Online Appendix

Vertical Integration and Plan Design in Health Insurance Markets

José Ignacio Cuesta, Carlos Noton, and Benjamín Vatter

A Data and Setting Appendix

This appendix provides additional details about our data, preliminary steps taken to estimate risk and extract negotiated prices, and form choice sets for the plan demand analysis. In addition, we provide details about the Chilean healthcare market and regulatory environment to complement those in the main text.

A.1 Data Construction

A.1.1 Enrollment and Claims Data. Our analysis relies on two main administrative datasets. First, an enrollment dataset including the universe of private plan policyholders, their plan choices, household composition, and a vector of sociodemographic characteristics. Second, a claims dataset that collects each claim processed by private insurers, along with details about the hospital, the services rendered, the prices charged, and plan reimbursement. We restrict our attention to 2013–2016 throughout the paper. This section describes how we move from these raw datasets to the data we use in our analysis.

The raw enrollment dataset includes 1,484,897 policyholders and 61,450 different plans from the five insurers in our sample. Many of these plans were held by a few policyholders, with a median enrollment of 12. Two reasons explain the large number of plans with low enrollment. First, guaranteed renewability implies that enrollees may keep their plan even if not offered in the spot market. In practice, only a minority of plans are ever offered in the spot market in our data. Our analysis mostly focuses on plans offered on the spot market, though we account for legacy plans in plan demand estimation, as discussed in Section 5. Second, insurers offer many nearly identical plans under different codes to limit their exposure to regulation, as discussed in Section A.2. Motivated by these features of the setting, we group the 61,450 reported plan codes by their characteristics: insurer, inpatient and outpatient coverage, main preferential provider, and deciles of the base premium. This grouping captures the key financial elements of plan design, competition, and insurance value relevant to the analysis of VI. Previous work on the Chilean market has adopted similar strategies (Atal, 2019; Dias, 2022). We rely on the claims data to recover plans' preferential providers. The data indicate whether a claim originates at a preferential provider, which we use together with information on consumer cost-sharing to infer coverage levels by tier. In practice, we observe claims for almost all plans in the enrollment data, minimizing the loss from this procedure.

We impose two additional restrictions on the enrollment data. First, whenever a policyholder switches plans within a year, we keep the one held for the longest. Second, we only keep policyholders between 25 and 64 years old. The final sample includes 1,247,125 policyholders and 4,110 plan-years, of which 1,431 are offered in the spot market.

The raw claims dataset includes 18,957,306 inpatient service claims by the insurers in our sample at hospitals in Santiago. We aggregate individual claims into admission events identified in the data. For each admission, we observe ICD-10 diagnosis codes. Ten percent of admissions (eight percent of revenue) lack enough information to classify their diagnosis. Moreover, whenever an admission features services from multiple providers, we allocate the admission to the hospital that provides the most services. We separately identify admissions at the 11 main hospitals in the market, accounting for 74 percent of events. All remaining admissions occur at other small private hospitals or public hospitals, which we collect in the outside option in our analysis. We then aggregate prices and reimbursement across services within an admission. Each admission is linked to the plan covering the patient and the corresponding policyholder.

We clean the claims data and restrict the sample as follows. First, we drop a small number of admissions for which policyholders are not recorded in the enrollment data, or that do not have diagnosis information. Second, we restrict our attention to 16 diagnosis groups that include infections and parasites, neoplasms, blood diseases, endocrine diseases, nervous system diseases, ocular diseases, ear diseases, circulatory diseases, respiratory diseases, digestive diseases, skin diseases, musculoskeletal diseases, genitourinary diseases, pregnancy, perinatal treatments, and congenital malformation. Finally, we only keep admissions by patients below age 65. The final sample includes 773,264 admissions.

A.1.2 Diagnosis Risk. We aggregate diagnosis at the ICD-10 chapter level throughout the analysis. Using our claims data, we compute each beneficiary's average yearly risk of having a medical event within each diagnosis. To do so, we count the average number of medical events per enrollee-diagnosis-year within age-gender groups. We bin ages between 5 and 65 in 5-year intervals. We create separate bins for those under two and those between 2 and 5.

A.1.3 Negotiated Prices and Resource Intensity Weights. We compute negotiated prices and resource intensity weights based on observed total payments to providers per diagnosis following recent work in the literature (e.g., Gowrisankaran *et al.* 2015; Cooper *et al.* 2018). We regress log total hospital payments per medical event on insurer-hospital-year and gender-diagnosis-age fixed effects:

$$\log p_{ehjt}^{\text{Total}} = \log p_{m(j)ht} + \log \omega_{\kappa(i)d(e)} + \varepsilon_{ehjt}$$

where $\log p_{m(j)ht}$ are insurer-hospital-year fixed effects capturing the log negotiated prices, and $\log \omega_{\kappa(i)d(e)}$ are the log resource intensity weights associated with consumer i's type $\kappa(i)$ and the medical condition $d(e)$ associated with medical event e. We normalize weights relative to delivery for a woman aged 25–40. Finally, we recenter the estimated negotiated prices such that the mean predicted total payment equals the observed mean. In estimation, we drop the top 99*th* and bottom 1*st* percentiles of spending to reduce the impact of outliers. The regression R^2 is 0.299.

A.1.4 Insurance Plan Choice Sets and Publicly Insured Consumers. To form plan choice sets, we select the top 70 percent most popular insurance plans sold by each insurer in each market segment and year. On average, this captures 27 percent of the plans but 95.2 percent of the enrollment. Each household's choice set includes these popular plans and their previous year's plan, regardless of whether it is currently offered in the spot market, to conform with Chile's guaranteed renewability regulation.

For each plan, we compute the expected network surplus for each household member. This involves using estimates of consumer preferences for hospital care and their estimated risk to integrate potential medical needs and expected utilities. Based on equation (5) in the main text, we define $\delta_{ihdt|j}^{H} = u_{ihdt|j}^{H} - \epsilon_{ihdt|j}^{H}$ and compute each member's surplus as $WTP_{ijt} = \sum_{d \in D} r_{id} \ln \sum_{h \in H} \delta_{ihdt|j}^{H}$. Given the number of diagnoses and available options, this is a lengthy process. To reduce computational costs, we use a random 30 percent sample of households when estimating plan demand.

We also incorporate data on public insurance enrollees to build a complete picture of the insurance market. We obtain information on the number of public enrollees by age group, gender, dependents, and income quartile from the CASEN survey (CASEN, 2015). We use the waves of 2013, 2015, and 2017 and linearly interpolate the share of public insurance enrollees by group for the gap years. We use the data to create representative public enrollees and assign them to private market segments. We match representative consumers to the modal neighborhood among similar private enrollees and draw representative characteristics for their dependents. We then compute their expected utility from each private plan. Sampling weights are introduced in all subsequent estimations to account for the number of public enrollees represented by each added representative household.

A.2 Details about Health Insurance in Chile

A.2.1 Public and Private Sector Interaction. The public and private health systems operate in a notably isolated manner: Private enrollees account for 97 percent of private hospital revenue and only 3 percent of public hospital revenue (Galetovic and Sanhueza, 2013). Research on sorting across sectors highlights differences in premiums as the key driver of enrollment decisions (Pardo and Schott, 2012).

During our period of study, private and public insurance enrollment have remained steady at around 18 and 76 percent, respectively. While we incorporate substitution between private and public insurance, substitution across the markets is low (Duarte, 2011). The evidence suggests that the public sector is a safety net for private enrollees, who utilize it primarily when unemployed.

A.2.2 The Regulatory Environment. In May 2005, the government introduced Law 20,015 known as *Ley Larga de Isapres*, imposing several regulations on the private insurance sector. This law limited risk pricing by enforcing risk-rating functions on insurers and capping annual premium increases. It also prohibited direct VI by banning insurers from providing healthcare services.

However, Chile's law code has enabled industry players to find legal ways to risk-price consumers and vertically integrate with hospitals. To bypass the pricing step-function rules, insurers often duplicate their plans, selling different versions to different consumers, thus segmenting the market without significant regulatory constraints. To capture these strategies, we aggregate plans based on coverage, premium, and preferential providers rather than identifiers reported by insurers. Similarly, insurers have circumvented the VI ban by forming holdings that own both insurers and hospitals. This ability to bypass regulations has been the focus of much political debate over the last decade.

In September 2004, the government enacted Law 19,966, mandating coverage for a list of medical conditions. Effective June 2005, public and private insurers must provide adequate treatment and insurance for these conditions, known as AUGE-GES guarantees. The law ensures access to adequate treatment and requires hospitals to certify their care quality. It also mandates private plans cover 80 percent and public insurance 100 percent for these conditions. However, private enrollees must file forms and request approval to use these benefits, leading to significantly lower utilization among them (Alvear-Vega and Acuña, 2022). In our analysis, we consider this regulation as an overall shifter of the value of public insurance relative to private insurance and as a determinant of the resource intensity weight multiplier identified in the price regressions.

Regulation also dictates that private plans must offer coverage at least as generous as the public insurance system. However, the enforcement of this requirement is incomplete, and some plans fall below the 60 percent effective coverage of the public option. Due to this imperfect enforcement of regulatory constraints, we model the regulatory environment in plan design as a function to be estimated rather than imposing strict constraints that would not align with the data.

A.2.3 Risk Pricing and Selection. The public insurer distinguishes consumers primarily based on income and household size but does not offer plans that vary based on other factors. In contrast, private insurers offer a variety of plans targeted to specific population subgroups, such as women without dependents, men without dependents, and families. The target group is often made explicit in the plan's name and is (imperfectly) reported to the regulator, which we observe in our data. As noted in the previous section, insurers face regulatory constraints on their ability to charge consumers different premiums for the same plan and are not allowed to reject a consumer based on gender. Nevertheless, they effectively select across groups by carving out coverage for gender-specific treatments (i.e., pregnancy). Thus, by offering different plans to different age and gender groups, insurers circumvent the regulatory pricing constraints. We capture this behavior by treating each gender and age group as a market segment.

The regulation allows private insurers to reject consumers based on pre-existing conditions. We do not incorporate this margin of selection in our analysis as insurers rely on consumer disclosure of pre-existing conditions, which is not in our data. It is also unclear whether consumers have incentives to disclose such conditions or to what extent insurers can monitor them. In recent years, after our sample, one of the largest insurers in the market conducted a marketing campaign to attract enrollees by drastically limiting their consideration of pre-existing conditions (La Tercera, 2020). The effects on enrollment, however, appeared to be negligible, suggesting that selection on pre-existing conditions might not be a first-order problem in the market.

The vast number and heterogeneity of plans available in the private market suggest that insurers use plan design for risk selection. We highlight this force in our main analysis, quantifying the degree to which selection on networks contributes to overall adverse selection in the market. It also plays an implicit role in determining the plan design equilibrium in counterfactuals.

Finally, it is worth noting that this market has only a minimal risk adjustment mechanism. Insurers create a shared fund, pooling expenses related to AUGE-GES conditions and redistributing flows based on spending. The pool is managed and funded exclusively by private insurers. However, as noted in the previous section, AUGE-GES conditions represent a small share of inpatient care spending, making this risk-adjustment pool irrelevant to our analysis. Nevertheless, AUGE-GES represents a relevant source of spending in outpatient care, which is why we model the non-inpatient care cost of plans (which we call administrative cost) as plan-specific rather than varying across consumers.

A.2.4 Coverage and Networks. Hospitals in Chile are paid on a fee-for-service basis, with charges submitted to insurers who then split the bill with the patient according to the plan's cost-sharing terms. These structures include a deductible, a coinsurance rate, and a coverage cap. Deductibles, as usual, require consumers to pay the full bill until the deductible is met. However, in Chile, deductibles are relatively small compared to inpatient prices and thus play a minor role in our analysis. Coverage caps, which limit how much insurers will pay for care in a year, are the counterpart to the maximum out-of-pocket structure in the U.S., which is not present in the Chilean setting. Coverage caps, however, are quite high, and the government provides emergency catastrophic insurance that can often cover the remainder of the bill. Therefore, coinsurance is the main factor determining cost sharing, which we model and report in our analysis.

By regulation, plan networks come in one of three forms. First, there are unrestricted network plans, which provide the same coverage for all hospitals, similar to PPO plans in the U.S. Second, there are preferential provider plans, which offer two-tiered networks with differentiated coverage across different sets of private hospitals, akin to HMO plans in the U.S. Third, there are restricted provider networks, which only cover care at specific providers. Our analysis centers on the first two types of networks as the restricted network plans are very rare, representing less than eight percent of all plan codes and less than four percent of enrollment. It is worth noting that the preferential and non-preferential coverages we use in our analysis summarize plan benefits. In practice, two plans might have the same preferential coverage, but one might fully cover medical imaging and only partially surgeon fees, while the other might cover imaging only partially but fully cover surgeon fees. Importantly, however, coverage across treatment is uniform within each coverage tier. In our analysis, we use the average effective coverage of each plan, which closely matches what consumers observe when purchasing plans. This simplification implies that we cannot speak to the extent to which risk selection is conducted by differential coverage across medical treatments.

Hospitals cannot deny care to insured patients or emergency care regardless of insurance. In effect, all consumers can access all hospitals, though they may receive limited coverage from their plan. Overall, private insurers offer higher coverage in private hospitals, which are generally perceived as being of higher quality than public hospitals in terms of waiting time, medical resources, and medical outcomes.

A.3 Additional Descriptive Statistics

A.3.1 Insurance Market. Table A.2-A summarizes the variation in paid premiums and market shares across insurers. We document substantial variation in premiums across insurers, with the difference between the highest and the lowest insurer average premiums being 36 percent of the average premium in the market. Furthermore, we observe significant variation in premiums within an insurer, likely driven by a combination of preference heterogeneity leading to compositional differences and the corresponding risk-rating and cost differences.

A.3.2 Hospital Market. Table A.2-B documents substantial price and market share variation across hospitals. The industry is moderately concentrated: no hospital has a market share higher than 13 percent, five hospitals have market shares between 8 and 13 percent, while the rest have market shares below 5 percent. The overall HHI is 1,387. The outside option—smaller private hospitals and all public ones—has a market share of 26 percent. Dispersion in admission prices across hospitals is substantial. For example, hospitals h_1 and h_6 charge average prices more than double those charged by h_4 and h_{11}. Differences in location, infrastructure, and real and perceived quality explain this price dispersion.

A.3.3 Plan Tiering. Table A.3-A describes the structure of plan preferential tiers by documenting the share of plans that have each hospital in its preferential tier. In most cases, VI insurers are more likely to place their integrated hospitals in the preferential tier of their plans relative to rival hospitals. For instance, m_a places its integrated hospitals in the preferential tier of between 37 and 69 percent of its plans, and m_b does so between 49 and 96 percent. In contrast, averaging across insurer-hospital combinations shows hospitals are preferential in only 20 percent of plans. Notably, regardless of vertical linkages, most hospitals feature in the preferential tier of plans offered by all insurers.

A.3.4 Admission Flows. Table A.3-B displays the breakdown of hospital admissions per insurer and shows that this pattern holds when looking at individual firms. For each integrated hospital, its integrated insurer is the dominant source of admissions, accounting for between 43 and 71 percent. Nevertheless, all integrated hospitals receive a substantial share of patients from integrated and non-integrated rival insurers, and all non-integrated hospitals receive patients from integrated insurers.

