -0.8 to -12. An immediate implication is that the demand-discipline component of the bargaining FOC (the consumer price sensitivity term Ω in (11)) is likely to be more important in outpatient markets than in inpatient markets.

Figure 6 shows that elasticities vary systematically with hospital price levels: demand is more price sensitive at higher-priced hospitals. More importantly, elasticities vary across insurers primarily through differences in coverage. When coverage is high (around 90 percent), the mean inpatient elasticity is close to -0.1, consistent with estimates in high-coverage ESI settings (Gowrisankaran, Nevo, & Town, 2015; Prager, 2020). As coverage declines, out-of-pocket exposure rises, and demand becomes more price-sensitive.

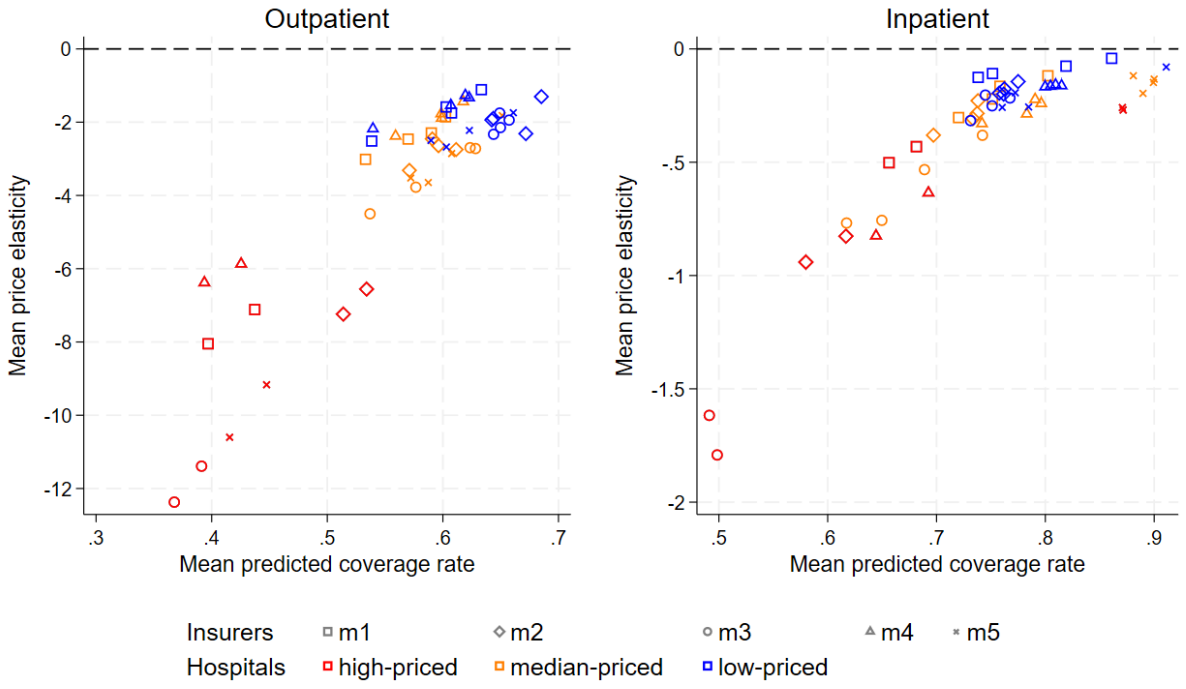


Figure 6: Relationship between insurance coverage and consumer elasticity to hospital price

[36] During the sample period, public insurance nominally requires 10-20% coinsurance for higher-income beneficiaries under age 60 at public hospitals, though implementation is limited and contested. For simplicity, I treat public hospitals as free of charge and set their travel cost to zero. Together with $v_{public} = 0$, this implies the public option delivers mean utility zero.

4.4 Why do some consumers buy PP plans yet use non-preferred hospitals?

Insurers with greater enrollment in PP plans have lower expected coverage at high-priced hospitals because a nontrivial share of PP enrollees nonetheless obtain care at hospitals placed on the non-preferred tier. Non-preferred-tier use is central for the mechanism in the paper: it is what generates tiered insurer exposure even when insurer-level networks are complete. This subsection provides suggestive evidence that such behavior is consistent with specialty-contingent hospital preferences that are difficult to forecast at the time of plan choice.

