NBER WORKING PAPER SERIES

VERTICAL INTEGRATION AND PLAN DESIGN IN HEALTHCARE MARKETS

José Ignacio Cuesta
Carlos E. Noton
Benjamin Vatter

Working Paper 32833
http://www.nber.org/papers/w32833

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
August 2024, Revised May 2025

This paper supersedes a previous version titled "Vertical Integration between Hospitals and Insurers". We thank our discussants, Christian Cox, Martin Gaynor, Kate Ho, and Shoshana Vasserman for their valuable suggestions. We also thank Mert Demirer, Joe Doyle, David Dranove, Liran Einav, Gautam Gowrisankaran, Igal Hendel, Ali Hortaçsu, Gastón Illanes, Robin Lee, Victoria Marone, Elena Prager, William Rogerson, Pietro Tebaldi, Robert Town, Michael Whinston, Ali Yurukoglu, and seminar and conference participants for helpful comments. We also thank Aaron Banks for superb research assistance. Finally, we thank the Superintendence of Health of Chile for data access and useful discussions, particularly Claudia Copetta, David Debrott, and Eduardo Salazar; and Alejandra Benítez and Manuel Inostroza for useful conversations about the institutional framework. Noton acknowledges financial support from Fondecyt 1170029 and the Institute for Research in Market Imperfections and Public Policy, ICM IS130002. This research was partially supported by the NLHPC (ECM-02) computing infrastructure. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

Vertical Integration and Plan Design in Healthcare Markets
José Ignacio Cuesta, Carlos E. Noton, and Benjamin Vatter
NBER Working Paper No. 32833
August 2024, Revised May 2025
JEL No. I11, L13, L40

ABSTRACT

We study the impacts of vertical integration between insurers and hospitals. We develop a bilateral oligopoly model of competition with endogenous plan design and hospital prices, and estimate it using administrative data from Chile. Integrated insurers offer plan networks that favor integrated hospitals, and non-integrated insurers restrict access to integrated hospitals to limit the raising rivals' costs threat. These skewed networks restrict substitution across hospitals, softening price competition. Whereas vertical integration reduces double marginalization, plan design distortions and weaker hospital competition more than compensate. Vertical integration leads to a misallocation of demand across hospitals, higher healthcare spending, and lower welfare.

José Ignacio Cuesta
Stanford University
Department of Economics
and NBER
jicuesta@stanford.edu

Benjamin Vatter
Massachusetts Institute of Technology
MIT Sloan School of Management
and NBER
bvatter@mit.edu

Carlos E. Noton
Pontifical Catholic University of Chile
carlos.noton@uc.cl

A online appendix is available at http://www.nber.org/data-appendix/w32833

1 Introduction

Vertical integration (VI) in healthcare markets has become a significant concern for policymakers and regulators (DOJ, 2024; FTC, 2024). As of 2022, four out of five physicians in the U.S. are employed by hospitals, insurers, or other corporate entities (PAI, 2024). Seventy percent of commercial drug coverage is provided by insurers integrated with pharmacy benefit managers (Guardado, 2023), and about half of all inpatient care is provided at hospitals that also sell insurance (AHRQ, 2023). Despite the prevalence of VI in healthcare, the evidence quantifying its impacts is limited (Handel and Ho, 2021).

Hospitals and insurers interact in a vertical market. Downstream insurers sell plans to consumers, offering access to upstream hospitals. VI between these players has theoretically ambiguous welfare effects. On the positive side, VI aligns insurer and hospital incentives to eliminate double marginalization and limit wasteful spending (Spengler, 1950; Williamson, 1971). On the negative side, it grants market power to integrated firms, creates incentives for VI hospitals to increase rival insurers' costs or foreclose their access, and for VI insurers to limit patient access to rival hospitals (Salop and Scheffman, 1983; Hart and Tirole, 1990; Ordover *et al.*, 1990).

We study the equilibrium effects of VI between hospitals and insurers in the Chilean private healthcare market. The setting offers useful variation in VI and rich administrative data on claims, enrollment, ownership, and cost-sharing rules. We combine these data with a bilateral oligopoly model with endogenous cost-sharing, networks, premiums, and hospital prices. Comparing the status quo to a counterfactual without VI, we find that VI firms steer demand to their hospitals by charging low premiums and low out-of-pocket prices at integrated hospitals, and offering limited coverage at rival hospitals. The generous coverage offered at VI hospitals induces non-VI insurers to provide high coverage at non-VI hospitals. These differences in networks vertically segment the market, softening hospital competition, raising prices, and lowering consumer surplus and total welfare. Crucially, an evaluation of the impacts of VI that only accounted for effects on prices and premiums but not on plan design would fail to rationalize the stark plan design heterogeneity across insurers, and deliver an opposite prediction on welfare effects.