A.3.5 Comparison Between VI and non-VI patients. One explanation for outcome differences across admissions covered by insurers that are VI and non-VI with a hospital is that patients may differ in observables. To study this possibility, we estimate different versions of equation (2) in the main text, using patient observables as the dependent variable. Finding differences in observables depending on VI would suggest that the patient groups we are comparing are not balanced. Table A.4 shows the results.

The unconditional comparison in column (1) shows that patients from VI insurers at their integrated hospital are almost three years older, 6 percentage points less likely to be female, equally likely to be employed, and have 16 percent lower income than other patients. However, once we control for the set of fixed effects in equation (2), those differences get much closer to zero and become less statistically significant. Moreover, once we include an insurer-hospital fixed effect, VI and non-VI insurer patients only differ in age. The former are around one year younger on average and are otherwise balanced. This suggests that our comparison is based on groups of patients with similar characteristics. Regardless, we control for these observables in our analysis in Section 3.[1]

A.4 Additional Evidence on VI and Admission Outcomes

We complement the analysis of Section 3 by exploring whether VI affects the provision of services for which physicians enjoy some discretion. We focus on C-sections and ultrasounds during pregnancy, hemogram tests in digestive diagnoses, and chest X-rays and cross-section imaging in respiratory diagnoses. C-sections provide a particularly convenient setting for this test as, upon delivery, a physician must either implement a C-section or proceed with natural birth—there is no extensive margin decision. Therefore,

[1]Note that the regression R^2 increases substantially once location, hospital, plan, and diagnosis controls are included in column (2). This suggests observables explain patient sorting across hospitals, which suggests accounting for observable heterogeneity along these dimensions in demand. Our demand estimates in Section 5 show substantial observable heterogeneity in preferences over plans and hospitals.

we know the exact alternatives for physicians in these cases, which is not as evident for the other services we study. Table A.6-A displays the results, which provide mixed evidence on the association between VI and the provision of these services. On the one hand, we find that admissions from VI insurers tend to provide fewer C-sections, consistent with evidence from the U.S. (Cutler *et al.*, 2000; Johnson and Rehavi, 2016). On the other hand, we find the opposite effect on imaging and no significant effect on ultrasounds, hemograms, and chest X-rays. Moreover, once we focus on hospital-insurer variation only by including hospital-insurer fixed effects, the coefficients become less statistically significant, as shown by Table A.6-B. Taken together, these results do not suggest a strong relationship between VI and hospital provision of services. Hence, we do not model these dimensions of hospital behavior explicitly. However, in our counterfactual analysis in Section 6, we consider the role of cost efficiencies associated with VI.

B Model Appendix

This appendix presents additional details about how we formulate and implement our model. It also includes the proofs for the statements made in the main text. Appendix C provides details about model estimation and simulation.

B.1 Discussion of Model Assumptions

B.1.1 Connection with the Regulatory Environment. This section outlines the connection between our model and the setting. The most fundamental issue distorting regulatory control in the market is insurers' ability to duplicate plans and selectively offer them to consumers. For example, the law requires insurers to set a base premium for each plan and establish up to two price schedules determining how each plan's premium changes with gender and age. The intent is to control age and gender discrimination by tying insurers' hands to the same schedule across plans. Nevertheless, because insurers can duplicate a plan and set base premiums freely, they can simply sell the same plan under different codes and thus different prices across gender and age groups without meaningful constraints. In fact, many of the plans in our sample are branded with gender names and family labels. Thus, rather than follow the regulatory model for pricing, we adapted it to what insurers use. We examined insurers' websites and marketing materials and approximated their age, gender, and family segments.

Our model also simplifies other regulatory margins. First, regulation caps the extent to which insurers can increase premiums every year. The regulation's enforcement, however, requires enrollees to submit formal complaints during a specific time window.[2] Either because of difficulties in submitting the complaints or due to lax enforcement, insurers in our study period do not appear to be constrained by these caps.[3] In effect, in more recent years, enrollees resorted to suing insurers for violating these limits.

Second, guaranteed renewability regulation stipulates plans have an indefinite duration. The law dictates that insurers must outline conditions under which they can change coverage and access and must notify consumers within 60 days of changes. Conversely, enrollees must notify insurers 30 days before exiting an insurance contract. This guaranteed renewability, in principle, would allow enrollees to retain their plans indefinitely. In practice, however, insurers can increase premiums for enrollees they wish to displace. In our model, we accommodate guaranteed renewability in two simplified ways. When estimating demand, we allow policyholders to enroll in their previous year's plan regardless of whether it was offered in the spot market. We exclude this behavior when analyzing the equilibrium effects of VI using our model, treating all policyholders as new enrollees. In addition, we incorporate legal penalties for disagreements between insurers and hospitals to capture violations of plan coverage guaranteed renewability.

Third, we also simplify the extent to which insurers may reject enrollees based on risk. In practice, insurers ask enrollees to report preexisting medical conditions and can reject applicants or deny coverage on existing conditions. However, the impact of this selection is unclear, as we discuss in Section A.2. Thus, we choose not to model the margin of insurance rejection in our model.

Finally, as we note in Section A.2, there is a limited risk pool that compensates insurers for spending differences on a select list of conditions. Our analysis does not model risk adjustment as few of these conditions are treated in an inpatient setting. These conditions, however, are an important source of outpatient spending, which justifies modeling the non-inpatient care cost (or administrative cost) as plan-specific and invariant across individuals.

B.1.2 Value of Insurance. We impose two assumptions about the value of insurance. First, we only model allocative moral hazard spending. Consumers do not forgo care in response to higher out-of-pocket prices but rather seek care at the outside option. A cheap public option makes it the relevant alternative for patients. Second, we model the value of risk protection only through consumer network surplus. We do not model risk aversion

[2] For details, see Chile Atiende (2024).

[3] This limited enforcement is well documented in the press from the time (CIPER, 2013).

explicitly, as the regulatory pressure to provide generous plans limits the extent to which coverage changes in counterfactuals. Instead, regulation leads firms to use coverage to steer demand and affect negotiated prices. Therefore, we expect the first-order impacts of VI to be on prices, premiums, and choices rather than risk protection. Our model builds on previous work on price competition and healthcare consolidation, allowing us to benchmark our analysis against known results. Our model describes the effects of VI on a regulated insurance market, similar to managed competition (Tebaldi, 2024).

B.1.3 Exogenous Hospital Costs and Quality. As discussed in Section 3, the data do not suggest meaningful differences in costs, treatment, or quality at VI hospitals for patients covered by their integrated insurers. Additionally, VI hospitals in our setting are neither the highest nor the lowest quality providers. Therefore, we chose not to endogenize cost or quality. We discuss the impact of alternative assumptions regarding cost efficiencies and quality improvements in Section 6.4 of the main text.

B.2 The Pricing Subgame

We formulate the pricing subgame through two nested fixed-point conditions to improve computational efficiency and respect our equilibrium refinement. The outer layer corresponds to the stacked conditions associated with the optimality of negotiated and VI-optimal prices, while the inner loop operationalizes the premium setting condition. In what follows, we ignore time indices when possible to reduce the notational burden.

B.2.1 Premiums. We use the approach of Morrow and Skerlos (2011) to form a fixed point formulation of the optimality condition associated in equation (10). This results in:

$$\Phi \quad = \Lambda(\Phi)^{-1}[\Gamma(\Phi)\Phi + C(\Phi)]$$

where Φ is the stacked vector of all plan premiums, Λ is a diagonal matrix with j-th element equaling $\sum_{i\in I}|F_i|^2\alpha_i^M D_{ij}^M(\Phi)$, and Γ is a sparse matrix with (j,j')-th entry equal to $\sum_{i\in I}|F_i|^2\alpha_i^M D_{ij}^M(\Phi)D_{ij'}^M(\Phi)$ if plan j and j' belong to the same insurer and zero otherwise. Finally, $C(\cdot)$ is a vector capturing all cost and VI elements:

$$C(\Phi)[j'] = \sum_{i\in I}|F(i)|D_{ij'}^M(\Phi)\left[\alpha_i^M(ec_{ij'} + \eta_j - \sum_{j\in J_m} D_{ij}^M(\Phi)(ec_{ij} + \eta_j)) - 1 - \alpha_i^M\Big(\pi_{ij'}^{H|m(j')} - \sum_{j\in\mathcal{J}} D_{ij}^M(\Phi)\pi_{ij}^{H|m(j')}\Big)\right]$$

where ec_{ij} is consumer i's expected inpatient cost as described in equation (7) and $\pi_{ij}^{H|m(j')}$ corresponds to insurer $m(j')$'s weighted profits at integrated hospitals from claims paid by plan j if individual i enrolls in it: $\pi_{ij}^{H|m(j')} = \sum_{h\in H} \theta_{m(j')h} \sum_{i'\in F_i} \sum_{d\in D} r_{i'd} D^H_{i'hd|j} \omega_{i'd}(p_{m(j)h} - k_{hm(j)})$.

As noted by Morrow and Skerlos (2011), while there are no theoretical guarantees that this formulation results in a contraction mapping, it appears to be so in practice. We obtain consistent, rapid convergence to the same stable set of premiums, regardless of the starting point, initial price, and coverage values.

B.2.2 Prices. For any hospital system $s \subset H$, we follow the logic of Gowrisankaran *et al.* (2015) to express the optimality conditions associated with negotiated prices and VI-optimal prices (if any) as:

$$\boldsymbol{p}_s = \boldsymbol{k}_s - (\nabla_{p_s} D_s^{Ht} \cdot diag(\boldsymbol{\theta}_s) + diag(\chi_s) \cdot \Lambda_s)^{-1}(diag(\boldsymbol{\theta}_s) \cdot \tilde{D}_s + \nabla_{p_s} \pi^M_{m(s)} + diag(\chi_s) \cdot \Gamma_s^{VI})$$

where $\boldsymbol{p}_s$ and $\boldsymbol{k}_s$ are the system's vector of prices and hospital costs across the system's hospitals and all insurers, D_s^H is the vector of total resource-weighted expected hospital demand from each insurer, $diag(\cdot)$ is the diagonalization operator, and χ_s is an indicator vector that equals one when the hospital and insurer in each row are not VI. Finally, Λ and Γ are defined as:

$$\begin{aligned}
\Lambda_s[i,j] &= \left(\frac{\tau_h}{1-\tau_h} \frac{\frac{\partial \pi^M_m}{\partial p_{h,m}} + \sum_{h'\in H} \theta_{h'm} \frac{\partial \pi^H_{h'}}{\partial p_{h,m}}}{\Delta_{m,s} V^M_m}\right)[i] \times \theta_{h'}[j] \Delta_{m[i],s} \tilde{D}_{m'h'}[j]^H \\
\Gamma_s^{VI}[i] &= \Big(\frac{\tau_h}{1-\tau_h} \frac{\frac{\partial \pi^M_m}{\partial p_{hm}} + \sum_{h'\in H} \theta_{h'm} \frac{\partial \pi^H_{h'}}{\partial p_{h,m}}}{\Delta_{m,s} V^M_m} \Delta \pi^M_{m(s)}\Big)[i] \\
\Delta_{m,s} V^M_m[i] &= \Delta_{m[i],s} \pi^M_{m[i]} + \sum_{h'\in H} \theta_{h'm[i]} \Delta_{m[i],s} \pi^H_{h'} + l_{m[i]h[i]} \Delta_{m[i]s} WTP_{m[i]s}
\end{aligned}$$

where $\Delta_{ms} WTP_{ms}$ is the total expected loss in network surplus from removing hospital system s's members from all of m's networks. When computing this value, we assume the courts would compute consumers' expected network surplus using the average prices negotiated by rival insurers with the removed hospitals. We also assume they will use the realized insurance demand (under disagreement) to compute total losses. These assumptions mostly affect the identified distribution of multiplier L. Our estimates remain largely unchanged if we assume the courts use the full-agreement insurance demand or previous hospital prices as the basis for computing network surplus losses.

Unlike the premium fixed point formulation, we find that the expression above is a contraction mapping only locally around the equilibrium. To address this, we use a globally convergent Anderson Acceleration routine (Zhang *et al.*, 2020). This results in an efficient pricing subgame equilibrium solver, with the outer loop solving the more challenging price fixed-point problem and the inner loop providing equilibrium premiums conditional on the current guess of hospital prices.

B.3 Solving the Plan Design Problem

We begin by proving Proposition 1 from the main text, which is the central result for solving the plan design problem.

Proof of Proposition 1. We begin by restating the original problem for completeness. Denote J_m the set of plans to be optimized and the firm's objective $V(C) = \tilde{\pi}_m(\boldsymbol{\phi}^*(C), \boldsymbol{p}^*(C), C) - \sum_{j \in J_m} M_j(K^r_m(C) + K^o_m(C))$, where $C \in [0,1]^{|J_m| \times |H|}$ is the matrix of coverages from each plan at each hospital. Denote the set of tiered coverages as $\mathcal{C}_m = \{C \in [0,1]^{|J_m| \times |H|} | C_{(j,\cdot)} \in \mathcal{C}\}$ where $C_{(j,\cdot)}$ is the row associated with plan j and $\mathcal{C}$ is defined as in equation (12). Thus, the original combinatorial optimization problem established in equation (12) can be expressed as $\mathcal{P}^* : \max_{C \in \mathcal{C}_m} V(C)$.

Let $\mathcal{P}^0$ denote the following continuous-control non-convex optimization problem:

$$\max_{\underline{\boldsymbol{c}}, \bar{\boldsymbol{c}} \in [0,1]^{|J_m|}, W \in [0,1]^{|J_m| \times |H|}} V(diag(\underline{\boldsymbol{c}})(1-W) + diag(\bar{\boldsymbol{c}})W) \tag{1}$$

$$s.t \quad W_{jh}(1 - W_{jh}) = 0 \quad \forall j \in J_m, h \in H \tag{2}$$

Denote $\mathcal{C}^0_m = \{C \in [0,1]^{|J_m| \times |H|} \,\big|\, \exists \underline{\boldsymbol{c}}, \bar{\boldsymbol{c}} \in [0,1]^{|J_m|}, W \in [0,1]^{|J_m| \times |H|}$, s.t. $C = diag(\underline{\boldsymbol{c}})(1-W) + diag(\bar{\boldsymbol{c}})W\}$. Note that $\mathcal{C}_m \subset \mathcal{C}^0_m$ as any element of $\mathcal{C}_m$ corresponds to an element of $\mathcal{C}^0_m$ where the weight matrix W is an extreme point of its domain. Moreover, note that any element of $\mathcal{C}_m$ has a representation $(\underline{\boldsymbol{c}}, \bar{\boldsymbol{c}}, W)$ that satisfies the tiering constraint $W_{jh}(1 - W_{jh}) = 0$ and any element of $\mathcal{C}^0_m$ that satisfies the constraint is an element of $\mathcal{C}_m$.

As $V(\cdot)$ is continuous and the set $\mathcal{C}^0_m$ is compact, the Weierstrass extreme value theorem guarantees a solution to $\mathcal{P}^0$ exists. Therefore, $\arg\max \mathcal{P}^* = \arg\max \mathcal{P}^0$ necessarily and neither are empty. By contradiction, suppose that $\exists \tilde{C} \in \arg\max \mathcal{P}^*$ not in $\arg\max \mathcal{P}^0$. Then there exists $\hat{C} \in \arg\max \mathcal{P}^*$ such that $V(\hat{C}) > V(\tilde{C})$ and moreover $\hat{C}$ satisfies the tiering constraint, hence it is an element of $\mathcal{C}_m$. Analogously, if there exists $C^0 \in \arg\max \mathcal{P}$ not in $\arg\max \mathcal{P}^*$, then there exists $C' \in \mathcal{C}_m$ such that $V(C') > V(C^0)$ that satisfies the tiering constraint. Therefore, problem $\mathcal{P}^*$ and $\mathcal{P}^0$ are equivalent.