Figure 7 plots the joint distribution of monetized latent valuations for the two largest private hospitals, net of negotiated prices and predicted coverage (so the remaining variation reflects non-price hospital attractiveness). The diagonal line in each panel traces the indifference locus between the two hospitals, and the vertical and horizontal red lines at zero mark the public option benchmark. Darker shading indicates higher density.

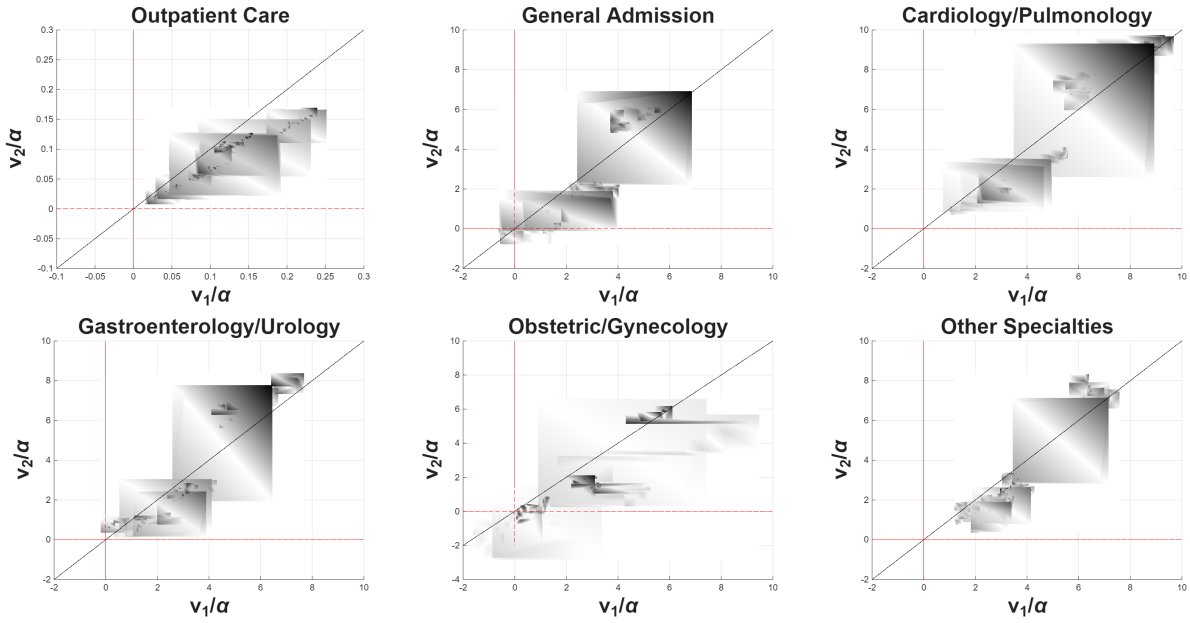


Figure 7: Preferences for h_1^{high} and h_2^{medium}: joint distribution of v_1 and v_2

Notes: Valuations are divided by α, so axes are in million pesos (1 million Chilean pesos $\simeq$ \$1.4-2.0 thousand U.S. dollars). Shades in each graph represent density. The diagonal line is a 45-degree line, where the two hospitals are considered equal; the red dashed line is the normalized public-hospital benchmark $v_{public} = 0$.

Two patterns stand out. First, both private hospitals are valued higher than the public option across service areas, with the gap particularly large for inpatient specialty care. For outpatient care, differences are modest – outpatient valuations are higher at the star hospital h_1^{high} but on the order of 0.1 million pesos. By contrast, inpatient specialty admissions command much larger valuations: in cardiology/pulmonology, implied willingness to pay reaches about 10 million pesos at both hospitals, relative to the public alternative. Second, preferences are strongly specialty-contingent. Although the two hospitals are similarly valued for general admissions, h_1^{high} is valued more highly

for obstetric/gynecology, while h_2^{medium} is valued more highly for gastroenterology/urology.