The Chilean private healthcare market provides a well-suited setting for our study. It comprises a handful of hospitals and five insurers competing for enrollees in a regulated market. We study the 2013–2016 period, which features a well-established VI segment accounting for half of hospital capacity, as well as the de-integration of a VI firm. In this market, plans feature tiered hospital networks, each with a base tier with moderate cost-sharing, and a preferential tier of hospitals with more generous cost-sharing. Tiered networks and access regulation imply that VI hospitals admit patients from all insurers, and access to care is mediated only by cost-sharing and prices.

We develop a descriptive analysis that reveals clear differences between VI and non-VI firms. VI insurer plans offer 5.8 percentage points lower cost-sharing at integrated hospitals and are 34 percentage points more likely to have them in their preferential tiers than rival plans. Moreover, VI hospitals charge 13 percent lower prices to their integrated insurer's enrollees, conditional on diagnosis, demographics, and complexity. Comparing similar patients arriving at VI hospitals, we find no evidence that VI reduces hospital spending, impacts clinical treatment decisions, or affects medical quality outcomes. These findings hold when comparing across all VI hospitals and insurers in the cross-section and when comparing over time within the hospital-insurer pair that de-integrated. Moreover, while the average VI insurer has less than a quarter market share, it accounts for two-thirds of admissions at its integrated hospitals. These results suggest that VI strongly affects market outcomes. However, they do not inform whether consumers benefit from these impacts, as VI affects all firms in equilibrium. To answer this question, we combine data and theory to recover the primitives governing the equilibrium effects of VI.

We model the interaction between hospitals, insurers, and consumers as the subgame perfect equilibrium of a four-stage game. In the first stage, insurers design plans, choosing hospital tiers and coverage. In the second stage, non-VI hospitals and insurers bargain over hospital prices, VI firms set prices to maximize joint profits, and insurers set premiums. In the third stage, households choose insurance plans. Finally, health risks are realized, and consumers choose hospitals given out-of-pocket prices and quality. This model bridges work on insurer-hospital competition (Gowrisankaran *et al.*, 2015; Ho and Lee, 2017, 2019) and VI (Crawford *et al.*, 2018). Its main novelty is to endogenize plan design.

Our model captures key incentives induced by VI. First, VI hospitals internalize that lower prices increase their integrated insurer's profits, mitigating double marginalization. Second, VI hospitals account for how increasing prices to rival insurers leads enrollees to switch to their integrated insurer. Third, VI insurers internalize that lower coverage at rival hospitals steers demand toward their hospitals. Finally, VI insurers account for the value of their plans in attracting demand to their hospitals. These forces incentivize VI firms to reduce out-of-pocket prices for patients within their system, increase prices to rival insurers, skew plan coverage in favor of integrated hospitals, and set lower premiums. Non-VI rival insurers and hospitals do not face these incentives directly but compete in their presence.

We identify and estimate the primitives driving the impacts of VI. On the demand side, we identify preferences over hospitals by examining how consumers trade off out-of-pocket prices, distance, and quality. We use data on contracts between hospitals and an isolated insurance market segment to form instruments to address price endogeneity, and leverage the de-integration event to identify how VI firms may steer demand through means beyond prices and quality (e.g., referrals). Moreover, we use data on enrollment choices and plan attributes to estimate preferences over plans.

On the supply side, we develop general results for the identification of Nash-in-Nash bargaining models in bilateral oligopoly. Following Grennan (2013), we show that bargaining weights and costs cannot be separately identified without imposing additional structural restrictions. We prove a necessary and sufficient condition for additional structure to achieve identification, providing conditions under which data on hospital average or total costs are sufficient to identify these objects; a strategy we follow in our application. Finally, we use optimal plan tiering conditions to formulate moment inequalities and identify plan design costs.

We estimate that demand is more elastic to premiums than to out-of-pocket hospital prices. This limits VI hospitals' ability to benefit from raising rivals' costs as insurers limit the pass-through of costs to premiums by allocating some of them to consumers' out-of-pocket prices, mitigating demand losses. We also find that hospitals are meaningfully horizontally differentiated. The combination of heterogeneous quality across hospitals and diagnoses with heterogeneous medical risk within households reduces the appeal of narrow networks. This incentivizes VI insurers to include non-VI hospitals in their networks, and rival insurers to include VI hospitals in theirs. Moreover, we find that patient steering by VI insurers and cost-sharing lead to the misallocation of nearly half of all medical spending relative to a theoretical first-best in which hospital care is priced at cost and accessed without distortions. This misallocation largely benefits VI hospitals.