Let $\lambda > 0$ and $G(\cdot)$ be a positive, continuous, strictly monotonic function. Without loss, normalize $G(0) = 0$. Define $\mathcal{P}^1(\lambda)$ the original statement of the proposition:

$$\max_{\underline{c},\bar{c}\in[0,1]^{|J_m|},W\in[0,1]^{|J_m|\times|H|}} V(diag(\underline{c})(1-W) + diag(\bar{c})W) - \lambda \sum_{j\in J_m}\sum_{h\in H} G(W_{jh}(1-W_{jh})) \tag{3}$$

Again, this is a continuous optimization problem over a compact domain. Hence, a solution exists. Moreover, by Berge's maximum theorem, $\arg\max \mathcal{P}^1(\lambda)$ is continuous in λ. Note that any solution to $\mathcal{P}^0$ attains the minimum tiering penalty $\sum_{j\in J_m}\sum_{h\in H} G(W_{jh}(1 - W_{jh})) = 0$ and that for any convergent sequence of solution to $\mathcal{P}^1$, $C^1(\lambda)$, the tiering penalty at $C^1(\lambda)$ must be weakly decreasing in λ. Therefore, by the upper hemicontinuity of the solution to $\mathcal{P}^1(\lambda)$ established by the maximum theorem, $\lim_{\lambda\to\infty} \arg\max \mathcal{P}^1(\lambda) = \arg\max \mathcal{P}^0 = \arg\max \mathcal{P}^*$. □

The proof above contains a simple idea. It states that whether insurers face a strict requirement to submit tiered networks or a penalty for submitting plans with untiered networks is equivalent. The proposition establishes a sequence of design problems under increasing penalties, which allows us to trade off exploration versus optimality when solving the problem. The rate at which these problems converge as the penalty λ increases is inherently tied to the value firms have from offering intermediate tiers of coverage. In our setting, the value of doing so is low because insurers use tiers to steer consumers to specific hospitals. To do so effectively, they must introduce a large wedge between the out-of-pocket price at the hospitals they want consumers to visit and those they do not. Designing plans with coverages between the preferential and base tiers counteracts the efforts to steer consumers. In Appendix C, we discuss how we implement the convergent sequence of design problems and find intersections of insurers' best responses.

B.4 Identification of Price and Premium Parameters

This section formalizes the identification of the price- and premium-setting parameters, discussed in Section 5. Before proceeding with the statements, we introduce some additional notation. Denote the total weighted demand to hospital h from insurer m as $D^H_{h|m} = \sum_{j\in J_m}\sum_{i\in I} D^M_{ij} \sum_{i'\in F_i}\sum_{d\in D} r_{i'd}\omega_{i'd} D^H_{i'hd|j}$. Denote the vector of total weighted demand from each insurer to hospital h as $\boldsymbol{D}^H_h$. Denote the total demand for plan j, and $\boldsymbol{D}^M_m$ the vector of demand for each plan of insurer m as $D^M_j = \sum_i D^M_{ij}$. Denote the difference in demand from hospital h from insurer m' under full agreement and under disagreement between the pair (m, h) as $\Delta_{mh} D^H_{h|m'}$. Denote the matrix of demand for hospital h from each insurer (columns) under every potential disagreement (rows) as $\Delta_{h|M}\boldsymbol{D}_h^H$. Denote the vector of premiums of insurer m as $\boldsymbol{\phi}_m$, the vector of hospital prices for h from all insurers by $\boldsymbol{p}_h$, and the vector of prices between insurer m and its integrated hospitals (if any) as $\boldsymbol{p}_m^{VI}$. Given this notation, we impose the following regularity assumptions on the problem:

Assumption 1.

(a) *(substitution): For every h, m, $\nabla_{\boldsymbol{\phi}_m}\boldsymbol{D}_m^M$ and $\nabla_{\boldsymbol{p}_m}\boldsymbol{D}_h^H$ are negative definite, $\Delta_{h|M}\boldsymbol{D}_h^H$ is positive definite.*

(b) *(gains from trade): For any pair (h, m), $\Delta_{mh}\pi_m^M > 0$, $\frac{\partial \pi_m^M}{\partial p_{hm}} \leq 0$, $\Delta_{mh}\pi_h^H > 0$, and $\frac{\partial \pi_h^H}{\partial p_{mh}} \geq 0$.*

(c) *(rank): For any VI insurer m, the gradient matrix $\nabla_{[\boldsymbol{\phi}_m, \boldsymbol{p}_m^{VI}]}[\boldsymbol{D}_m^M, \boldsymbol{\pi}_m^M]$ is column rank.*

Assumption 1(a) establishes basic requirements on substitution patterns, which are common in demand theory (Kihlstrom *et al.*, 1976). Inherently, they require strong substitution away from products as prices and premiums increase and strictly positive spillovers to competing products. Note that the negative definiteness of plan demand is required separately by market segment and, hence, does not require plans in other market segments to be affected by competition for a separate set of consumers. Assumption 1(b) requires negotiation between hospitals and insurers to be individually rational for each party and to feature tension in optimal payments. Negotiations failing individual rationality should not occur, and those without tension do not require bargaining to be resolved. Finally, Assumption 1(c) establishes basic rank requirements analogous to those found in the analysis of linear regressions. It is essentially a statement about the data featuring sufficient variation for identification.

In addition to these assumptions, we introduce two additional assumptions required only for our specialization of the bargaining model. To do so, we introduce the additional structure of our model and disagreement penalties. The following notation describes the key identified elements of demand and hospital profits that are at stake during negotiation:

$$
\begin{aligned}
v_{hm}^H &= \frac{\sum_{m'} \Delta_{mh} D_{h|m'}^H p_{hm'} - \frac{\sum_{m''} \Delta_{mh} D_{h|m''}^H}{\sum_{m''} D_{h|m''}^H} k_h^o}{\sum_{m'} \partial_{mh} D_{h|m'}^H p_{hm'} - \frac{\sum_{m''} \partial_{mh} D_{h|m''}^H}{\sum_{m''} D_{h|m''}^H} k_h^o} \\
D_{hmm'}^{\partial} &= D_{h|m'}^H \sum_{m''} \partial_{mh} D_{h|m''}^H - \partial_{mh} D_{h|m'}^H \sum_{m''} D_{h|m''}^H \\
D_{hmm'}^{\Delta} &= D_{h|m'}^H \sum_{m''} \Delta_{mh} D_{h|m''}^H - \Delta_{mh} D_{h|m'}^H \sum_{m''} D_{h|m''}^H
\end{aligned}
$$

where k_h^o is the observed average hospital cost and ∂_{mh} is shorthand for the partial derivative with respect to p_{mh}.

Assumption 2.

(a) (rank): Let X be a matrix with row elements $[v_{hm}^H, D_{hm1}^{\partial}, \ldots, D_{hmM}^{\partial}, D_{hm1}^{\Delta}, \ldots, D_{hmM}^{\Delta}]$*, for every pair* (h, m)*. X is column rank.*

(b) (independence): The disagreement penalty multiplier l_{hm} *is drawn independently across pairs* (h, m) *from a distribution L with finite variance. Moreover,* l_{hm} *is independent of hospital and insurance demand and costs.*

As before, Assumption 2(a) is essentially a requirement of sufficient variation. Assumption 2(b) requires strong independence of the regulatory penalty from other relevant random factors. This assumption is only used to establish a regression residual in Lemma 3, and hence can be weakened to incorporate heteroskedasticity or endogeneity using standard results from the literature on non-linear regressions (Amemiya, 1983). We can now state our key results, which we prove below.

Proposition 1. *Let Assumptions 1 and 2 hold, and hospital costs be decomposed as* $k_{hmt} = k_{ht}^H + \tilde{k}_{mt}^H$*, then* $(\theta_{mht}, k_{hmt}, \tau_h, \eta_{jt})$ *are point identified and L is non-parametrically identified.*

The proof proceeds in three steps. First, we show that elements specific to VI firms are point-identified up to VI hospital costs. Therefore, we can focus attention on identifying these objects in a setting without VI, as the added complexity of VI firms is resolved through the optimality condition of VI prices and premiums. Second, we show that if there were no disagreement penalties for insurers, then a more general version of our model is exactly identified from the optimality conditions associated with prices and our additional data on total hospital costs. Third, we show that these general identification results imply that our more constrained model is also identified. Finally, the key proposition follows as a corollary of these three lemmas. Our identification arguments below are fully cross-sectional. Hence, we omit the time subscript from all equations.

We begin by resolving the identification of insurer administrative costs η_{jt} and VI firms' objective weights θ_{hmt}. The following lemma is a direct consequence of the monotonicity of premiums and prices in these parameters.

Lemma 1. *Let Assumption 1 hold. For any non-VI insurer,* η_{jt} *is point-identified from the premium setting condition. For any VI insurer, given VI-hospital costs* k_{hmt}*, administrative costs* η_{jt}*, and VI weights* θ_{hmt} *are point-identified.*

Proof. Denote a plan's profit from insurance net of administrative cost as:

$$v_j = D^M_{ij}(|F_i|\phi_j - \sum_{i'\in F_i}\sum_{d\in D} r_{i'd}\sum_{h\in H} D^H_{i'hd|j}(\boldsymbol{p}_m, \boldsymbol{c}_j)c_{jh}\omega_{i'd}p_{mh})$$

And hence, we can write a firm's objective as:

$$\pi^{VI}_m = \sum_{j\in J_m}(v_j - D^M_j\eta_j) + \sum_{h\in H}\pi^H_h(\boldsymbol{k}_h)\theta_{hm}$$

We can use this expression to evaluate the insurer's premium optimality condition and VI firm's price optimality condition. Stacking these conditions, we get

$$\nabla^t_{\phi_m}\boldsymbol{v}_m - \nabla^t_{\phi_m}\boldsymbol{D}^M_m\boldsymbol{\eta}_m + \nabla^t_{\phi_m}\boldsymbol{\pi}^H_m(\boldsymbol{k_m})\boldsymbol{\theta}_m = 0 \tag{4}$$

$$\nabla^t_{p^{VI}_m}\boldsymbol{v}_m - \nabla^t_{p^{VI}_m}\boldsymbol{D}^M_m\boldsymbol{\eta}_m + \nabla^t_{p^{VI}_m}\boldsymbol{\pi}^H_m(\boldsymbol{k_m})\boldsymbol{\theta}_m = 0 \tag{5}$$

where bold letters denote vectors, $\boldsymbol{\pi}^H_m$ denotes the vector of hospital profits for hospitals integrated with insurer m, and $\nabla_{p^{VI}_m}$ is the gradient with respect to prices between insurer m and its integrated hospitals. Equation (4) establishes a linear system of equations in $\boldsymbol{\theta}_m$ and $\boldsymbol{\eta}_m$. For non-VI insurers, equation (4) resolves the identification of $\boldsymbol{\eta_m}$, as $\nabla^t_{\phi_m}\boldsymbol{D}^M_m$ is square and invertible by Assumption 1(a). For VI firms, equation (5) provides an additional linear system of equations with as many equations as hospitals are integrated with insurer m. Moreover, by Assumption 1(c), the system has column rank. As π^H_h is linear in costs, this establishes that, conditional on hospital costs, there is a unique solution to $\boldsymbol{\theta}_m, \boldsymbol{\eta}_m$. □

The previous lemma states that the VI-specific parameters are point-identified conditional on identifying hospital costs. Hence, we can focus on identifying costs and safely ignore any complexities imposed by VI in this problem. The following lemma establishes that, given our additional data on total hospital costs, the standard Nash-in-Nash bargaining model is exactly identified.

Lemma 2. *Let Assumption 1 hold. If there are no disagreement penalties ($l = 0$) and no-VI firms, then (τ_h, k_{hm}) are identified for every h, m.*

Proof. Denote $\lambda_h = (1-\tau_h)/\tau_h$ and note that Nash bargaining optimality conditions imply that for any pair h, m, we have:

$$\frac{|\partial\pi^M_m/\partial p_{hm}|}{\Delta_{mh}\pi^M_m} = \lambda_h\frac{\partial\pi^H_h/\partial p_{hm}}{\Delta_{mh}\pi^H_h}$$

where, by Lemma 1, the left-hand side is identified, and hence we can treat it as a known component $\hat{\alpha}_{hm} = \frac{|\partial\pi^M_m/\partial p_{hm}|}{\Delta_{mh}\pi^M_m}$. Therefore, we can rewrite the equation above as:

$$\begin{aligned}
\hat{\alpha}_{hm} &= \lambda_h \frac{\sum_{m' \in M} \partial D^H_{h|m'}/\partial p_{hm}(p_{hm'} - k_{hm'}) + D^H_{h|m}}{\sum_{m \in M} \Delta_{mh} D^H_{h|m'}(p_{hm'} - k_{hm'})} \\
\hat{\alpha}_{hm} \sum_{m' \in M} \Delta_{mh} D^H_{h|m'}(p_{hm'} - k_{hm'}) &= \lambda_h (\sum_{m' \in M} \partial D^H_{h|m'}/\partial p_{hm}(p_{hm'} - k_{hm'}) + D^H_{h|m}) \\
-\lambda_h D^H_{h|m} + \sum_{m' \in M} (\hat{\alpha}_{hm}\Delta_{mh} D^H_{h|m'} - \lambda_h \partial D^H_{h|m'}/\partial p_{hm}) p_{hm'} &= \sum_{m' \in M} (\hat{\alpha}_{hm}\Delta_{mh} D^H_{h|m'} - \lambda_h \partial D^H_{h|m'}/\partial p_{hm}) k_{hm'}
\end{aligned}$$

which we stack over all negotiations of hospital h to obtain:

$$\begin{aligned}
-\lambda_h \boldsymbol{D}^H_h + (\hat{\boldsymbol{\alpha}}_h \Delta_{h|M} D^H_h - \lambda_h \nabla^t_{\boldsymbol{p}_h} D^H_h)\boldsymbol{p}_h &= (\hat{\boldsymbol{\alpha}}_h \Delta_{h|M} D^H_h - \lambda_h \nabla^t_{\boldsymbol{p}_h} D^H_h)\boldsymbol{k}_h \\
-(\lambda_h \hat{\boldsymbol{\alpha}}_h \Delta_{h|M} D^H_h - \nabla^t_{\boldsymbol{p}_h} D^H_h)^{-1} \boldsymbol{D}^H_h + \boldsymbol{p}_h &= \boldsymbol{k}_h
\end{aligned} \tag{6}$$

where invertibility in the last row follows from Assumption 1(a). In this expression, D^H_h denotes the vector of total weighted demand for hospital h from each insurer, $\boldsymbol{\alpha}_h$ is a diagonal matrix with elements α_{hm}, and $\Delta_{h|M} D^H_h$ corresponds to the matrix of differences in hospital demand from each insurer (columns) under each negotiation (rows). Note that Gowrisankaran *et al.* (2015) derived an equivalent condition, albeit for a different purpose.