Because diagnoses and specialty needs are (largely) difficult to forecast at the time of insurance purchase, hospital tiering introduces an additional risk dimension. Consumers may rationally select lower-premium PP plans ex ante, yet ex post face high cost sharing when the realized specialty is better matched to a non-preferred hospital. This specialty-contingent heterogeneity also helps explain the limited appeal of EP plans: excluding high-quality, high-priced hospitals imposes large expected utility losses in the specialties where those hospitals are most valued.

4.5 Which derivative drives the negative inpatient coverage-price slope?

With the estimated demand primitives in hand, I compute hospital quantities, consumer surplus, insurer spending, and the derivative terms that enter the bargaining FOC.[37]

I then ask which primitives can account for the negative relationship between expected coverage and negotiated prices, focusing on inpatient care at high-priced hospitals. This is the setting where demand is relatively inelastic, and the coverage-price slope is steepest, so a local linear approximation to the (nonlinear) bargaining FOC is most informative. For each insurer-hospital-year cell, I compute $\partial Q/\partial p$, $\partial W/\partial p$, and $\partial C/\partial p$ and test whether controlling for these objects absorbs the reduced-form coverage-price slope.[38]

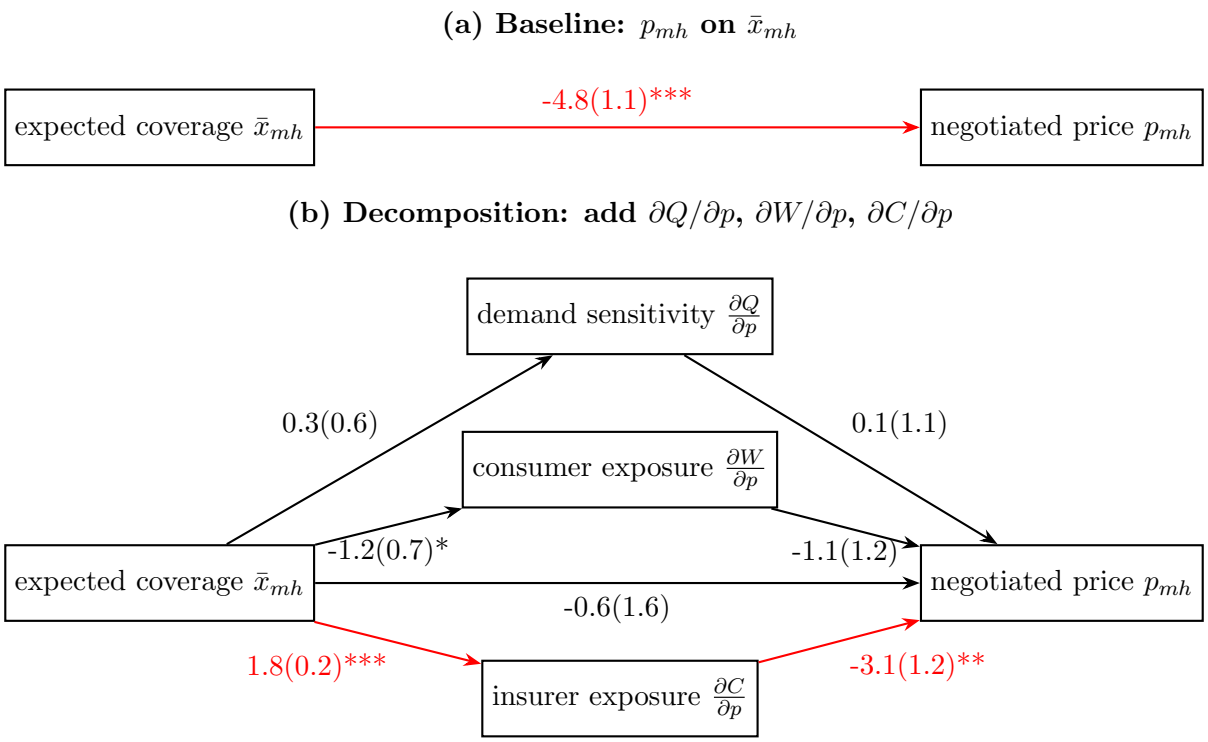
Figure 8 provides evidence consistent with the $\tau < 1$ case in Proposition 2. Panel (a) reports the baseline within-hospital relationship between expected coverage $\bar{x}_{mh}$ and negotiated prices p_{mh} across insurers and years. Panel (b) adds the three derivative terms as parallel covariates. Once insurer cost exposure, $\partial C/\partial p$, is included, the baseline coverage coefficient becomes small and statistically indistinguishable from zero: expected coverage strongly predicts insurer cost exposure, and higher insurer exposure is associated with lower negotiated prices.[39]