Using our model, we quantify the equilibrium effects of VI by simulating a VI ban. Removing VI eliminates the incentive for VI insurers to steer consumers toward their hospitals. As a result, formerly VI insurers adjust their plans to reduce their generosity at previously integrated hospitals and improve access to high-quality non-VI hospitals. Similarly, banning VI eliminates VI hospitals' incentives to raise rivals' costs, inducing non-VI insurers to increase their plans' generosity at formerly VI hospitals, particularly high-quality ones. Ultimately, banning VI intensifies competition among high-quality hospitals, driving average paid prices down by 11 percent. VI hospitals see profits decline by 43 percent, underscoring how VI segments the market vertically, insulating VI hospitals from price competition. These price reductions lead to a 2 percent decrease in premiums, outweighing the reintroduction of double marginalization. In terms of welfare, consumer surplus increases by \$62.7 million, and total welfare rises by \$41.7 million. These gains accrue primarily to non-VI insurers and their enrollees, with insurers gaining 19 percent in profits and consumers an equivalent of 5 monthly premiums in surplus.

The results highlight a critical role for plan design in shaping the effects of VI. To explore this, we simulate a VI ban in which prices and premiums adjust and consumers reallocate between insurers and hospitals, but plan networks and cost-sharing remain fixed. In this scenario, banning VI negatively impacts consumer surplus, firm profits, and overall welfare. Intuitively, absent plan design responses, removing VI reintroduces double marginalization and eliminates incentives to raise rivals' costs, but the skewed structure of plans offered under VI constrains insurers' ability to reallocate demand effectively. Our finding thus indicates that the welfare losses from double marginalization are greater than those from the raising rivals' cost incentives. Overall, this result suggests the welfare losses from VI primarily arise from plan design distortions that impact hospital access. Moreover, absent regulatory intervention, capturing the beneficial aspects of VI without inducing plan design distortions appears infeasible since insurers face strong incentives to redesign their networks.

To inform the broader discussion on insurer-hospital integration, we explore the role of factors driving the effects of VI. First, we show that VI harms more consumers when higher-quality hospitals are integrated, since that limits access to quality care. Second, we find that to overturn the negative impact of VI, cost efficiencies would have to exceed 18 percent, an order of magnitude larger than recent estimates (Schmitt, 2017; Demirer and Karaduman, 2024). Finally, we find that for quality gains to overturn the welfare losses, VI hospitals would have to provide an additional \$473 worth of quality per admission, much higher than recent estimates of quality effects of VI (Cho, 2025).

Our paper contributes to the literature on healthcare competition by providing new evidence on the effects of consolidation (Handel and Ho, 2021). We build upon the literature on horizontal mergers of hospitals (Dafny, 2009; Gowrisankaran *et al.*, 2015; Craig *et al.*, 2021) and insurers (Dafny *et al.*, 2012; Chorniy *et al.*, 2020; Ho and Lee, 2017). While there is recent work on physician-hospital integration (Baker *et al.*, 2016; Capps *et al.*, 2018; Cutler *et al.*, 2020; Koch *et al.*, 2021) and insurer-pharmacy benefit manager integration (Brot-Goldberg *et al.*, 2022; Gray *et al.*, 2023), our paper fills a gap in the literature on insurer-provider integration. We contribute a model to evaluate the equilibrium impacts of VI and evidence from a compelling setting. Our work complements research comparing VI and non-VI plans' costs and quality (Johnson *et al.*, 2017; Parekh *et al.*, 2018; Park *et al.*,

2023) and recent work examining integrated care delivery (Cho, 2025).[1]

Our study also contributes to the literature on contract design. First, relative to previous work examining single-agent design problems (Chade *et al.*, 2022; Marone and Sabety, 2022; Tilipman, 2022; Ho and Lee, 2023), we study oligopolistic competition with endogenous plan design. This connects to recent work on insurer competition over non-price attributes (Miller *et al.*, 2023; Vatter, 2025), although our model accounts for the impact of plan design on both downstream and upstream outcomes, including upstream prices. Second, by focusing on the network aspect of plan design, we link the literature on contract design with that on network formation (Lee and Fong, 2013; Ho and Lee, 2019; Liebman, 2022; Shepard, 2022; Serna, 2024). This literature has, with few exceptions (e.g., Ghili 2022), focused on single-agent problems and exclusionary network formation. We build on this literature by studying competitive intensive margin adjustments in networks, where insurers choose cost-sharing at each hospital.[2]

Finally, our work connects the literature on vertical integration and healthcare competition. We build on extensive work on vertical contracting (Lee *et al.*, 2021) and vertical mergers (Hortaçsu and Syverson, 2007; Atalay *et al.*, 2014; Crawford *et al.*, 2018; Luco and Marshall, 2020; Chen *et al.*, 2024), adapting the framework of Crawford *et al.* (2018) to the healthcare setting. We contribute to the study of VI by allowing firms to change their input choices (i.e., hospital networks and coverage) in response to VI. This enables firms to reposition in response to changes in vertical market structure. By accounting for this channel, we contribute to the broader literature on competition with endogenous product attributes (Eizenberg, 2014) and on firm repositioning in response to mergers and policy changes (Fan, 2013; Wollmann, 2018; Fan and Yang, 2022; Bontemps *et al.*, 2023).