Our data includes observations of total hospital costs expressed as $k^o_h = \boldsymbol{D}^{Ht}_h \boldsymbol{k}_h$. We can rewrite equation (6) as:

$$\boldsymbol{D}^{Ht}_h (\lambda_h \hat{\boldsymbol{\alpha}}_h \Delta_{h|M} D^H_h - \nabla^t_{\boldsymbol{p}_h} D^H_h)^{-1} \boldsymbol{D}^H_h = \boldsymbol{D}^{Ht}_h \boldsymbol{p}_h - k^o_h \tag{7}$$

By Assumption 1(b), $\hat{\boldsymbol{\alpha}}_h$ is a diagonal matrix of positive entries, and by Assumption 1(a) $\Delta_{h|M} D^H_h$ is positive definite. Hence, the first term on the left is monotonic, decreasing in λ_h. Therefore, equation (7) has a unique solution for λ_h. Substituting this solution in equation (6) establishes the uniqueness of $\boldsymbol{k}_h$. □

Lemma 2 is the central identification lemma as it establishes that the Nash-in-Nash pricing condition, augmented with partial data on hospital costs as in Ho and Lee (2017), delivers identification of the model. Moreover, the identification relies exclusively on cross-sectional variation. Hence, some of the structural assumptions made in our model can be relaxed to leverage time-series information. For example, the argument implies we could identify a time-varying hospital-specific bargaining parameter, τ_{ht}. The following corollary follows directly from combining the two previous lemmas.

Corollary 1. *If there are no disagreement penalties ($l = 0$), then $(\tau_h, k_{hm}, \theta_{hm}, \eta_j)$ are point-identified.*

Finally, we can establish the identification of our specialized model without VI.

Lemma 3. *Let Assumptions 1 and 2 hold, and assume there are no VI firms, and $k_{hm} = k_h + \tilde{k}_m$. Then $(\tau_h, k_h, \tilde{k}_m)$ are point-identified and L is non-parametrically identified.*

Proof. First, note that the observation of k_h^o implies $k_h = \frac{k_h^o - \sum_m D_{hm}^H \tilde{k}_m}{\sum_m D_{hm}^M}$, hence we only need to identify L, $\tilde{k}_m$, and τ_h from the bargaining optimality condition, which is now given by:

$$\frac{|\partial \pi_m^M / \partial p_{hm}|}{\Delta_{mh}\pi_m^M + l_{hm}\Delta_{mh} WTP_{mh}} = \frac{1-\tau_h}{\tau_h} \frac{\partial \pi_h^H / \partial p_{hm}}{\Delta_{mh}\pi_h^H}$$

and can be rewritten as:

$$\frac{\Delta_{mh}\pi_m^M}{\Delta_{mh} WTP_{mh}} = -l_0 + \frac{\tau_h}{1-\tau_h} \frac{\Delta_{mh}\pi_h^H}{\partial \pi_h^H / \partial p_{hm}} \frac{|\partial \pi_m^M / \partial p_{hm}|}{\Delta_{mh} WTP_{mh}} - \tilde{l}_{hm} \tag{8}$$

where $l_0 = \mathbb{E}[l]$ and $\tilde{l}_{hm}$ is the residual which, by Assumption 2(b), satisfies $\mathbb{E}\Big[\tilde{l}| \frac{|\partial \pi_m^M / \partial p_{hm}|}{\Delta_{mh} WTP_{mh}}\Big] = 0$.Therefore, equation (8) is a non-linear regression equation. With some manipulation, it is easy to see that this non-linear regression equation is of second order and has an equivalent regression:

$$y_{mh} = \alpha\Big(1 + \sum_{m' \in M} \tilde{k}_{m'} D_{hmm'}^{\partial}\Big) + \beta_h x_{hm}\Big(v_{hm}^H + \sum_{m' \in M} \tilde{k}_{m'} D_{hmm'}^{\Delta}\Big) - y_{mh}\Big(\sum_{m'} \tilde{k}_{m'} D_{hmm'}^{\partial}\Big) + \eta_{hm} \tag{9}$$

where $y_{hm} = \frac{\Delta_{mh}\pi_m^M}{\Delta_{mh} WTP_{mh}}$, $x_{hm} = \frac{|\partial \pi_m^M / \partial p_{hm}|}{\Delta_{mh} WTP_{mh}}$, $\alpha = -l_0$, $\beta_h = \frac{\tau_h}{1-\tau_h}$. To see that α, β_h, and $\tilde{k}_m$ are identified from the above regression, it is sufficient to look at how the variation in the covariates $(D^{\Delta}, v, D^{\partial}, x)$ translates into variation in the dependent (y). In particular, it is easy to verify that $\frac{\partial y_{mh} / \partial D_{hmm'}^{\Delta}}{\partial y_{mh} / \partial v_{hm}^H} = \tilde{k}_{m'}$ and thus $\partial y_{mh} / \partial v_{mh}^H$ identifies β and $\partial y_{mh} / \partial D_{hmm'}^{\partial}$ identifies α. L is identified from the distribution of the implied residual $\tilde{l}_{hm}$. □

Finally, Proposition 1 follows as a direct consequence of Lemma 1 applied to Lemma 3. The results derived in this section imply that the identification of premium-setting and pricing parameters in our model can be cast into the framework of identification in linear and quadratic regressions. Lemma 2 shows that identifiability is not determined by the simplifying assumptions we imposed on the model, namely the iid disagreement penalty and the block structure on hospital costs. Rather, the identification of the Nash-in-Nash bargaining parameters follows from the optimality of the bargaining solution augmented with observation of hospital total costs.

C Methodological Appendix

This appendix provides additional details about the estimation and counterfactual simulation methodology.

C.1 Demand Instruments

To estimate consumer preferences for hospitals and insurance plans, we must address the potential endogeneity of firms' choices to unobserved preferences. As the main text notes, we rely on two sets of instruments for this purpose, which we describe next.

C.1.1 Hospital Demand Instruments. The main endogeneity problem in identifying preferences over hospitals is that prices might be negotiated with knowledge of unobserved preferences for care. We instrument prices using supplementary data on claims paid by *closed* insurers—large employers who have formed insurance companies exclusively to cover their employees and who do not sell insurance on the open private market we study. For each medical event of a closed-insurer enrollee, we regress the total bill amount (i.e., the total amount paid to the hospital) on diagnosis-year, hospital fixed effects, age, and gender fixed effects. We use the estimated parameters to predict total closed-insurer payments for each option in our hospital demand panel. We multiply the predicted price by the enrollee's coinsurance rate to account for differential coverage. We use this predicted out-of-pocket price as our instrument.

Table A.7-A shows the first stage of the instrument on private enrollees' out-of-pocket hospital prices within the hospital demand panel. To match the estimating equation, the regression controls for distance and year-diagnosis-insurer-hospital fixed effects. As shown, the instrument has a strong and significant positive relationship with the endogenous price variable. While Table 2 in the main text shows the effect with the instruments included, Table A.7-B shows the impact of ignoring them. As usual, failing to instrument for endogenous prices results in a lower predicted price elasticity.

C.1.2 Plan Demand Instruments. The main endogenous variable in our plan demand estimation corresponds to plan premiums. As explained in the main text, we instrument this variable using the public hospital system's list prices to compute each household's expected spending. Adjusting spending by each plan's coverage, we construct a predicted cost metric for each insurer, which we call the actuarially fair premium. We instrument each plan's premium with the average of its rivals' actuarially fair premiums. Table A.7-A shows the first stage regression of the instrument, while Table A.7-B shows the effect of ignoring the instruments on estimated preferences. As with hospital demand, ignoring the endogeneity of premiums results in a lower predicted elasticity.

C.2 Price and Premium Setting Parameter Estimator

We estimate our model's price and premium setting parameters using an iterative two-step procedure. The outer loop takes a guess of hospital costs and applies a linear inversion to recover hospital plan administrative cost (η_{jt}) and VI weights (θ_{hmt}) from the optimality conditions associated with plan premiums and VI prices. Conditional on plan administrative cost and VI weights, the inner loop solves a constrained maximum likelihood problem to recover hospital costs, bargaining weights, and the distribution of the disagreement penalty multiplier. The estimator can be described as follows:

$$
\begin{aligned}
\max_{\boldsymbol{\tau}\in[0,1]^{|H|},\boldsymbol{k}^H,\tilde{\boldsymbol{k}}^H,\mu_l,\sigma_l} \sum_{h,m,t\in\mathcal{B}} & \ln \mathcal{L}(p^*_{hmt}|\boldsymbol{\beta}^\lambda, \mu^L, \sigma^L, \boldsymbol{\beta}^c) \\
\text{s.t.} \quad & k^H_h + \tilde{k}^H_{mt} \geq 0 \quad \forall h,m,t && (C1) \\
& \Delta_{mh}\tilde{\pi}_{ht} \geq 0 \quad \forall h,m,t && (C2_H) \\
& \Delta_{mh}\tilde{\pi}_{mt} + l_{hmt}\Delta_{mh} WTP_{mht} \geq 0 \quad \forall h,m,t && (C2_M) \\
& \frac{\partial \Delta_{mh}\tilde{\pi}_{ht}}{\partial p_{hmt}} \geq 0 \quad \forall h,m,t && (C3) \\
& k^o_{ht} = \sum_{m'} D^H_{h|m}(k^H_h + \tilde{k}^H_{mt}) \quad \forall h,t && (MATCH) \\
& \nabla_{[\phi_m, \boldsymbol{p}^{VI}_m]}\tilde{\pi}_{mt} = 0 \quad \forall \text{ VI-}m && (INV)
\end{aligned}
$$

where the likelihood of hospital prices is evaluated over all hospital-insurer price negotiations. The first four constraints match the requirements of our identification results: $C1$ imposes that hospital costs are positive, $C2_H$ and $C2_M$ that negotiations are individually rational, and $C3$ that hospital gains from trade are increasing in price at the observed prices. These constraints impose bounds on hospital costs, as neither bargaining parameters nor the penalty distribution enter them. The $MATCH$ condition incorporates our additional hospital cost data, where $D^H_{h|m}$ denotes the demand for hospital h from insurer m, weighted by resource intensity ω_{id}. Finally, the INV constraint captures the linear inversion procedure for recovering plan administrative costs and VI weights from VI firms' optimality conditions.

Following the Nash bargaining model formulated in Section 4, the likelihood of observing p^*_{hmt} for the negotiation between hospital h and insurer m in year t is the likelihood with which observed price satisfy the first-order optimality condition:

$$\begin{aligned}\mathcal{L}(p^*_{hmt}|\boldsymbol{\beta}^{\lambda}, \mu^L, \sigma^L, \boldsymbol{\beta}^c) &= \mathbb{P}\left(\lambda_h \frac{\partial \tilde{\pi}_{mt}/\partial p_{hmt}}{\Delta_{mh}\tilde{\pi}_{mt} + l_{mht}\Delta_{mh}WTP_{mt}} + \frac{\partial \tilde{\pi}_{ht}/\partial p_{hmt}}{\Delta_{mh}\tilde{\pi}_{ht}} = 0\right) \\ &= \mathbb{P}\left(l_{hmt} = -\frac{\lambda_h}{\Delta_{mh}WTP_{mt}} \frac{\partial \tilde{\pi}_{mt}}{\partial p_{hmt}} \frac{\Delta_{mh}\tilde{\pi}_{ht}}{\partial \tilde{\pi}_{ht}/\partial p_{hmt}} - \frac{\Delta_{mh}\tilde{\pi}_{mt}}{\Delta_{mh}WTP_{mt}}\right)\end{aligned}$$

where $\lambda_h = \frac{\tau_h}{1-\tau_h}$. The analytic form for this expression is easily derived from the assumption that l is normally distributed.

A relevant advantage of splitting the estimation into two nested loops is that the resulting constraints on the maximum likelihood problem are linear in parameters. Moreover, only cost parameters ($\boldsymbol{k}^H, \tilde{\boldsymbol{k}}^H$) feature in them, appearing only in hospital profits, which are linear in costs.[4] Hence, all constraints are linear, and the feasible set of cost solutions is a regular convex set. It is easy to see that because $\tilde{\pi}_{mt}$ is linear in administrative costs (η_{jt}) and VI weights (θ_{hmt}), these are uniquely determined from the optimality condition of VI prices and premiums, conditional on a guess of hospital costs (see the proof of Proposition 1). Hence, this estimator is simple and efficient when implemented in a two-step procedure, and the feasible region of parameters is easy to explore. None of the first four constraints is binding, and the other two can be absorbed within the likelihood objective. Therefore, the asymptotic properties of the estimator are well described by those of standard maximum likelihood. Figure A.7 shows the convergence rate of the solution in our estimation. It shows that within two iterations, we attain a stable solution for supply-side parameters, which takes only a few minutes to solve.

C.3 Plan Design Cost Estimator

The cost of designing plan coverage is estimated in two steps. First, we estimate the regulatory cost component, which is a function of the continuous coverage level of each plan. We specify the cost as $K^r_m(\boldsymbol{c}_{jt}) = \exp(c^K(\boldsymbol{c}_{jt})) + \underline{c}_{jt}\underline{\zeta}_{jt} + \bar{c}_{jt}\bar{\zeta}_{jt}$, where $c^K(\cdot)$ is a flexible polynomial of coverage and $(\bar{\zeta}, \underline{\zeta})$ are mean zero iid normal shocks of unknown variance. We estimate this component of the model by using the optimality condition associated

[4]In practice, we do not impose the MATCH constraint directly but add it as a large quadratic cost to the solver. This is to avoid issues with empty feasible sets on single hospitals for some intermediate solver iterations. On convergence, the solution satisfies the constraint.

with coverage levels of each plan $(\underline{c}_{jt}, \bar{c}_{jt})$:

$$
\begin{aligned}
\frac{\partial \tilde{\pi}_{mt}(\phi^*(\boldsymbol{c}), \boldsymbol{p}^*(\boldsymbol{c}), \boldsymbol{c})}{\partial \underline{c}_{jt}} - M_{jt}\frac{\partial K^r_m(\boldsymbol{c}_j)}{\partial \underline{c}_{jt}} &= 0 \\
\frac{\partial \tilde{\pi}_{mt}(\phi^*(\boldsymbol{c}), \boldsymbol{p}^*(\boldsymbol{c}), \boldsymbol{c})}{\partial \bar{c}_{jt}} - M_{jt}\frac{\partial K^r_m(\boldsymbol{c}_j)}{\partial \bar{c}_{jt}} &= 0
\end{aligned}
$$

Given the pricing subgame, the primary challenge in this estimation procedure is computing the derivative of equilibrium profits. With the average insurer offering 81 plans with two tiers each over four years of data, this necessitates calculating 648 derivatives of the subgame equilibrium for each insurer. This is made feasible by recent advances in GPU-accelerated linear algebra and automatic differentiation(Bradbury *et al.*, 2018).

Having estimated the continuous regulatory cost, we turn to estimating the fixed cost of tiering. We define $\tilde{V}_{mt} = \tilde{\pi}_{mt} - \sum_{j\in J_{mt}} M_{jt}K^r_m(\boldsymbol{c}_j)$, and $\Delta_{jht}\tilde{V}_{mt}$ denotes the change in $\tilde{V}_{mt}$ when the tiering position of hospital h in plan j is inverted. We implement equations (14) and (15) following the approach of Canay *et al.* (2023). We form a grid of nearly two thousand potential tiering costs, ranging from -3 to 3 million dollars per hundred thousand enrollees. We use the test of Chernozhukov *et al.* (2019) to admit points into the identified set. To evaluate the associated confidence interval, we identify the minimum and maximum points that satisfy the inequality and refine the cost grid to pinpoint the exact values where the inequality binds. We apply the bootstrap approach described in Chernozhukov *et al.* (2019), taking 300 random samples of our data with replacement.