4.6 Why is $\tau < 1$?

Table 4 provides a diagnostic for insurer internalization by comparing insurer revenue from financial reports, consumer WTP implied by the estimated demand system, and hospital payments decomposed into insurer and patient coinsurance shares. It is noteworthy that a large share of insurer revenue comes from the mandatory 7% income contribution, which is pinned to taxable income rather than to the valuation of any particular plan or provider network. This creates a sizable fixed-revenue component, $\bar{R}_m$, in insurer surplus (6) and weakens the marginal link between consumer surplus and insurer revenue. This mandatory contribution could instead be paid into public insurance, which provides access to the public provider network (normalized to utility zero). Consistent with this limited marginal internalization, voluntary contributions are substantially below the consumer WTP implied by demand for access to the private provider network, suggesting that insurers capture only a fraction of the surplus they generate. The gap is most pronounced

[37] Under Logit demand, these derivatives have closed-form expressions; Online Appendix C derives the formulas.
[38] Online Appendix C illustrates how these objects relate to expected coverage.
[39] A similar pattern holds for medium-priced hospitals, though the implied magnitudes are smaller.

Figure 8: Linear accounting decomposition of the coverage-price slope



Notes: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. The sample is inpatient care at high-priced hospitals. All regressions include hospital and year fixed effects and insurer size (enrollment). Standard errors are robust. The three derivative terms $\partial Q/\partial p$, $\partial W/\partial p$, and $\partial C/\partial p$ are standardized (mean zero, standard deviation one) across insurer-hospital-year cells so that coefficients can be read as changes in p_{mh} per one-standard-deviation change in each derivative term.

for the biggest insurers, m_1 and m_2, which generate the highest consumer WTP yet have the lowest voluntary-contribution-to-WTP ratios, consistent with weaker marginal internalization among high-reputation insurers whose revenues are less tightly tied to consumer surplus.

This benchmark is not intended as a literal estimate of τ. Rather, it provides a transparent diagnostic: when a sizable portion of revenue is mechanically determined by income-based contributions, the marginal revenue benefit from bargaining down hospital prices can be substantially smaller than the consumer-side gains that accrue through lower out-of-pocket spending. In the bargaining framework, this corresponds to $\tau < 1$, the condition under which coinsurance can weaken insurers' direct exposure, and shift negotiated prices toward non-negotiated benchmarks (Proposition 2).

Table 4: Insurer revenue, consumer willingness-to-pay and hospital payments

yearly average 2012-2019	m_1	m_2	m_3	m_4	m_5
Insurer revenue (financial reports)					
Total	591	505	399	457	301
Mandatory contribution (7%)	374	431	282	307	241
Voluntary contribution	217	74	117	150	60
Consumer willingness-to-pay for the private provider network					
Total	688	336	128	303	144.5
Outpatient	19	6	-4	29	0.5
Inpatient	669	330	132	274	144
Insurers' hospital costs					
Total	426	354	265	304	201
Outpatient	158	161	136	117	95
Inpatient	268	193	129	187	106
Consumers' coinsurance					
Total	162	142	119	131	59
Outpatient	97	79	70	76	47
Inpatient	65	63	49	55	12

Notes: All numbers are in billion pesos (1 billion Chilean pesos $\simeq$ \$1.4-2.0 million U.S. dollars). The mandatory contribution is set at 7% of taxable income.