The remainder of the paper is organized as follows. Section 2 provides a conceptual framework that clarifies the key mechanisms we seek to uncover. Section 3 describes our setting and data. Section 4 provides descriptive evidence for how VI shapes market outcomes. Section 5 introduces our model, and Section 6 discusses its identification, estimation, and main estimates. Section 7 discusses our counterfactual analysis of the impacts of banning VI. Finally, Section 8 concludes.

---

[1] There is also some theoretical work on foreclosure in insurer-hospital contracting (Gal-Or, 1999).

[2] Much of the network formation literature focuses on how networks shape bargaining outcomes. Our model allows for network structure to affect bargaining leverage, but not through threats of replacements.

2 Conceptual Framework

To illustrate the main economic forces at play in our setting, we develop a simple 2×2 bilateral oligopoly example, illustrated in Figure 1a. Upstream, two hospitals $h \in \{h_1, h_2\}$ provide services at negotiated prices p_{mh} to the enrollees of two downstream insurers $m \in \{m_a, m_b\}$. Insurers sell access to care, offering consumers to cover a fraction $c_{mh} \in [0, 1]$ of hospital prices in exchange for a premium ϕ_m. In this framework, coverage levels determine hospital network structure; a coverage rate of zero effectively means that the hospital is excluded from the network. VI is an agreement between an insurer and a hospital to maximize joint profits. To describe the incentives created by VI, we consider how the integration of h_1 and m_a may impact coverage and prices.

Figure 1b illustrates the integration of h_1 and m_a and the comparative statics of the VI firm's choices, holding all other variables at their status quo. First, VI leads insurer m_a to internalize the profits of its hospital. This incentivizes m_a to retain demand within the integrated system by increasing coverage at its hospital ($\uparrow c_{a1}$) and decreasing it at rivals ($\downarrow c_{a2}$).[3] Effectively, VI introduces an incentive to skew the insurer's hospital network, by *self-preferencing* integrated hospitals and *narrowing the network* access to rivals.

Similarly, VI leads hospital h_1 to internalize profits accrued at its integrated insurer m_a. This prompts the hospital to lower prices for enrollees of m_a ($\downarrow p_{a1}$) to incentivize downstream enrollment, eliminating *double marginalization* within the VI firm. In addition, VI also creates an incentive for h_1 to intensify price negotiations with the rival insurer, as higher out-of-pocket prices for m_b's enrollees induce them to substitute in favor of the integrated m_a. Effectively, VI creates an incentive for *raising rivals' costs* ($\uparrow p_{b1}$).[4]

Central to our analysis, VI also distorts the incentives of rival insurers and hospitals, as shown by Figure 1c. First, VI implies that insurer m_b faces a negotiation with a more aggressive hospital h_1, which benefits from increasing prices beyond their impact on upstream profits. Second, VI implies that h_2 faces decreased demand from m_a's enrollees due to network skewing, increasing its dependence on m_b for revenue. These forces prompt m_b to skew coverage away from h_1 ($\downarrow c_{b1}$) and in favor of the non-VI h_2 ($\uparrow c_{b2}$).

[3]We focus on a setting in which upstream demand is sensitive to prices. In settings with low intensive margin price-sensitivity but high sensitivity to extensive-margin changes in coverage structure, equivalent shifts in coverage may attain the same effects. For example, a decrease in coverage can be replicated by removing a hospital from the network for certain services.

[4]With price-sensitive upstream demand, a hospital foreclosing an insurer is a limiting case of raising rivals' costs. Similarly, in the limit, self-preferencing leads to full coverage within the VI system, eliminating the need for an internal transfer price. Consumption moral hazard and regulation, however, might dissuade VI insurers from offering full coverage at their hospitals, which is the case in our empirical application.