Evaluating the impact of tiering decisions on equilibrium conditions is computationally intensive. With 11 hospitals to tier, we need to evaluate approximately 17,820 different subgame equilibria to form the estimator. Again, this is facilitated by optimizing the subgame equilibrium fixed-point solver and employing GPU acceleration.

C.4 Solving Counterfactual Equilibria

Simulating counterfactual equilibria involves two primary challenges: solving the combinatorial plan design problem and identifying a set of equilibrium strategies for all firms.

We tackle the first problem by relying on the result of Proposition 1 and XLA/JAX. We define a grid of tiering penalties $\lambda \in \{0, 1, 10, 100\}$ and a penalty function $G(x) = x^2$. Given that our costs are measured in millions, these choices imply a maximum penalty of 68.75 million per plan for violating the tiering constraint. Fixing rival designs, we solve each firm's design problem in parallel, determining the solution $\tilde{c}^*(\lambda)$ along the grid. The initial condition at $\lambda = 0$ and the first iteration of the solver are always set to the status quo, with subsequent iterations updating from the previous solution. If at any $\lambda < 100$ the single-firm optimum has a tiering violation ($\sum_{j \in J_{mt}} \sum_{h \in H} G(w_{jh}(1 - w_{jh}))$) below 10^{-4}, we accept it as optimal for the best response iteration.

To illustrate the performance of this approach, we compare it to solving the problem by brute force using grid search. We select two firms at random and choose a single plan for each. We consider a grid of potential coverage levels, taking values in $[0.65, 0.7, 0.75, 0.8, 0.85]$, and evaluate the profitability of every potential tier design while keeping the design of all other plans in the market fixed. In total, each firm evaluates 20,480 configurations. As the effect of a single plan on equilibrium prices and premiums is small, we shut down this channel for this exploration, which vastly reduces the computational cost of grid search. To illustrate its stability, we run our regular convergent solver from 30 random starting points. Figures A.8a and A.8b show the results. On average, our algorithm attains between 95 and 101 percent of the maximum value of the grid search.[5] For insurer 1, the objective curve is slightly smoother, likely contributing to improved performance. Our worst solution attains 98.7 percent of the maximum grid value. For insurer 2, the objective is slightly more jagged, likely contributing to a reduced worst-case performance of 91 percent.[6] Figure A.8c shows the difference in computing times for both approaches. While grid search is about as fast as our approach for a single plan, it is about two orders of magnitude slower for two plans and nearly seven for three plans. When evaluating these results, it is worth remembering that this class of problems is NP-hard; insurers are likely relying on commercial mixed-integer packages to solve these design problems, which are not guaranteed to obtain the optimum. As noted by Murray and Ng (2010), this approach is often as good or better than commercial solvers.

The second challenge involves finding the intersection of firms' best responses. We employ a Gauss-Seidel approach: starting with status quo designs, we identify each firm's best response to the current state and update the state accordingly. We iterate until the change in plan coverages is less than 10^{-5} in the Euclidean norm.

To illustrate the performance of the approach, we return to the two single-product firms described above. We solve the equilibrium between the two firms by grid search and through our regularized approach, starting from 30 random points as before. Denoting

[5] Our algorithm can improve upon the result of the grid because it is not constrained to the same coarse coverage options.

[6] The figures also show that the grid optimum is isolated from other designs. This suggests that firms' best responses might have a unique solution.

the equilibrium profit of firm j under the brute-force grid search as π_j^B, and under the regularized approach as π_j^R, we define the relative fit of our approach as $\varepsilon_j = (\pi_j^R - \pi_j^B)/\pi_j^B$. We describe equilibrium fit based on $\varepsilon^* = (\sum_{j=1}^2 \varepsilon_j^2)^{1/2}$. Figure A.8d shows the distribution of ε_j and ε^* across the different starting points. The relative difference in insurer profits (ε_j) is minuscule, with an absolute value below 0.3 percent for insurer 1 and below 0.1 percent for insurer 2. The equilibrium fit (ε^*) displays an interquartile range below 0.25 percent. Most of the mismatch between the approaches is due to improved best responses found by the regularized approach rather than from meaningful differences in plan design.

Finally, below is an illustration of our full process for one counterfactual equilibrium exercise, which converged in four iterations of best response intersections:

```
Outer [Relaxed][0] Time: 839.7 seconds
 BR: [764.778, 501.524, 210.011, 874.270, 4.082]
 NI: [506.341, 294.806, 89.630, 632.665, −264.353]
 Penalties: [5.329∗10^−14, 3.908∗10^−14, 5.684∗10^−14, 1.279∗10^−13, 1.778∗10^−8]
 Delta prices 10.0
 Delta premiums 2.032∗10^−1

Outer [Relaxed][1] Time: 793.041 seconds
 BR: [748.513, 491.870, 210.517, 862.170, 1.267]
 NI: [748.512, 489.813, 210.513, 862.134, 0.017]
 Penalties: [5.329∗10^−14, 5.329∗10^−14, 6.040∗10^−14, 1.386∗10^−13, 4.718∗10^−4]
 Δ prices 1.815∗10^−4
 Δ premiums 3.250∗10^−4

Outer [Relaxed][2] Time: 670.903 seconds
 BR: [748.501, 492.114, 210.460, 862.138, 1.242]
 NI: [748.501, 492.113, 210.453, 862.138, 1.241]
 Penalties: [4.619∗10^−14, 5.329∗10^−14, 5.329∗10^−14, 1.386∗10^−13, 4.718∗10^−04]
 Δ prices 2.295∗10^−6
 Δ premiums 1.270∗10^−7

Outer [Strict][3] Time: 352.612 seconds
 BR: [747.797, 491.862, 208.375, 853.319, 1.176]
 NI: [747.795, 491.861, 208.375, 853.319, 1.177]
 Penalties: [8.171∗10^−14, 4.263∗10^−14, 1.741∗10^−13, 5.933∗10^−13, 5.149∗10^−13]
```

These iterations display several key properties of our approach and the problem. First, they illustrate how convergence occurs through various margins. At each iteration, we present the best response profits (BR) versus the profits implied by all the best responses (NI, for Nikaido-Isoda) for each firm. Formally, for an initial condition vector $(\boldsymbol{c}_1, \ldots, \boldsymbol{c}_5)$ and best responses $(\boldsymbol{c}_1^*(\boldsymbol{c}_{-1}), \ldots, \boldsymbol{c}_5^*(\boldsymbol{c}_{-5}))$, BR profits correspond to $\tilde{V}_m(\boldsymbol{c}_m^*(\boldsymbol{c}), \boldsymbol{c}_{-m})$ and NI profits are $\tilde{V}_m(\boldsymbol{c}_m^*(\boldsymbol{c}), \boldsymbol{c}_{-m}^*(\boldsymbol{c}))$, where $\tilde{V}_m$ is the total profit of insurer m net of underwriting cost. A vector of coverages $\boldsymbol{c}^*$ is a Nash equilibrium if and only if BR and NI profits match. The iterations reveal the rate at which this convergence is attained.

Additionally, the iterations show how insurer tiering penalties converge. They indicate that throughout our iterations, all but one firm satisfy the tiering constraint almost exactly. Our algorithm identifies that at the third iteration (marked [Relaxed][2]), the coverage vector is stable, but tiering constraints are not all satisfied. This triggers a new cycle (marked [Strict]) in which any violating coverage vector is adjusted to its nearest tiered configuration, serving as a starting point, and the convergence tolerance on firms' best response solver is substantially increased. As shown, the subsequent iteration results in an equilibrium with insignificant tiering violations. Intuitively, this improvement occurs because the profit objective of the last insurer is nearly flat on its coverage decision, necessitating a shift in its initial condition and convergence criteria to attain the optimum.

Finally, the iterations illustrate how changes to the coverage structure affect the subgame price equilibrium at the end of every iteration. Changes are shown in the Euclidean norm, indicating that as coverage converges, equilibrium prices and premiums also converge. Importantly, this reflects the local stability of the equilibrium: small changes in coverage do not induce large changes in prices or premiums.

It is worth noting that the counterfactual presented above was chosen for its few iterations to demonstrate full convergence. Our main VI ban counterfactual takes considerably longer to solve (approximately 8 hours) but shares the same properties.

References

Alvear-Vega, S. and Acuña, M. (2022). Determinantes Sociales que Influyen en el Acceso en Chile al Plan GES, según CASEN 2017. *Revista Médica de Chile*, pp. 70 – 77.

Amemiya, T. (1983). Chapter 6 Non-linear Regression Models. In *Handbook of Econometrics*, *Handbook of Econometrics*, vol. 1, Elsevier, pp. 333–389.

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

Bradbury, J., Frostig, R., Hawkins, P., Johnson, M. J., Leary, C., Maclaurin, D., Necula, G., Paszke, A., VanderPlas, J., Wanderman-Milne, S. and Zhang, Q. (2018). JAX: Composable Transformations of Python+NumPy Programs.

Canay, I. A., Illanes, G. and Velez, A. (2023). A Users Guide for Inference in Models Defined by Moment Inequalities. *Journal of Econometrics*, p. 105558.

CASEN (2015). *Encuesta de Caracterización Socio-Econmíca Nacional*. Tech. rep., Ministerio de Desarrollo Social.

Chernozhukov, V., Chetverikov, D. and Kato, K. (2019). Inference on Causal and Structural Parameters using Many Moment Inequalities. *The Review of Economic Studies*, **86** (5 (310)), pp. 1867–1900.

Chile Atiende (2024). Reclamo por alza en planes de Isapres. `https://shorturl.at/e2JSF`, accessed: 2024-07-25.

CIPER (2013). Reajuste de precios de Isapres: una historia de abuso, imposición, desigualdad y lucro. `https://shorturl.at/JySsr`, accessed: 2024-07-25.

Cooper, Z., Craig, S., Gaynor, M. and Reenen, J. V. (2018). The Price Ain't Right? Hospital Prices and Health Spending on the Privately Insured. *Quarterly Journal of Economics*.

Cutler, D. M., McClellan, M. and Newhouse, J. P. (2000). How Does Managed Care Do It? *The RAND Journal of Economics*, **31** (3), 526–548.

Dias, M. (2022). Selection in a Health Insurance Market with a Public Option: Evidence from Chile, manuscript.

Duarte, F. (2011). Switching Behavior in a Health System with Public Option, manuscript.

Galetovic, A. and Sanhueza, R. (2013). Un Análisis Económico de la Integración Vertical entre Isapres y Prestadores, manuscript.

Gowrisankaran, G., Nevo, A. and Town, R. J. (2015). Mergers When Prices Are Negotiated: Evidence from the Hospital Industry. *American Economic Review*, **105** (1), 172–203.

Ho, K. and Lee, R. (2017). Insurer Competition in Health Care Markets. *Econometrica*, **85** (2), 379–417.

Johnson, E. M. and Rehavi, M. M. (2016). Physicians Treating Physicians: Information and Incentives in Childbirth. *American Economic Journal: Economic Policy*, **8** (1), 115–141.

Kihlstrom, R., Mas-Colell, A. and Sonnenschein, H. (1976). The Demand Theory of the Weak Axiom of Revealed Preference. *Econometrica*, **44** (5), 971–978.

La Tercera (2020). Banmédica recibirá a nuevos afiliados con preexistencias y le pone presin a la industria y al regulador. `https://shorturl.at/MWNSU`, accessed: 2024-07-25.

Morrow, W. R. and Skerlos, S. J. (2011). Fixed-point Approaches to Computing Bertrand-Nash Equilibrium Prices under Mixed-logit Demand. *Operations Research*, **59** (2), 328–345.

Murray, W. and Ng, K.-M. (2010). An algorithm for nonlinear optimization problems with binary variables. *Computational Optimization and Applications*, **47** (2), 257–288.

Pardo, C. and Schott, W. (2012). Public versus Private: Evidence on Health Insurance Selection. *International Journal of Health Care Finance and Economics*, **12**, 39–61.

Tebaldi, P. (2024). Estimating Equilibrium in Health Insurance Exchanges: Price Competition and Subsidy Design under the ACA. *Review of Economic Studies*.

Zhang, J., O'Donoghue, B. and Boyd, S. (2020). Globally Convergent Type-I Anderson Acceleration for Nonsmooth Fixed-point Iterations. *SIAM Journal on Optimization*, **30** (4), 3170–3197.

Figure A.1: Location of hospitals and enrollees in the market

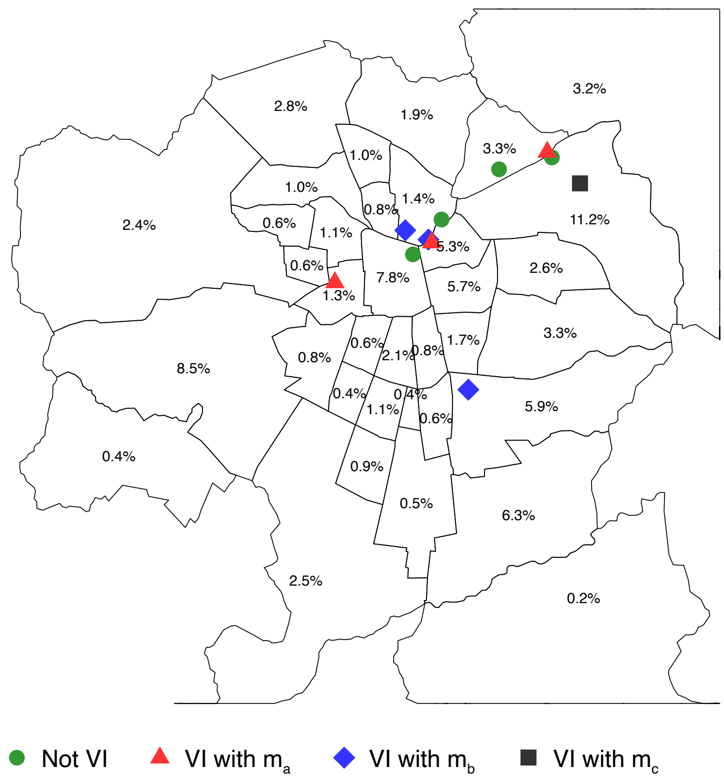

Notes: This figure shows the location of hospitals in the market. The map covers most of urban Santiago, our market of interest. Green circles indicate independent hospitals, red triangles indicate hospitals vertically integrated with insurer m_a, blue diamonds indicate those vertically integrated with insurer m_b, and black squares indicate those integrated with m_c. Finally, numbers indicate the share of enrollees from each county on the map. Some counties located further away are omitted from the plot for convenience but are included in the analysis. Relative to the full population distribution, there is a noticeable enrollment concentration in the city's wealthier areas (from the center of the figure to the northeast). However, there is clear dispersion in the location of enrollees relative to hospitals.