5 Discussion and conclusion

5.1 Limitation

The main limitation of this paper is that it is agnostic about *why* enrollee allocation across plans differs across insurers. The variation partly reflects brand reputation and marketing strategies, but also indicates market frictions in insurance, which are widely documented in the economic literature. Given the myriad of plans on offer, three types of frictions may have contributed to this variation. The first is choice friction. Individual insurance buyers often make sub-optimal choices by placing much higher weight on plan premiums than on expected out-of-pocket costs (Abaluck & Gruber, 2011, 2022), and fail to re-optimize when cost-saving plans become available in the choice set (Abaluck & Gruber, 2016) due to behavioral inertia (Handel, 2013; Polyakova, 2016), information frictions and hassle costs (Handel & Kolstad, 2015), or regulatory lock-in (Atal, 2019). The second is comparison frictions (Kling et al., 2012). Most of my sample period was before the advent of digital plan-comparison tools, so most consumers purchased plans through insurance sales agents, whose incentives may not fully align with the consumers' interests. The third is incomplete price transparency, as many consumers are unaware that prices at the same hospital can vary across insurers.

5.2 Policy implications

Managed-competition regulation typically targets consumer risk protection and sorting, but my results highlight an additional margin: plan design also shapes provider prices through insurers' bargaining incentives. Two implications follow. First, in settings with differentiated hospitals and incomplete insurer internalization ($\tau < 1$), greater plan heterogeneity that involves aggressive tiering can raise negotiated prices by reducing insurers' exposure at high-priced hospitals. Policies that limit coinsurance for inpatient episodes or constrain tiering when specialty needs are difficult to forecast may therefore improve welfare both by reducing risk exposure and by strengthening insurers' incentives to bargain. Second, price transparency could reduce dispersion if consumers can compare insurer-hospital prices and switch accordingly (Klemperer, 1987). However, the model cautions that stronger downstream competition can attenuate the role of insurer incentives in bargaining (a lower ω; Ho and Lee (2017)), potentially compressing dispersion while increasing the overall level of negotiated prices.

5.3 Conclusion

In this paper, I highlight the importance of accounting for hospital market responses when evaluating insurance market regulations. The stylized model provides a valuable framework for understanding how competition among insurance plans influences price-setting for differentiated hospitals.

Variations in hospital prices incentivize insurers to screen consumers based on their hospital preferences. However, paradoxically, tiered coverage contributes to the very hospital price dispersion that drives these selection incentives. Cost-sharing through coinsurance misaligns bargaining incentives, as insurers prioritize cost containment while patients prioritize the value of their medical services. From a partial equilibrium perspective, heterogeneous plan choices may seem cost-saving, but in a general equilibrium, they may ultimately lead to higher aggregate medical spending.

A Appendix

A.1 Proof of Proposition 1

Proof. Fix negotiated prices with $p_{mA} > p_{mB}$. Consider any PP contract $x^{PP}_{q>1} = (x_A, x_B)$ with $x_A > x_B$. Its generosity is $g^{PP} = \max\{x_A, x_B\} = x_A$.

Pick any $x \in (x_B, x_A)$ and consider the FC contract $x^{FC} = (x, x)$, which has strictly lower generosity $g^{FC} = x < g^{PP}$. Under free entry and zero-profit pricing, x^{FC} can be offered at a premium equal to its expected cost for the consumers who select it.

By Lemma 1, lowering generosity makes x^{FC} relatively more attractive to types with low θ^X (low demand for risk protection), and removing the A-tilt makes it relatively more attractive to types with low θ^A (low likelihood of choosing A when sick). Hence, there exists a nonempty set of lower-cost types – simultaneously low θ^X and low θ^A – who strictly prefer x^{FC} to $x^{PP}_{q>1}$ at the corresponding break-even premiums.

When these lower-cost types leave $x^{PP}_{q>1}$, the remaining pool in $x^{PP}_{q>1}$ becomes more concentrated in high θ^X and/or high θ^A types, increasing the average cost of $x^{PP}_{q>1}$ and therefore its zero-profit premium. This further expands the set of types who prefer x^{FC}, leading to unraveling. Therefore, no PP contract with $x_A > x_B$ can be sustained under free entry.