Figure 1: Illustration of VI incentives and competitive responses

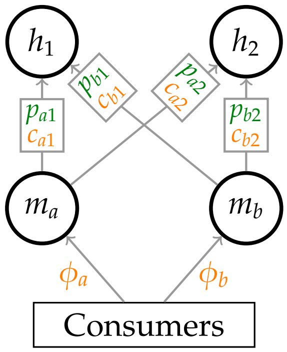

(a) Setup

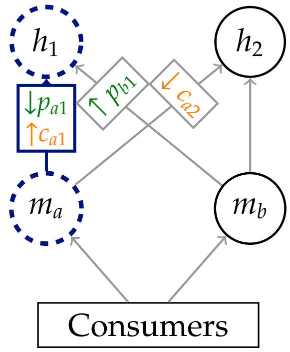

(b) Direct effects of VI

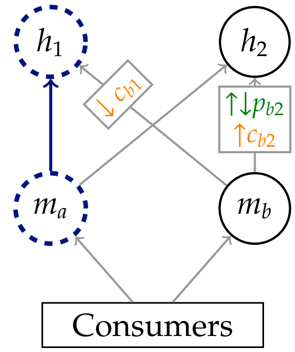

(c) Competitive effects of VI

Notes: These figures illustrate a simple vertical market with upstream hospitals (h_1, h_2) and downstream insurers (m_a, m_b). Figure (a) shows the market before VI, with negotiated prices (p), coverage rate (c), and premiums (ϕ). Prices (in green) are negotiated, while coverage and premiums (in orange) are determined by insurers. Figure (b) illustrates the direct effect of the integration of h_1 and m_a on their decision variables (in blue), holding all other variables fixed. It also marks the effect of the integration on the negotiated price between h_1 and m_b (in green). Figure (c) shows the competitive response of m_b to the raising rivals' costs incentive of h_1.

However, increased reliance on h_2 and decreased demand elasticity due to higher coverage increases the leverage of h_2 in its negotiations with m_b, leading to ambiguous price effects ($\uparrow\downarrow$ p_{b2}). Intuitively, the greater the value of h_2 relative to h_1 to consumers (e.g., by providing specialty services or higher quality), the less impacted it will be by decreased coverage from m_a, and the more its relative bargaining position with m_b will improve following the integration of h_1 and m_a, hence the more likely that p_{b2} will increase.

This framework clarifies that plan design enhances the ability of VI firms to retain demand within their system, and of their rivals to weaken the raising rivals' cost threat. At a high level, the economics of the problem are those of differentiated input choices by an oligopoly of consumer-facing platforms. Whether inputs are hospitals in insurance plans or television channels bundled in a cable provider's system, VI will incentivize input-skewing by downstream firms. Thus, accounting for input choice (or plan design) is key for capturing the full impacts of VI and identifying changes in upstream versus downstream market power. In our application, we quantify these channels. We find that network skewing softens price competition and limits access, and that the associated welfare losses exceed the gains from reduced prices and premiums within VI firms.

3 Setting and Data

3.1 The Chilean Healthcare Market

The Chilean health insurance system combines public and private provision. Public insurance is provided through a government-managed plan (*Fondo Nacional de Salud*, FONASA), funded by a mandatory contribution of 7 percent of taxable income by workers and retirees.[5] Private insurance is provided by a small number of firms (*Instituciones de Salud Previsional*, ISAPREs) offering plans in a regulated market, funded through premiums. Individuals can allocate their mandatory contributions to purchase private plans. These plans are often expensive, resulting in additional premium payments for most enrollees. During our study period, 77.3 percent of consumers were publicly insured, and 15.1 percent were privately insured. The remainder were either in the military or uninsured (CASEN, 2015). We limit our analysis to private and public enrollees aged 25 to 64.

Private insurers offer individual contracts in a regulated market. Insurers must set premiums based exclusively on age groups, gender, and the number of dependents. We refer to the groups defined by these community pricing rules as population segments throughout the paper. Insurers may offer a variety of plans, each required to feature a two-tier network: The base tier determines the default coverage (i.e., one minus the coinsurance rate) for inpatient hospital care, while the preferential tier includes a set of hospitals where coverage is more generous.[6] Insurers must disclose plan cost-sharing and tiers to potential enrollees. Private plan coverage must be at least as generous as the public insurer's at private hospitals, which is roughly 50 percent. Insurers must provide access to all large general acute care hospitals, even if at varying cost-sharing. This regulation implies that when an insurer and hospital reach a disagreement in negotiations, insurers might face penalties associated with the loss of access. Moreover, there are no out-of-pocket maximums and negligible deductibles, implying enrollees remain exposed to hospital prices regardless of spending and that cost-sharing is determined almost exclusively by plan coverage. Finally, the regulation also includes minimum solvency requirements for insurers, and other features we discuss in Appendix A.2.

We focus on the five insurers available to all workers, which account for 96 percent of the private market. We denote these insurers by m_a–m_e. The remainder of the private market is served by seven *closed* insurers associated with a few large companies that only enroll their employees.

The hospital market consists of a mix of public and private hospitals. The public network is broader than the private one, with nearly twice as many hospitals (Clínicas de

[5]Mandatory contributions are capped. In 2015, approximately only the first $2,000 of monthly earnings were subject to mandatory contribution.

[6]Single-tiered plans resemble PPOs in the U.S., offering broad access. Preferential networks resemble HMOs that offer in-network (preferential) and out-of-network (non-preferential) coverage. For comparison, only 6 percent of spending occurs out-of-network in the U.S. (Song *et al.*, 2020), while 34 percent occurs at non-preferential hospitals in Chile.