Figure A.2: Vertical integration, hospital choices, and expenditure (movers subsample)

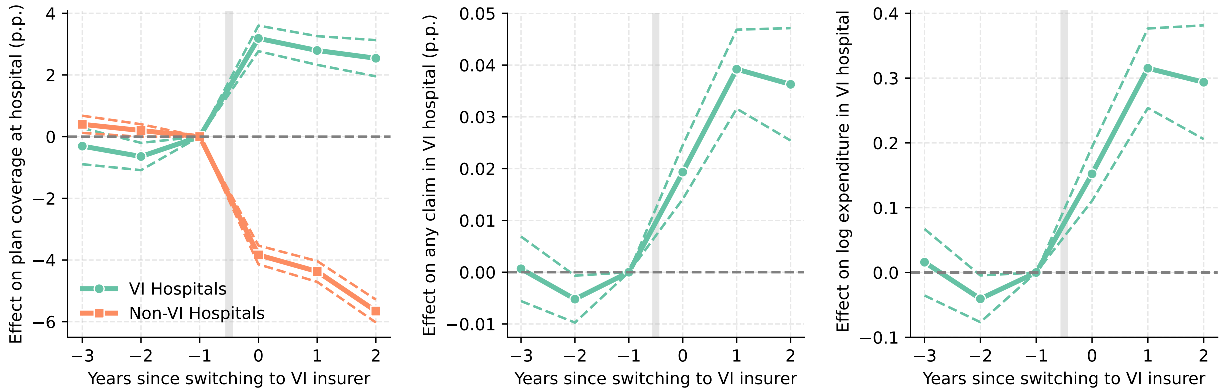

(a) Plan coverage **(b)** Any claims in VI hospital **(c)** Expenditure in VI hospital

Notes: This figure displays event study estimates from equation (4) in the main text for a subsample of enrollees that move across neighborhoods. This subsample includes 18 percent of the enrollees in the main analysis. The coefficient for the year before the patient switches is set to zero. Green dots and orange squares are estimates of β_τ and γ_τ in equation (4), respectively. Dashed lines indicate 95% confidence intervals. The dependent variable in Figure A.2c is $\log(1 + y)$ to accommodate zeros, but the results are similar when using expenditure in levels.

Figure A.3: Estimates of hospital specialization

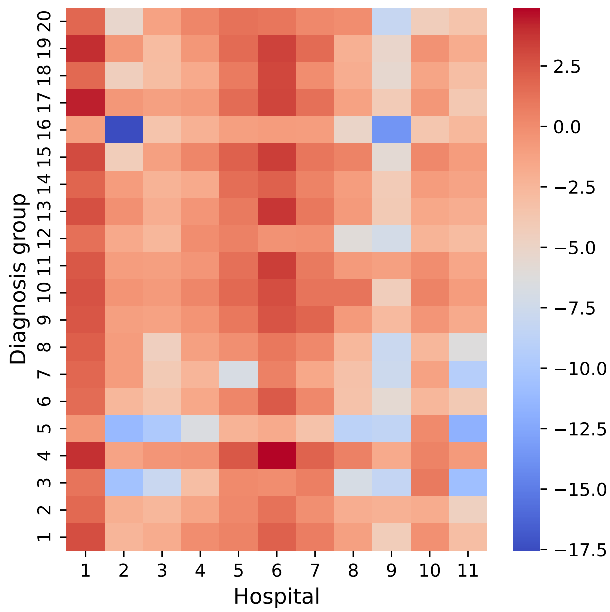

Notes: This figure presents estimates of consumers' average preference for hospital-diagnosis pairs. These correspond to χ^{H}_{hdt} in equation (5), averaged across years and normalized with respect to the outside option. Diagnosis groups correspond to ICD-10 diagnosis chapters, and hospitals match our anonymized identifiers. Hospitals 1 and 6 correspond to the star hospitals in our data. Hospitals 2, 3, and 8 are integrated with insurer m_a, and hospitals 4, 7, and 11 are integrated with insurer m_b. Hospital 10 is integrated with insurer m_c in the first year of the data.

Figure A.4: Counterfactual preferential structure, demand, and prices

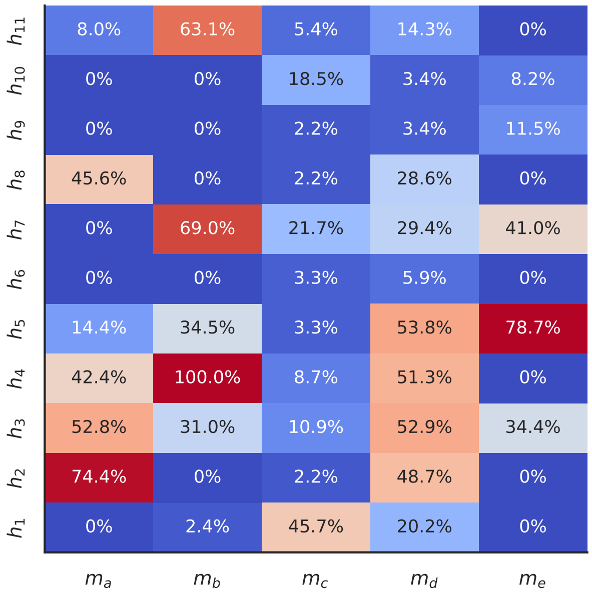

(a) Baseline Preferentials

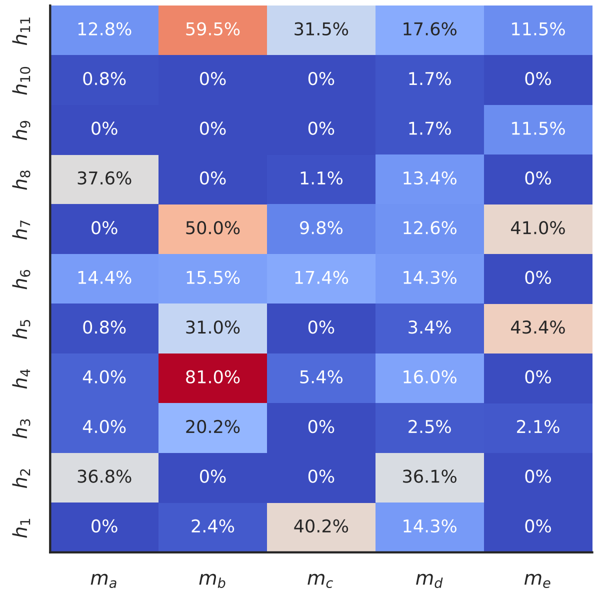

(b) Counterfactual Preferentials

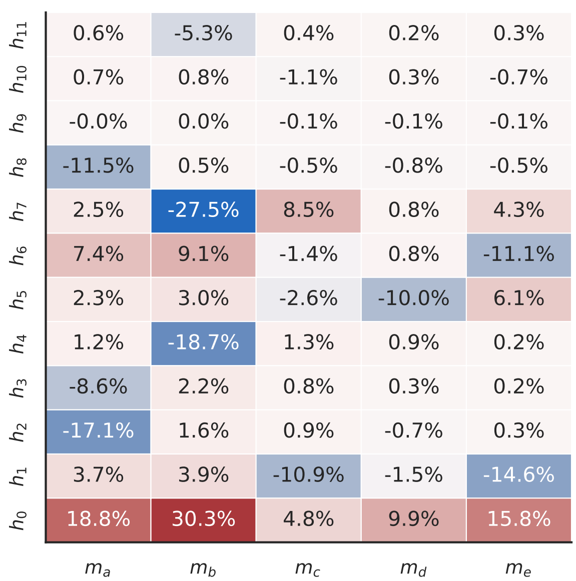

(c) Counterfactual demand change

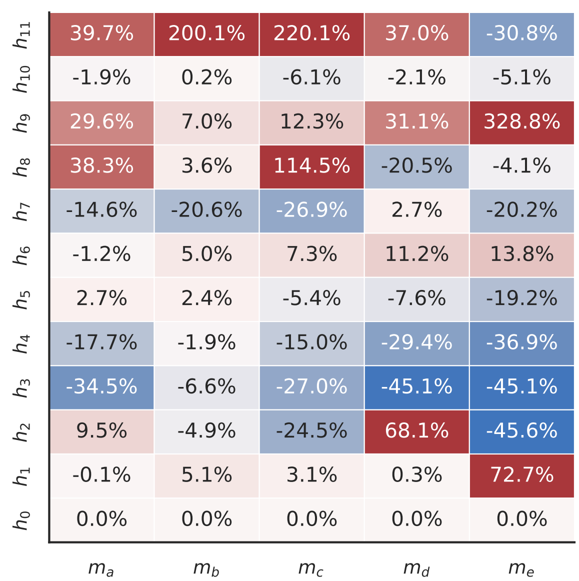

(d) Counterfactual price change

Notes: Figures (a) and (b) illustrate the number of plans of each insurer (rows) that have each hospital (columns) as preferential. Figure (c) shows the change in hospital demand within each insurer. Changes add up to zero within each column. Figure (d) shows the percent change in negotiated prices between each hospital and insurer. Hospitals 1 and 6 correspond to the highest-quality and highest-priced non-VI hospitals in our data. Hospitals 2, 3, and 8 are integrated with insurer m_a, and hospitals 4, 7, and 11 are integrated with insurer m_b.

Figure A.5: Effects of banning VI under cost efficiencies and quality gains

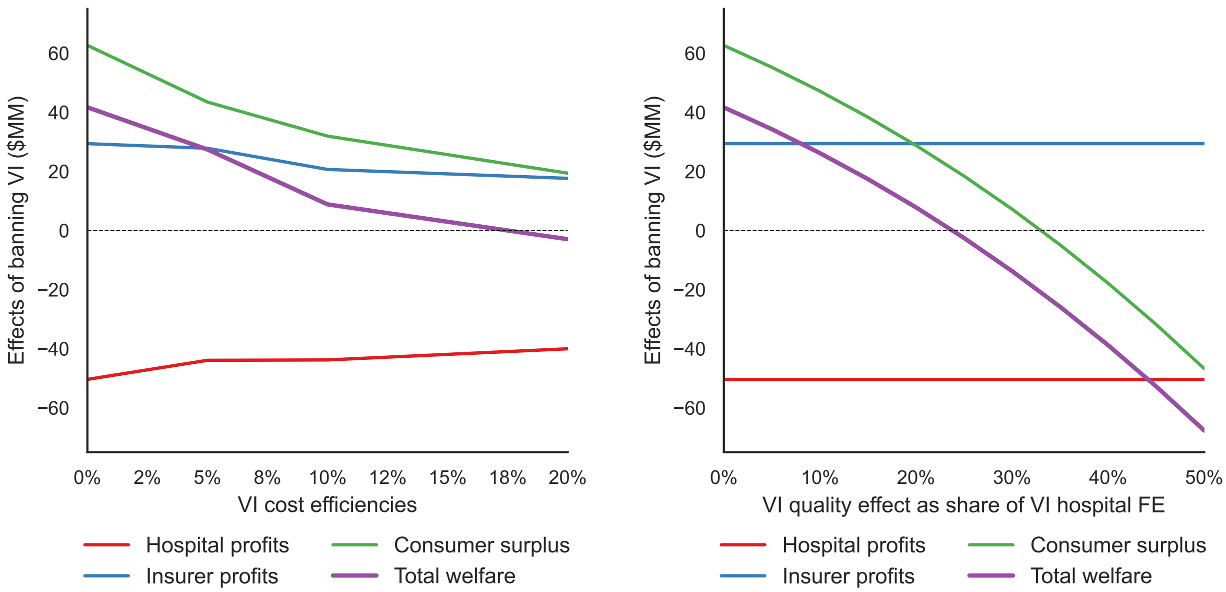

(a) Cost efficiencies

(b) Quality effect

Notes: Figure (a) shows the effect of banning VI in the presence of cost efficiencies. The x-axis represents the percent increase in hospital costs for formerly VI hospitals when serving former partners under the VI ban. Figure (b) shows the effect of banning VI under quality effects. This is computed by assuming that different shares of the VI marketing effect represent true quality differentials. This only affects the baseline surplus value as both marketing and quality changes disappear in the counterfactual, leading to the same equilibrium.

Figure A.6: Effects of banning VI under alternative elasticities

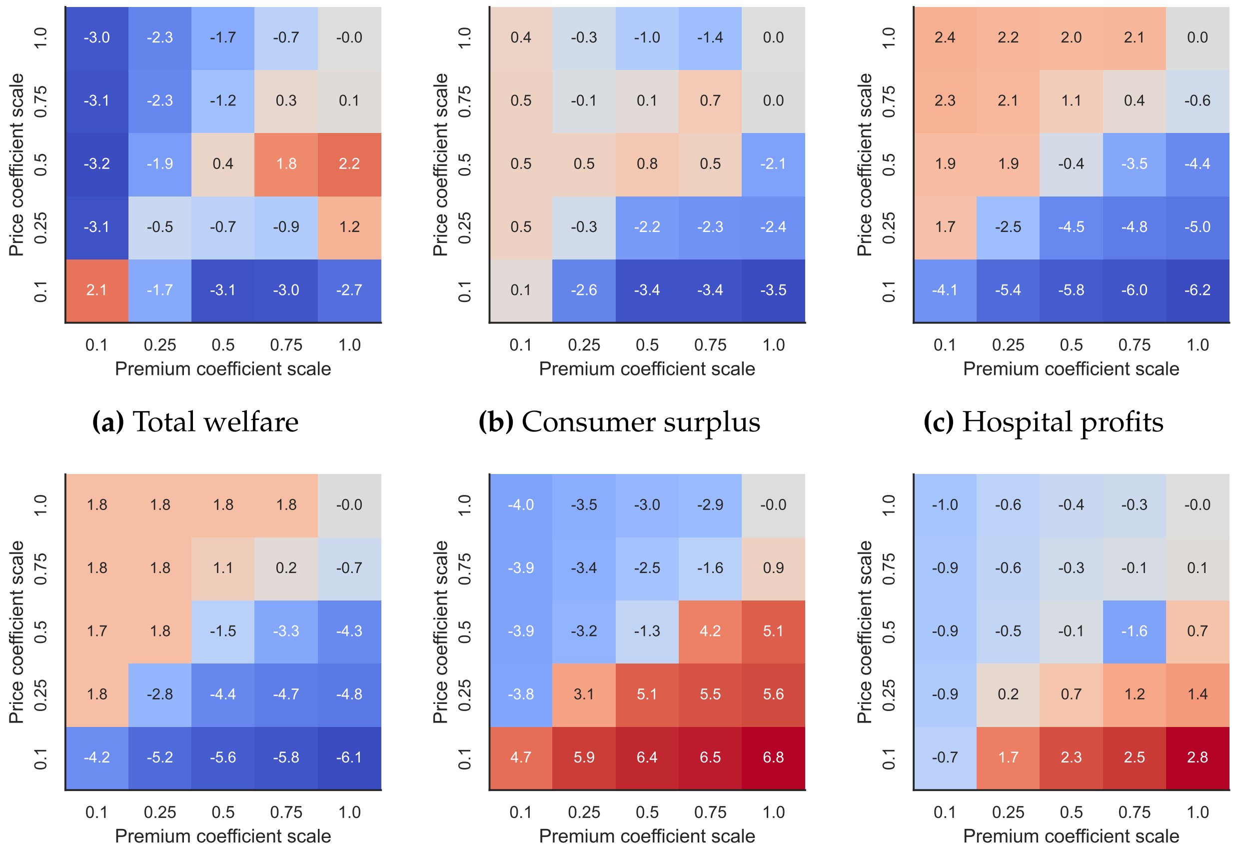

(a) Total welfare **(b)** Consumer surplus **(c)** Hospital profits

(d) VI hospital profits **(e)** Insurer profits **(f)** VI insurer profits

Notes: These figures show the change in welfare from banning VI under alternative price and premium elasticity in log scale relative to the results of the main analysis. Formally, letting (θ_x, θ_y) be the multipliers on the horizontal and vertical axes, and $\Delta TW(\alpha^H, \alpha^M)$ denote the change in total welfare produced from banning VI when consumers' price coefficient is α^H and premium coefficient is α^M, each cell consists of computing $v(\theta_x, \theta_y) = \Delta TW(\theta_x \alpha^H, \theta_y \alpha^M)/|\Delta TW(\alpha^H, \alpha^M)|$ and reporting $sign(v) log(|v|)$. Note that each cell requires computing two full (medium run) counterfactual equilibria, one for the simulated baseline and one for the counterfactual under a VI ban. In both cases, the starting point is the observed status quo to retain our equilibrium notion and consistent treatment of counterfactual simulations. This, however, implies that the VI ban counterfactual in each cell might not be the closest local equilibrium to its baseline comparison. Therefore, we adjust all values relative to our main estimates such that the welfare change in the top right corner equals the simulated baseline.