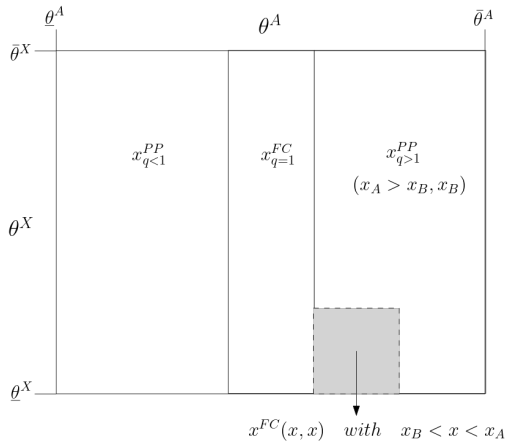


Figure A.1: Cream-skimming of A-preferred plans

□

A.2 Proof of Proposition 2

Proof. Since $\frac{\partial}{\partial \bar{x}_{mh}}\left(-\frac{\partial V_{m,\mathcal{H}}}{\partial p_{mh}}\right) = (1-\tau)Q_{mh}$, the derivative is negative when $\tau > 1$ and positive when $\tau < 1$ (as $Q_{mh} > 0$). Because greater coinsurance corresponds to lower $\bar{x}_{mh}$, the stated comparative statics follow. □

References

Abaluck, J., & Gruber, J. (2011). Choice inconsistencies among the elderly: Evidence from plan choice in the medicare part d program. *American Economic Review*, *101*(4), 1180–1210. https://doi.org/10.1257/aer.101.4.1180

Abaluck, J., & Gruber, J. (2016). Evolving choice inconsistencies in choice of prescription drug insurance. *American Economic Review*, *106*(8), 2145–84. https://doi.org/10.1257/aer.20130778

Abaluck, J., & Gruber, J. (2022). When less is more: Improving choices in health insurance markets. *The Review of Economic Studies*, *90*(3), 1011–1040. https://doi.org/10.1093/restud/rdac050

Akerlof, G. A. (1970). The Market for "Lemons": Quality Uncertainty and the Market Mechanism. *The Quarterly Journal of Economics*, *84*(3), 488–500.

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

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

Bernheim, B. D., & Whinston, M. D. (1998). Exclusive dealing. *Journal of Political Economy*, *106*(1), 64–103.

Berry, S., & Haile, P. (2021). Chapter 1 - foundations of demand estimation. In *Handbook of industrial organization, volume 4* (pp. 1–62, Vol. 4). Elsevier. https://doi.org/10.1016/bs.hesind.2021.11.001

Berry, S., & Haile, P. (2024). Nonparametric identification of differentiated products demand using micro data. *Econometrica*, *92*(4), 1135–1162. https://doi.org/10.3982/ECTA20731

Cohen, A., & Einav, L. (2007). Estimating Risk Preferences from Deductible Choice. *American Economic Review*, *97*(3), 745–788.

Collard-Wexler, A., Gowrisankaran, G., & Lee, R. S. (2019). "Nash-in-Nash" Bargaining: A Microfoundation for Applied Work. *Journal of Political Economy*, *127*(1), 163–195. https://doi.org/10.1086/700729

Cooper, Z., Craig, S. V., Gaynor, M., & van Reenen, J. (2019). The price ain't right? hospital prices and health spending on the privately insured. *The Quarterly Journal of Economics*, *134*(1), 51–107.

Craig, S. V., Ericson, K. M., & Starc, A. (2021). How important is price variation between health insurers? *Journal of Health Economics*, *77*, 102423. https://doi.org/10.1016/j.jhealeco.2021.102423

Crawford, G. S., & Yurukoglu, A. (2012). The welfare effects of bundling in multichannel television markets. *American Economic Review*, *102*(2), 643–85. https://doi.org/10.1257/aer.102.2.643

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

Cutler, D. M., & Zeckhauser, R. J. (2000). Chapter 11 - the anatomy of health insurance. In A. J. Culyer & J. P. Newhouse (Eds.), *Handbook of health economics* (pp. 563–643, Vol. 1). Elsevier. https://doi.org/10.1016/S1574-0064(00)80170-5

Diaz-Campo, C. (2023). *Dynamic moral hazard in nonlinear health insurance contracts* [Working Paper].