Figure 2: Market structure

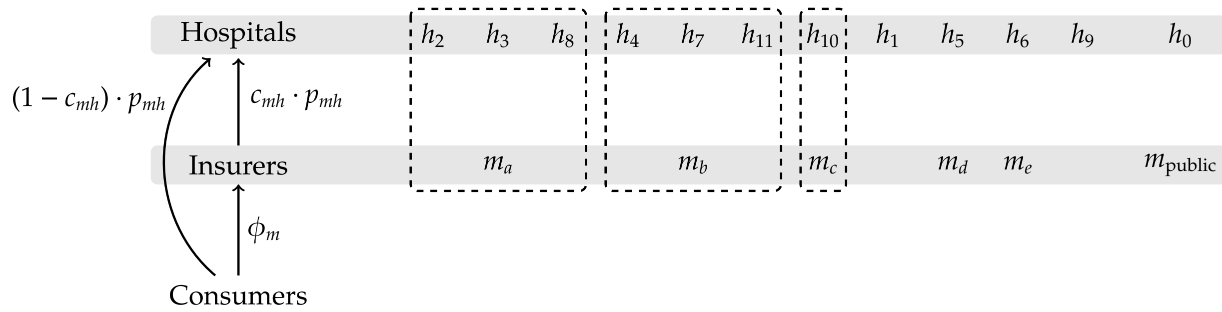

Notes: This figure displays the market structure in our settings. Downstream consumers pay premiums ϕ_m to insurers for a plan with coverage rate c_{mh}. Insurers negotiate over prices p_{mh} with hospitals. Dashed lines indicate VI hospitals and insurers.

Chile, 2012). Public hospitals provide care at highly subsidized and regulated prices, while private hospitals offer higher quality at a higher price. Differences in access, funding, and demand have led public hospitals to have longer wait times for procedures. As a result, private insurance enrollees utilize primarily private hospitals, accounting for 97 percent of all their spending (Galetovic and Sanhueza, 2013). There are significant differences among private hospitals in quality, specialization, and location.

We focus our analysis on inpatient care provided by general acute care hospitals in Santiago. This is the largest healthcare market in the country, accounting for more than one-third of private hospitals and around one-half of the capacity (Galetovic and Sanhueza, 2013). Inpatient care accounts for most medical spending and comprises fewer players, exacerbating the strategic concerns associated with VI. We limit our attention to the 11 leading private hospitals, which receive 74 percent of admissions in the market, the remainder captured by a large set of small hospitals. We denote these hospitals as h_1–h_{11} and h_0. Among these hospitals, h_1 and h_6 are broadly accepted as the highest quality. Given the demand these hospitals command, we refer to them as star hospitals (Ho, 2009).

VI is widespread, accounting for 48 percent of private hospital capacity (Galetovic and Sanhueza, 2013).[7] As illustrated in Figure 2, insurer m_a is integrated with hospitals h_2, h_3, and h_8, while insurer m_b is integrated with h_4, h_7, and h_{11}. Insurer m_c was integrated with h_{10} during the first year of our sample. These VI firms formed long before our period of study, and there were no horizontal mergers, entries, or exits during our sample that might confound the effects of VI. Importantly, VI hospitals accept patients from all insurers. Finally, note that star hospitals are not VI.

[7]Chilean law forbids insurers from owning hospitals. However, it does not forbid holding companies from owning both insurers and hospitals. We define VI firms as those for which the holding owns more than 50 percent of the hospital and 98 percent of the insurer, according to Copetta (2013).

3.2 Data

We use administrative data from the private insurance market regulator (*Superintendencia de Salud*) on all private plan enrollment choices and insurance claims for 2013–2016. These include 3,946,900 enrollment decisions linked to 773,264 inpatient admission events. We observe each plan's premiums, cost-sharing rules, and networks, including their preferential providers. We provide additional information about our data, our sample construction, and its connection with the regulatory environment in Appendix A.[8]

The data include detailed information on private enrollees and their inpatient claims. We observe members' age, gender, income, employment status, neighborhood of residence, and similar details about their dependents. The claims data includes the total hospital bill per line item, the amount paid by the insurer, and the consumer's out-of-pocket share. We observe detailed diagnoses, service codes, and associated ICD-10 codes. We use ICD-10 chapter codes to classify each medical event into one of 16 diagnosis groups.

To measure enrollment in the public plan, we use data from the CASEN survey on insurance enrollment (CASEN, 2015). We use the waves of 2013, 2015, and 2017 to compute the yearly enrollment share of FONASA for each quartile of the income distribution, age, gender, and dependents, linearly interpolating for the gap years. We also collect the public insurer's list prices paid to each public hospital and the cost-sharing and premium subsidy rules by income group.

Finally, we do not observe the underlying contracts negotiated between insurers and hospitals. We follow the literature and estimate negotiated prices based on an approach that rationalizes observed prices as the product of a negotiated price index and a resource intensity weight that scales payments according to diagnosis and patient characteristics (Gowrisankaran *et al.*, 2015; Ho and Lee, 2017; Cooper *et al.*, 2018). We scale our price index to reflect the price of an average delivery for a woman aged between 25 and 40. We treat these objects as data throughout the paper. For details, see Appendix A.1.3.