Figure A.7: Convergence rate of supply-side MLE parameters

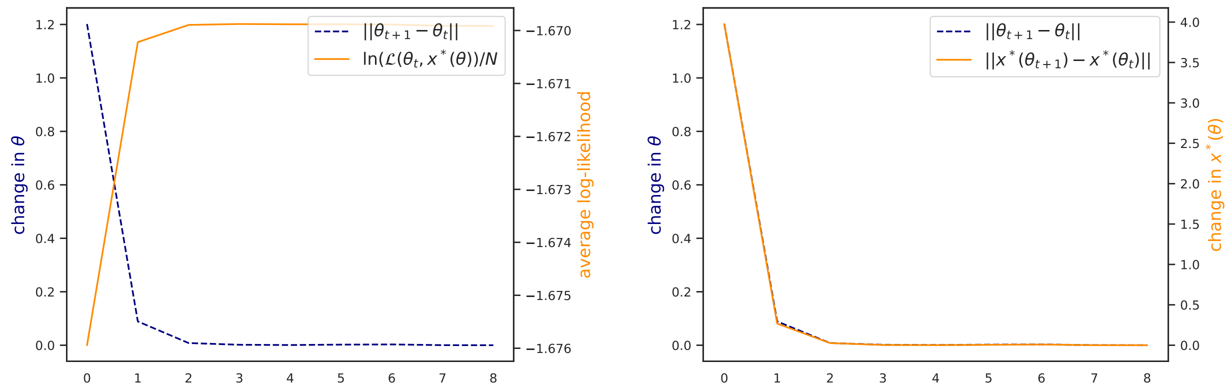

(a) Likelihood **(b)** Parameters

Notes: These figures show the convergence rate of model estimates governing price and premium setting. The x-axis shows the iterations of the nested algorithm. In the figures, θ_t stands for the vector of VI weights in the t-th iteration of the solver, and $x^*(\theta_t)$ is the vector of all MLE parameters conditional in the t-th iteration. Figure (a) shows the rate at which the log-likelihood and the VI weights converge. Figure (b) shows that the MLE parameters converge at a nearly identical rate as the VI weights.

Figure A.8: Performance of plan design solver

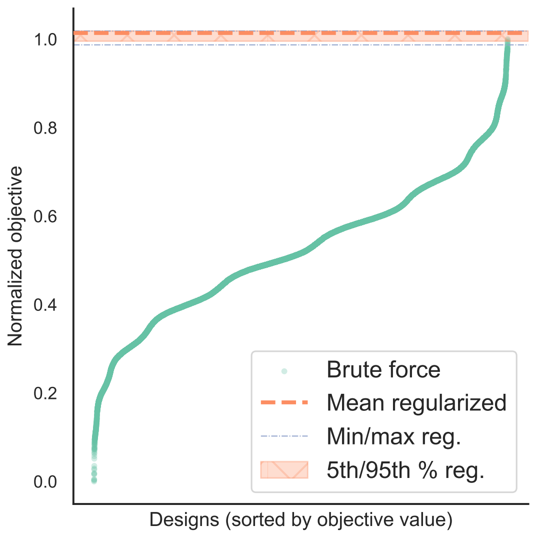

(a) Insurer 1 best response

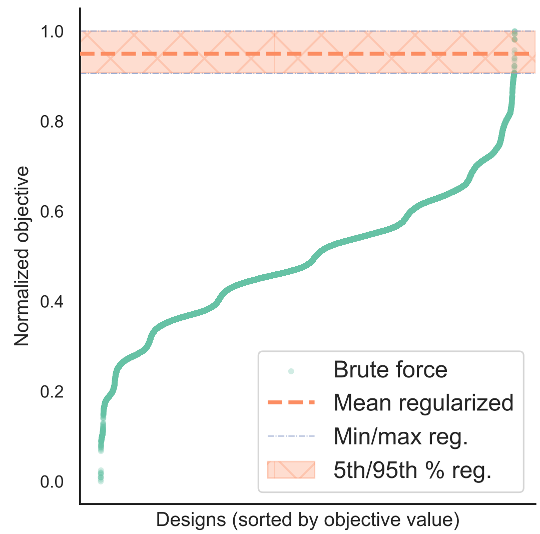

(b) Insurer 2 best response

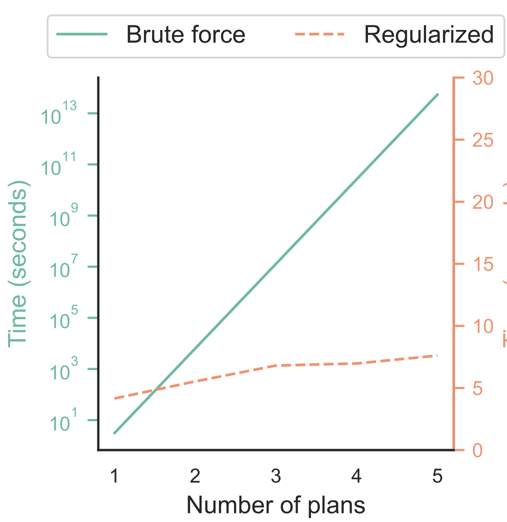

(c) Computing time

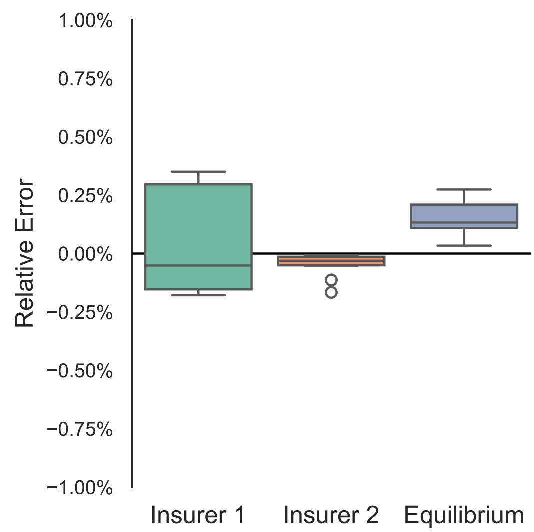

(d) Equilibrium fit

Notes: These figures show how the solution of our regularized convergent approach compares against grid-search. Panels (a) and (b) present the best responses of two firms. Green circles show the profit objective value at each of the 20,480 different designs evaluated. We run our regularized solver from 30 random starting points and report the range of results obtained. For ease of comparison, the vertical axis represents the normalized objective relative to the minimum and maximum grid-search value. Panel (c) shows the execution time associated with solving a firm's best response problem for different numbers of plans. The straight green line shows the time in seconds for the brute force grid search approach, inferred from the time it takes to compute a single objective evaluation, scaled by the number of evaluations associated with each problem. The dashed orange line is the mean compute time associated with solving the associated problem for Insurer 1 from five random starting points. All computations were done sequentially on an Nvidia A100 GPU. Panel (d) shows the fit between the equilibrium result derived from our approach relative to that found by grid search. In the first two columns, we plot the distribution of $\varepsilon_j = (\pi_j^R - \pi_j^B)/\pi_j^B$ where π_j^B is the brute-force equilibrium profit of firm j, and π_j^R is the regularized equilibrium profit of the same firm. The third column shows the distribution of $(\sum_{j=1}^{2} \varepsilon_j^2)^{1/2}$ across the 30 starting points.

Table A.1: Descriptive statistics

	N	Mean	SD	p10	p50	p90
A - Policyholder attributes						
Age	3,946,900	39.93	10.37	27.00	38.00	56.00
Single female	3,946,900	0.24	0.42	0.00	0.00	1.00
Single male	3,946,900	0.34	0.47	0.00	0.00	1.00
Number of dependents	3,946,900	0.81	1.17	0.00	0.00	3.00
Income	3,946,900	1.63	1.02	0.00	1.58	2.93
Paid premium	3,946,900	0.17	0.11	0.07	0.15	0.31
B - Plan network structure						
Tiered network plan	1,431	0.88	0.32	0.00	1.00	1.00
Tiered network plan | VI insurer	671	0.90	0.29	1.00	1.00	1.00
Tiered network plan | Non-VI insurer	760	0.86	0.35	0.00	1.00	1.00
Preferential coverage rate	1,431	77.33	12.79	60.00	80.00	90.64
Preferential coverage rate | VI insurer	671	74.37	11.96	58.32	76.70	89.11
Preferential coverage rate | Non-VI insurer	760	79.95	12.93	63.08	80.00	96.82
Base coverage rate	1,431	59.85	18.11	36.58	60.00	82.97
Base coverage rate | VI insurer	671	58.16	14.70	40.00	56.86	78.81
Base coverage rate | Non-VI insurer	760	61.34	20.55	32.69	64.82	84.74
C - Admission attributes at inside hospitals						
Full price	569,371	4.61	4.90	0.85	3.12	9.79
OOP share	569,371	0.24	0.24	0.00	0.12	0.63
Distance to hospital	569,371	7.47	6.27	0.00	5.96	15.13
Hospital is in preferential tier	569,371	0.64	0.48	0.00	1.00	1.00
Patient is enrollee of VI insurer | VI hospital	308,023	0.61	0.49	0.00	1.00	1.00
Patient is enrollee of VI insurer | Non-VI hospital	261,348	0.39	0.49	0.00	0.00	1.00
Full price at outside option	203,893	1.78	2.61	0.18	1.33	3.78

Notes: This table displays descriptive statistics for our estimating plans dataset. Panel A displays statistics across all policyholders in the sample. Panel B displays statistics for plan attributes across all plan-years in the sample offered in the spot market (i.e., excluding legacy plans held by consumers on guaranteed renewability). Panel C displays statistics across all admissions in the main hospitals included in our analysis unless otherwise noted. All monetary amounts are measured in thousands of U.S. dollars for December 30, 2014. Distance in miles.

Table A.2: Descriptive statistics for insurers and hospitals

	A - Insurer market shares and premiums					
	Market	Paid premium				
Insurer	share	Mean	SD	p10	p50	p90
m_a	19.87	0.14	0.08	0.06	0.12	0.24
m_b	29.62	0.18	0.11	0.07	0.15	0.32
m_c	17.26	0.20	0.13	0.08	0.17	0.37
m_d	22.59	0.18	0.11	0.07	0.15	0.31
m_e	10.66	0.16	0.10	0.07	0.14	0.29
	B - Hospital market shares and prices					
	Market	Full price				
Hospital	share	Mean	SD	p10	p50	p90
h_1	13.33	6.17	5.79	0.91	4.55	12.73
h_2	4.31	2.56	2.55	0.81	1.93	4.72
h_3	4.25	2.89	3.68	0.62	2.09	5.91
h_4	11.46	3.03	3.88	0.45	2.02	5.95
h_5	10.46	3.92	3.66	0.83	3.29	7.01
h_6	7.87	7.08	6.31	1.45	5.13	14.66
h_7	13.79	4.40	4.61	0.81	3.16	9.04
h_8	2.65	4.03	4.26	0.83	2.83	8.87
h_9	1.24	3.84	4.81	0.32	2.61	7.95
h_{10}	2.26	4.63	4.03	1.05	3.62	8.56
h_{11}	2.72	2.73	2.70	0.55	2.09	5.24
Other	25.66	1.75	2.53	0.18	1.31	3.74

Notes: This table displays descriptive statistics for our estimating admissions dataset. Only admissions on the hospitals in the sample are considered for these statistics. Panel A displays statistics across all hospitals in the sample. Panel B displays statistics for market shares and full prices by hospital.

Table A.3: Preferential tiering and admission flows between hospitals and insurers

	A - Percent of plans in which hospital is in preferential tier					**B - Percent of admissions by patient insurer**				
Hospital	m_a	m_b	m_c	m_d	m_e	m_a	m_b	m_c	m_d	m_e
h_1	0.99	4.30	43.81	15.08	0.75	5.17	30.74	31.42	25.55	7.11
h_2	69.17	3.49	0.00	48.24	0.00	57.13	10.93	7.77	21.49	2.68
h_3	49.21	20.16	6.35	43.97	41.04	63.84	7.31	5.57	21.86	1.43
h_4	46.84	95.70	4.44	53.27	0.00	11.48	70.76	5.21	10.88	1.67
h_5	15.02	45.43	1.59	52.51	81.34	8.40	18.29	22.06	24.28	26.97
h_6	0.00	4.03	0.95	8.29	0.00	4.84	35.81	30.96	20.55	7.84
h_7	0.59	74.73	23.17	31.66	46.27	4.10	56.44	15.86	17.46	6.14
h_8	36.76	0.00	0.32	38.19	0.00	42.93	18.38	16.15	18.51	4.03
h_9	0.00	0.00	0.00	0.50	10.45	19.91	0.73	8.78	69.84	0.74
h_{10}	0.00	0.27	17.46	0.00	9.33	2.91	17.94	58.87	14.86	5.42
h_{11}	6.32	48.92	2.86	7.79	0.00	16.53	66.70	5.99	8.40	2.38
Other	-	-	-	-	-	25.49	13.51	18.67	30.59	11.75

Notes: Panel A displays the share of plans offered in the spot market by each insurer that has a hospital in its preferential tier, with hospitals in the rows and insurers in the columns. The sample includes the set of plans that enter demand estimation and the supply-side analysis. Panel B displays a breakdown of the admission shares. Each cell in columns labeled $m_a - m_e$ displays the share of admissions that a given hospital received from each insurer. Cells that relate to VI firms are underlined. Recall that m_c and h_{10} are only VI in the first year of our sample.