Einav, L., & Finkelstein, A. (2011). Selection in insurance markets: Theory and empirics in pictures. *Journal of Economic Perspectives*, *25*(1), 115–38. https://doi.org/10.1257/jep.25.1.115

Einav, L., & Finkelstein, A. (2018). Moral Hazard in Health Insurance: What We Know and How We Know It. *Journal of the European Economic Association*, *16*(4), 957–982. https://doi.org/10.1093/jeea/jvy017

Einav, L., Finkelstein, A., & Cullen, M. R. (2010). Estimating welfare in insurance markets using variation in prices. *The Quarterly Journal of Economics*, *125*(3), 877–921. https://doi.org/10.1162/qjec.2010.125.3.877

Einav, L., Finkelstein, A., & Mahoney, N. (2021). Chapter 14 - the io of selection markets. In K. Ho, A. Hortaçsu, & A. Lizzeri (Eds.), *Handbook of industrial organization, volume 5* (pp. 389–426, Vol. 5). Elsevier. https://doi.org/10.1016/bs.hesind.2021.11.014

Enthoven, A. C. (1988). *Theory and practice of managed competition in health care finance*. North-Holland. https://doi.org/10.1016/C2009-0-08402-5

Ericson, K. M., & Starc, A. (2015). Measuring consumer valuation of limited provider networks. *American Economic Review*, *105*(5), 115–19. https://doi.org/10.1257/aer.p20151082

Frank, M. B., Hsu, J., Landrum, M. B., & Chernew, M. E. (2016). The impact of a tiered network on hospital choice. *Health Services Research*, *50*(5), 1628–1648. https://doi.org/10.1111/1475-6773.12291

Gaynor, M., Ho, K., & Town, R. J. (2015). The industrial organization of health-care markets. *Journal of Economic Literature*, *53*(2), 235–84. https://doi.org/10.1257/jel.53.2.235

Ghili, S. (2022). Network Formation and Bargaining in Vertical Markets: The Case of Narrow Networks in Health Insurance. *Marketing Science*, *41*(3), 501–527. https://doi.org/10.1287/mksc.2021.1331

Gowrisankaran, G., Nevo, A., & Town, R. (2015). Mergers when prices are negotiated: Evidence from the hospital industry. *American Economic Review*, *105*(1), 172–203. https://doi.org/10.1257/aer.20130223

Gruber, J. (2008). Covering the uninsured in the united states. *Journal of Economic Literature*, *46*(3), 571–606. https://doi.org/10.1257/jel.46.3.571

Gruber, J., & McKnight, R. (2016). Controlling health care costs through limited network insurance plans: Evidence from massachusetts state employees. *American Economic Journal: Economic Policy*, *8*(2), 219–50. https://doi.org/10.1257/pol.20140335

Handel, B. (2013). Adverse selection and inertia in health insurance markets: When nudging hurts. *American Economic Review*, *103*(7), 2643–82. https://doi.org/10.1257/aer.103.7.2643

Handel, B., & Ho, K. (2021). Chapter 16 - the industrial organization of health care markets. In K. Ho, A. Hortaçsu, & A. Lizzeri (Eds.), *Handbook of industrial organization, volume 5* (pp. 521–614, Vol. 5). Elsevier. https://doi.org/10.1016/bs.hesind.2021.11.016

Handel, B., & Kolstad, J. T. (2015). Health insurance for "humans": Information frictions, plan choice, and consumer welfare. *American Economic Review*, *105*(8), 2449–2500. https://doi.org/10.1257/aer.20131126

Hansen, L. P. (1982). Large sample properties of generalized method of moments estimators. *Econometrica*, *50*(4), 1029–1054.

Hausman, J. A. (1996). Valuation of new goods under perfect and imperfect competition. In T. F. Bresnahan & R. J. Gordon (Eds.), *The economics of new goods* (pp. 209–248). University of Chicago Press. https://doi.org/10.7208/9780226074184-007

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

Ho, K., & Lee, R. (2019). Equilibrium provider networks: Bargaining and exclusion in health care markets. *American Economic Review*, *109*(2), 473–522. https://doi.org/10.1257/aer.20171288

Ho, K., & Lee, R. (2023). *Contracting over rebates: Formulary design and pharmaceutical spending* [Working Paper].