The average private plan policyholder in our sample is 40 years old, has a monthly income of $1,631, and 0.81 dependents. For comparison, the median income in the country was around $540 (CASEN, 2015). Most policyholders do not cover dependents, with 34 percent being single males and 24 percent single females. The average policyholder pays $173 in premiums per month, with substantial variation across plans and insurers. While insurers offer several plans, enrollment is skewed, and 123 plans account for 90

[8]We group plans according to financial characteristics, as insurers duplicate plans under different codes to circumvent regulation (Atal, 2019; Dias, 2022). Appendix A.1.1 provides further detail.

percent of enrollment. As much as 88 percent of plans offer a preferential tier, with an average preferential coverage of 77 percent and an average base tier coverage of 60 percent. The lower end of base coverage often falls below the 50 percent minimum requirement, suggesting partial enforcement of coverage regulation.[9]

The average admission in the main hospitals in our analysis has a price of $4,610, more than twice that of the outside option. Patients pay 24 percent of the bill, and the insurer pays the remainder. Moreover, 64 percent of admissions are at preferential providers. VI hospitals receive 61 percent of their admissions from their VI insurer and are a preferential provider for 88 percent of those admissions. In contrast, non-VI hospitals only receive 39 percent of admissions from any VI insurer. Finally, while the average hospital is 9.2 miles from the patient's residence, the typical admission occurs within 7.5 miles.[10] We display summary statistics in Appendix Table A.1. Appendix A.3 further describes the market structure and the interaction between insurers and hospitals.

4 Descriptive Evidence

4.1 Differences between VI and non-VI Firms

To describe how VI shapes market outcomes, we leverage the variation in VI and detailed data available in our setting. We explore two dimensions of variation. First, we compare how outcomes within a hospital differ across insurers depending on whether the insurer is vertically integrated. This cross-sectional strategy exploits all the available variation in VI. Second, we exploit variation in VI within insurer-hospital over time, which is generated by insurer m_c selling its stake in hospital h_{10} in 2014. While this event directly impacts a small share of the population, it has the advantage of providing a cleaner quasi-experiment. In both cases, we attempt to isolate the role of VI by including rich controls and fixed effects.

We begin by exploring how VI affects insurers' incentives when designing hospital networks. In particular, we explore its effect on coverage and tiering choices by estimating:

$$y_{jht} = \beta \mathrm{VI}_{m(j)ht} + \eta_{m(j)} + \zeta_h + \rho_t + \varepsilon_{jht} \tag{1}$$

where y_{jht} is a network attribute of plan j at hospital h in year t, and $\mathrm{VI}_{m(j)ht}$ indicates whether $m(j)$, the insurer offering plan j, is VI with hospital h. The regression includes insurer and hospital fixed effects to account for persistent differences in plan generosity,

[9] We discuss additional evidence on partial regulatory enforcement in Appendix A.2.

[10] Appendix Figure A.1 displays hospital locations and the spatial distribution of policyholders.

Table 1: Vertical integration and market outcomes

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Plan design		**Admission outcomes**				**Hospital outcomes**	
	Preferential hospital	Coverage rate	log cost proxy	log # services	Re-admission	log price	Share of admissions	log revenue
A. Cross-sectional								
VI	0.336	5.843	0.029	-0.034	0.013	-0.137	0.382	1.875
	(0.049)	(1.009)	(0.026)	(0.022)	(0.006)	(0.022)	(0.025)	(0.172)
R^2	0.150	0.240	0.208	0.612	0.050	0.692	0.555	0.788
B. De-integration event								
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)
R^2	0.212	0.254	0.212	0.613	0.059	0.694	0.988	0.976
N	15,741	15,741	567,752	567,752	204,223	567,752	264	264
Mean non-VI	0.137	63.101	2.534	16.039	0.081	5.269	0.128	8,865.659
Insurer & Hospital FE	A	A	A	A	A	A	A	A
Insurer × Hospital FE	B	B	B	B	B	B	B	B
Year FE	A/B	A/B	N	N	N	N	A/B	A/B
Interacted FE	N	N	A/B	A/B	A/B	A/B	N	N
Plan FE	N	N	A/B	A/B	A/B	A/B	N	N
Cost proxy	N	N	N	A/B	A/B	A/B	N	N
Controls	N	N	A/B	A/B	A/B	A/B	N	N
Observation	plan-hospital-year		admission				insurer-hospital-year	

Notes: Panel A shows estimates from equations (1), (2) and (3). Panel B shows estimates from the same equations, but adding hospital-insurer FEs to focus on the de-integration event between insurer m_c and hospital h_{10}. The unit of observation is reported in the bottom panel. Columns (3)–(6) include the following additional controls: diagnosis fixed effects, patient age, gender, policyholder income and employment status, and county fixed effects. Columns (4)–(6) also include public system admission prices 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, 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.