Table A.4: Vertical integration and market outcomes (within insurer-hospital)

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	A - Plan design		**B - Admission outcomes**				**C - Hospital outcomes**	
	Preferential hospital	Coverage rate	log cost proxy	log # services	Re-admission	log price	Share of admissions	log revenue
VI	0.169	5.224	0.040	-0.012	0.005	-0.076	0.044	0.297
	(0.006)	(0.325)	(0.025)	(0.011)	(0.006)	(0.021)	(0.004)	(0.051)
N	15,741	15,741	567,752	567,752	204,223	567,752	264	264
R^2	0.212	0.254	0.212	0.613	0.059	0.694	0.988	0.976
Mean non-VI	0.137	63.101	2.534	16.039	0.081	5.269	0.128	8,865.659
H-I FE	Y	Y	Y	Y	Y	Y	Y	Y
Year FE	Y	Y	N	N	N	N	Y	Y
Interacted FE	N	N	Y	Y	Y	Y	N	N
Plan FE	N	N	Y	Y	Y	Y	N	N
Cost proxy	N	N	N	Y	Y	Y	N	N
Controls	N	N	Y	Y	Y	Y	N	N
Observation	plan-hospital-year		admission				insurer-hospital-year	

Notes: This table shows results from estimating equations (1), (2) and (3) with hospital-insurer FEs. This specification relies exclusively on the disintegration of insurer m_c and hospital h_{10} in 2014. The interaction between m_c and h_{10} provides a small case study covering only 1.8 percent of admissions. The unit of observation for each regression is reported in the bottom row. Regressions in columns (3)–(6) include the following controls: diagnosis fixed effects, patient age, gender, policyholder income, policyholder employment status, and county fixed effects, along with the fixed effects indicated in the table. Columns (4)–(6) also include admission prices in the public system, interacted with hospital dummies. Column (5) only includes admissions for circulatory, infections, pregnancy, and respiratory diagnoses. Mean non-VI indicates the mean of the dependent variable for non-VI observations, measured in levels. Interacted FE indicates diagnosis-hospital, diagnosis-year, and hospital-year fixed effects. Standard errors in parentheses are clustered at the insurer-hospital level.

Table A.5: Relationship between patient observables and VI

	(1)		(2)	(3)	(4)
			Regression results		
Patient attribute	Mean non-VI		No controls	Controls	Within I-H
Age	31.468	$\hat{\beta}_{VI}$	2.882	1.126	-0.986
		S.E.	(0.896)	(0.318)	(0.391)
		R^2	0.006	0.290	0.291
Female	0.590	$\hat{\beta}_{VI}$	-0.059	-0.011	0.019
		S.E.	(0.022)	(0.006)	(0.015)
		R^2	0.003	0.240	0.241
Employed	0.851	$\hat{\beta}_{VI}$	0.020	-0.004	-0.012
		S.E.	(0.014)	(0.003)	(0.022)
		R^2	0.001	0.045	0.046
log(Income)	2.069	$\hat{\beta}_{VI}$	-0.170	-0.039	0.018
		S.E.	(0.046)	(0.017)	(0.120)
		R^2	0.001	0.031	0.031
Interacted FE			N	Y	Y
Plan FE			N	Y	Y
County FE			N	Y	Y
I-H FE			N	N	Y

Notes: This table shows results from estimating equation (2) using patient observables as dependent variable, as indicated in the first column. The unit of observation is an admission. Mean non-VI indicates the mean of the dependent variable for non-VI observations, measured in levels. Each column is a different specification, as indicated in the bottom panel. For each regression, we report the coefficient associated to $\text{VI}_{m(j)ht}$. Interacted FE indicates diagnosis-hospital, diagnosis-year, and hospital-year fixed effects. Standard errors in parentheses are clustered at the insurer-hospital level.

Table A.6: Vertical integration and hospital treatment behavior

	(1)	(2)	(3)	(4)	(5)
	C-section	Ultrasound	Hemogram	Chest X-ray	Imaging
A - Main specification					
VI	-0.112	-0.001	-0.022	0.015	0.003
	(0.030)	(0.001)	(0.020)	(0.009)	(0.002)
Interacted FE	Y	Y	Y	Y	Y
Plan FE	Y	Y	Y	Y	Y
Cost proxy	Y	Y	Y	Y	Y
Controls	Y	Y	Y	Y	Y
B - Within hospital-insurer					
VI	0.116	0.003	0.009	-0.017	0.011
	(0.081)	(0.003)	(0.042)	(0.051)	(0.006)
H-I FE	Y	Y	Y	Y	Y
Interacted FE	Y	Y	Y	Y	Y
Plan FE	Y	Y	Y	Y	Y
Cost proxy	Y	Y	Y	Y	Y
Controls	Y	Y	Y	Y	Y
N	77,715	77,715	99,265	61,518	61,518
R^2	0.440	0.030	0.096	0.179	0.012
Mean non-VI	0.563	0.003	0.558	0.301	0.018

Notes: This table shows results from estimating equation (2). The unit of observation is an admission. The dependent variable is an indicator of whether a service that relies at least partially on physician discretion was delivered to the patient. These regressions include the following controls: diagnosis fixed effects, admission prices in the public system interacted with hospital dummies, patient age, gender, policyholder income, policyholder employment status, and county fixed effects, along with the fixed effects indicated in the table. Mean non-VI indicates the mean of the dependent variable for non-VI observations, measured in levels. Interacted FE indicates diagnosis-hospital, diagnosis-year, and hospital-year fixed effects. Standard errors in parentheses are clustered at the insurer-hospital level.

Table A.7: Instrument first stage and estimated consumer preferences without instruments

	(1)	(2)	(3)	(4)
	A - Healthcare		**B - Insurance**	
	Coef.	S.E.	Coef.	S.E.
I - Instrument first stage				
Predicted closed-insurer OOP	0.754	(0.000)		
Rival actuarially fair premiums			2.292	(0.042)
N	10,474,104		23,160	
R^2	0.605		0.856	
II - Estimates without instruments				
A: Price (α_i^H) / B: Monthly Premium (α_i^M)				
$\times$ Age $\in [25, 40)$	-1.410	(0.010)	-25.867	(0.051)
$\times$ Age $\in [40, 55)$	-1.178	(0.019)	-23.613	(0.051)
$\times$ Age $\in [55, 65]$	-1.139	(0.010)	-24.029	(0.052)
$\times$ Female $\times$ Single	0.243	(0.010)	10.134	(0.050)
$\times$ Has dependents	0.179	(0.009)	13.231	(0.047)
$\times$ High income	0.272	(0.005)	11.878	(0.024)
Distance (β^H)	-0.092	(0.001)		
VI Marketing (γ^H)	2.400	(0.006)		
Network (β^M)			0.921	(0.005)
Median elasticity	-0.67		-1.96	
N	261,857		163,034,142	

Notes: Panel I presents the key first-stage estimates associated with the demand instruments. Panel A presents the regression of consumer out-of-pocket prices for each option within their hospital demand choice set on the instrument. The regression includes distance and year-provider-diagonsis-insurer fixed effects, matching the covariates that are included in the demand model. Panel B presents the regression of plan premiums on the plan's rival actuarially fair premiums. The regression also includes the network utility instruments (average rival network utilities, share of rival plans with the same preferential hospitals, and share of other plans of the same insurer in the same segment with the same preferential hospitals). It also includes insurer-age fixed effects to match our specification of plan preferences. The sample is substantially smaller as this regression is at the plan-year level as premiums do not vary across consumers conditional on plan. Panel II shows the estimated preferences when estimated without the instruments. Panel A presents estimates of preferences for hospitals. The sample size is a 30 percent random sample of non-emergency inpatient events used to estimate preferences. The model includes insurer-hospital-diagnosis-year fixed effects, omitted from the table. Panel B presents estimates of preferences for plans. The model includes an insurer-year fixed effect, omitted from the table. Heterogeneity in price and premium preferences depend on policyholder attributes, where high income indicates those above the median income. Prices, premiums, and network surplus are measured in thousands of dollars. Network surplus is measured based on yearly risk and spending. Distance is measured in miles from neighborhood centroids to hospitals. The reported elasticities are the median own-price in Panel A and own-premium in Panel B.

Table A.8: Heterogeneity in effects of vertical integration consumer surplus

	A - Full effect	**B - Decomposition**	
		Short run	Medium run
	Change	Change	Change
Policyholder neighborhood			
Center	6.94	-2.85	9.79
North	35.53	-23.74	59.27
Northeast	16.02	-34.46	50.48
Northwest	8.72	-9.90	18.62
Periphery	45.10	-27.13	72.23
South	25.88	-46.30	72.18
Southeast	39.24	-46.61	85.85
Southwest	22.29	-20.59	42.89
Policyholder household type			
Single female	1.27	-8.09	9.36
Single male	5.86	-3.03	8.90
Has dependents	32.57	-45.90	78.46
Policyholder income			
Below median	10.88	-7.47	18.35
Above median	36.23	-80.86	117.09
Policyholder age			
Below 45	9.75	-13.99	23.74
Above 45	20.79	-28.16	48.94
Policyholder insurance program			
Private insurance	28.73	-34.78	63.51
Public insurance	11.19	-17.12	28.32

Notes: This table displays average changes in consumer surplus per policyholder by population groups in dollars relative to the status quo. Panel A displays the Full effect of banning VI. Panel B displays partial changes: Short run keeps coverage fixed, and Medium run shows the additional impact of coverage adjustments. Their sum is the Full effect.

Table A.9: Effects of single-firm vertical integration on plan design and hospital prices

	(1)	(2)	(3)	(4)	(5)	(6)
	Baseline		**A - Ban m_a's VI**		**B - Ban m_b's VI**	
	Raw	Weighted	Raw change	Weighted change	Raw change	Weighted change
Hospital prices (by hospital / insurer)						
Within m_a	3.43	3.43	0.27	-0.18	16.18	13.49
m_a-VI / m_b	4.07	4.07	-0.16	-0.25	-0.19	0.01
m_a-VI / Non-VI	3.77	3.77	-0.29	-0.25	-0.47	-0.48
Within m_b	3.88	3.88	-0.79	-0.80	-0.22	-0.30
m_b-VI / m_a	3.59	3.59	1.02	-0.01	45.52	0.92
m_b-VI / Non-VI	3.61	3.61	1.17	-0.15	0.49	-0.54
Non-VI / m_a	5.71	5.71	0.16	0.09	1.95	2.74
Non-VI / m_b	6.10	6.10	0.09	1.16	0.55	0.96
Non-VI / Non-VI	5.44	5.44	1.10	0.24	1.83	0.76
Public / Non-VI	1.58	1.58	0.00	-0.03	0.00	-0.03
Public / m_a	1.57	1.57	0.00	-0.02	0.00	-0.05
Public / m_b	1.57	1.57	0.00	0.00	0.00	-0.06
Total	4.32	4.32	0.45	-0.03	4.72	0.64
Premiums (by insurer)						
Insurer m_a	1.11	1.11	0.11	0.07	0.02	-0.53
Insurer m_b	1.40	1.40	0.55	0.49	0.17	0.10
Non-VI	1.38	1.38	-0.02	-0.08	-0.07	-0.12
Total	1.31	1.31	0.12	0.09	-0.00	-0.18
Plan design (by insurer)						
m_a - Base coverage	58.52	58.90	-6.08	-8.54	-25.95	-29.15
m_a - Preferential coverage	81.54	83.46	-0.71	-1.92	3.69	4.52
m_a - Self-preferencing	57.60	53.04	-31.73	-30.35	0.27	19.83
m_a - Other-VI-preferencing	16.80	14.78	-12.53	-6.73	10.40	12.75
m_a - Non-VI-preferencing	2.88	2.45	0.00	-0.87	8.00	1.22
m_a - Star-hospital-preferencing	0.00	0.00	6.80	3.76	27.20	9.18
m_b - Base coverage	43.63	43.07	-13.60	-13.17	-9.36	-7.73
m_b - Preferential coverage	73.12	72.51	13.12	9.52	11.26	9.00
m_b - Self-preferencing	77.38	77.82	2.78	9.39	-8.33	-15.79
m_b - Other-VI-preferencing	10.32	11.73	-10.32	-11.73	-0.40	0.32
m_b - Non-VI-preferencing	7.38	6.27	30.71	30.65	2.38	-1.01
m_b - Star-hospital-preferencing	1.19	0.68	94.05	91.63	7.14	3.80
Non-VI - Base coverage	58.88	58.67	-14.26	-14.75	-19.95	-18.74
Non-VI - Preferential coverage	85.75	84.94	0.78	0.91	-1.57	-1.77
Non-VI - Other-VI-preferencing	22.12	21.40	-8.00	-7.84	-8.62	-8.00
Non-VI - Non-VI-preferencing	16.91	17.98	-6.36	-9.28	-6.87	-8.51
Non-VI - Star-hospital-preferencing	13.97	18.22	2.39	-4.22	2.02	-2.47
Healthcare spending						
Actuarial value	0.66	0.66	-0.02	-0.03	-0.01	-0.02
Inpatient spending | private plan	0.82	1.04	0.03	0.01	0.18	0.20
Inpatient spending	1.16	0.50	0.03	-0.00	0.18	0.06
Total household spending	3.02	1.10	0.16	0.04	-0.00	-0.06

Notes: Prices in thousands of dollars per unit of resources, premiums in thousands per year, coverage in percentages, and healthcare spending in thousands per household. Actuarial value is the share of expected payments covered by insurers. VI Self-preferencing is the likelihood that a VI hospital is preferential in a VI plan. Other-VI-preferencing and Non-VI-preferencing are analogous for other-VI and non-VI hospitals. Odd columns display raw averages: for prices, it is across insurer-hospital; for premiums and coverages, it is across plans. Even columns display weighted averages by demand: for prices, it is by demand per unit of resources; for premiums, coverage, and spending, it is by plan demand. We omit unweighted spending since it is necessarily linked to plan enrollment probabilities. Panels A and B display the Full effect of banning VI for m_a and m_b, respectively.

Table A.10: Effects of single-firm vertical integration ban on choices and welfare

	(1)	(2)	(3)
		A - Ban m_a's VI	**B - Ban m_b's VI**
	Baseline	Change	Change
Efficiency			
Moral hazard spending	44.10%	-2.26%	-14.72%
Market shares			
m_a-VI hospital	11.11%	-5.44	3.93
m_b-VI hospital	27.56%	-4.45	-12.90
Non-VI hospital	30.95%	3.28	-9.85
Insurer m_a	3.69%	-0.84	2.04
Insurer m_b	7.15%	-1.50	-2.39
Non-VI insurer	13.18%	1.50	1.66
Admission shares (by hospital / insurer)			
Within m_a	55.84%	-43.71	5.03
m_a-VI / m_b	6.42%	-1.29	3.43
m_a-VI / Non-VI	37.74%	45.00	-8.45
Within m_b	78.51%	-14.37	-52.81
m_b-VI / m_a	2.69%	1.08	1.77
m_b-VI / Non-VI	18.80%	13.29	51.05
Non-VI / m_a	9.96%	0.15	0.95
Non-VI / m_b	12.99%	16.31	12.44
Non-VI / Non-VI	77.05%	-16.46	-13.39
Profits (in millions)			
m_a hospitals	34.900	-17.367	385.295
m_b hospitals	70.141	-45.474	-39.674
Non-VI hospitals	109.570	10.941	23.416
Insurer m_a	159.950	-25.287	-455.203
Insurer m_b	394.733	109.736	-103.417
Non-VI insurers	713.817	58.776	79.479
Consumer surplus			
m_a enrollees (per member)	-	0.335	0.669
m_b enrollees (per member)	-	-0.029	-0.230
Non-VI enrollees (per member)	-	5.773	7.541
Total consumer surplus (millions)	-	257.418	154.284
Share better off	-	0.603	0.806
Share better off relative to m_a ban	-	-	0.860
Total welfare (in millions)	-	348.74	44.18

Notes: Moral hazard spending is relative to the first best inpatient spending. Profits and total consumer surplus are measured in millions of dollars per year. Consumer surplus for VI enrollees is the average surplus conditional on enrolling in a VI plan, unweighted by demand. Non-VI consumer surplus is defined analogously. All values are in thousands of dollars unless stated otherwise. Panels A and B display the Full effect of banning VI for M_a and m_b, respectively.