Ho, K. (2009). Insurer-provider networks in the medical care market. *American Economic Review*, *99*(1), 393–430. https://doi.org/10.1257/aer.99.1.393

Horn, H., & Wolinsky, A. (1988). Bilateral monopolies and incentives for merger. *RAND Journal of Economics*, *19*(3), 408–419.

Hotelling, H. (1929). Stability in competition. *The Economic Journal*, *39*(153), 41–57.

Klemperer, P. (1987). The competitiveness of markets with switching costs. *The RAND Journal of Economics*, *18*(1), 138–150.

Kling, J. R., Mullainathan, S., Shafir, E., Vermeulen, L. C., & Wrobel, M. V. (2012). Comparison friction: Experimental evidence from medicare drug plans *. *The Quarterly Journal of Economics*, *127*(1), 199–235. https://doi.org/10.1093/qje/qjr055

Liebman, E., & Panhans, M. T. (2021). Why do narrow network plans cost less? *Health Economics*, *30*(10), 2437–2451. https://doi.org/10.1002/hec.4385

Manning, W. G., Newhouse, J. P., Duan, N., Keeler, E. B., & Leibowitz, A. (1987). Health insurance and the demand for medical care: Evidence from a randomized experiment. *The American Economic Review*, *77*(3), 251–277.

Marone, V. R., & Sabety, A. (2022). When should there be vertical choice in health insurance markets? *American Economic Review*, *112*(1), 304–42. https://doi.org/10.1257/aer.20201073

McFadden, D. (1974). Conditional logit analysis of qualitative choice behavior. *Frontiers in Econometrics*, 105–142.

McPake, B., & Noramnd, C. (2008). Parallel systems. In *Health economics: An international perspective* (pp. 245–257). Routledge.

Mussa, M., & Rosen, S. (1978). Monopoly and product quality. *Journal of Economic Theory*, *18*(2), 301–317.

Pauly, M. V. (1968). The economics of moral hazard: Comment. *The American Economic Review*, *58*(3), 531–537.

Polyakova, M. (2016). Regulation of insurance with adverse selection and switching costs: Evidence from medicare part d. *American Economic Journal: Applied Economics*, *8*(3), 165–95. https://doi.org/10.1257/app.20150004

Prager, E. (2020). Healthcare demand under simple prices: Evidence from tiered hospital networks. *American Economic Journal: Applied Economics*, *12*(4), 196–223. https://doi.org/10.1257/app.20180422

Raval, D., Rosenbaum, T., & Tenn, S. A. (2017). A semiparametric discrete choice model: An application to hospital mergers. *Economic Inquiry*, *55*(4), 1919–1944. https://doi.org/10.1111/ecin.12454

Rothschild, M., & Stiglitz, J. (1976). Equilibrium in competitive insurance markets: An essay on the economics of imperfect information. *The Quarterly Journal of Economics*, *90*(4), 629–649.

Serna, N. (2025). *Determinants of provider networks: Risk selection vs. fixed costs* [Working Paper].

Shepard, M. (2022). Hospital network competition and adverse selection: Evidence from the massachusetts health insurance exchange. *American Economic Review*, *112*(2), 578–615. https://doi.org/10.1257/aer.20201453

Spengler, J. J. (1950). Vertical integration and antitrust policy. *Journal of Political Economy*, *58*(4), 347–352. https://doi.org/10.1086/256964

Starc, A., & Swanson, A. (2021). Preferred pharmacy networks and drug costs. *American Economic Journal: Economic Policy*, *13*(3), 406–46. https://doi.org/10.1257/pol.20180489

Tebaldi, P. (2024). Estimating equilibrium in health insurance exchanges: Price competition and subsidy design under the aca. *The Review of Economic Studies*, *92*(1), 586–620. https://doi.org/10.1093/restud/rdae020

Tilipman, N. (2022). Employer Incentives and Distortions in Health Insurance Design: Implications for Welfare and Costs. *American Economic Review*, *112*(3), 998–1037. https://doi.org/10.1257/aer.20181917

Veiga, A., & Weyl, E. G. (2016). Product Design in Selection Markets. *The Quarterly Journal of Economics*, *131*(2), 1007–1056. https://doi.org/10.1093/qje/qjw007