and year fixed effects to account for time trends in plan design. The results in columns (1) and (2) in Table 1-A suggest that vertical incentives play a relevant role in shaping plan design. VI hospitals are 33.6 percentage points more likely to be preferential in their integrated insurer's plans, and the coverage for such hospitals is 5.8 percentage points higher. Table 1-B shows similar results for the specification based on variation within insurer-hospital, which enriches equation (1) with insurer-hospital fixed effects. To the extent patients are responsive to hospital prices, these results suggest VI insurers skew coverage toward their hospitals to steer demand their way. We note that the effect on preferential tiering is limited not because VI insurers fail to tier their own hospitals, but rather because non-VI insurers also offer some preferential access to the same hospitals. This stems from differentiation: certain VI hospitals hold specific value for certain demographics and locations. Hence, insurers that wish to compete in those markets must offer plans that provide preferential access to VI hospitals, regardless of their integration.

VI firms may also differ from non-VI firms in their hospital costs, quality, and prices.

To study these margins, we leverage variation in admission outcomes within VI hospitals across patients insured by their integrated insurer and by rivals. We estimate:

$$y_{idjht} = \beta \text{VI}_{m(j)ht} + \boldsymbol{x}'_{ih}\gamma + \tau_{dh} + \rho_{dt} + \zeta_{ht} + \delta_j + \varepsilon_{idjht} \tag{2}$$

where y_{idjht} is an outcome for admission i for diagnosis group d under plan j in hospital h, and τ_{dh}, ρ_{dt}, ζ_{ht}, and δ_j are diagnosis-hospital, diagnosis-year, hospital-year, and plan fixed effects, respectively. To account for differences in complexity and costs across patients, $\boldsymbol{x}_{ih}$ includes gender, age, income, employment status, number of dependents, and neighborhood of residence. Using data on the services provided to the patient and the public system's list prices, we construct a cost proxy by computing the predicted total public hospital price. We interact this cost proxy with hospital dummies to account for differential cost pass-through and include these variables as additional controls. The parameter of interest is β, which is identified from differences in the admission outcomes of similar patients within a hospital-diagnosis-year, comparing those who are insured with the hospital's VI insurer to those who are not.[11]

A proposed advantage of VI is an increased scope for insurers to control hospital spending by improving care coordination and aligning incentives (Johnson *et al.*, 2017). For example, VI insurers might reduce information frictions by coordinating electronic health records, making labs and reports more easily accessible within the VI system (Dix *et al.*, 2024). We explore this by estimating equation (2) using our cost proxy and the number of services provided within an admission as dependent variables. This analysis informs whether patients with similar diagnoses, characteristics, and complexity are treated differently by the same hospital, depending on their insurance. Evidence that patients from VI insurers are treated at a lower cost or receive fewer services than comparable patients from rivals would suggest that VI induces cost-control or treatment changes. Columns (3) and (4) of Table 1-A show no significant differences in these outcomes. This lack of effects, however, might stem from VI changing these outcomes for all patients regardless of integration status. To explore this, we turn attention to the de-integration event and estimate equation (2) adding an insurer-hospital fixed effect. The results are shown in columns (3) and (4) of Table 1-B and also point to null effects. In addition, we study whether VI is associated with treatment quality by estimating the same regression using a 30-day readmission indicator as the dependent variable.[12] Column (5) shows no evidence

[11] We show that VI and non-VI insurers' enrollees are mostly balanced in observables in Appendix A.3.5.

[12] We follow the Centers for Medicaid and Medicare Services' methodology and code readmissions to inpatient care within 30 days of discharge from the original admission, focusing on diagnoses less likely to

of quality improvements in either the cross-sectional or de-integration specification.

Taken together, these results suggest that, in our setting, VI might not be producing meaningful cost efficiencies, treatment choice changes, or quality improvements.[13] These results are in line with recent evidence on the impacts of integrated care on efficiency and quality that has found mixed results (Johnson *et al.*, 2017).

Nevertheless, hospital prices faced by patients do vary depending on their insurer. Column (6) shows that the full admission price at a VI hospital is 13 percent lower for patients enrolled with the insurer integrated with the hospital. The evidence from the de-integration event in Table 1-B points in the same direction, showing that prices are directly affected by integration. This pattern is consistent with the elimination of double marginalization, lowering prices within VI firms, or with VI hospitals' incentive to negotiate higher prices with rival insurers. To quantify whether these price differences are effective at steering demand towards their hospitals, we estimate:

$$y_{mht} = \beta \mathrm{VI}_{mht} + \eta_m + \zeta_h + \rho_t + \varepsilon_{mht} \tag{3}$$

using admission flows from insurer m to hospital h as the dependent variable. Column (7) in Table 1-A shows that the share of admissions VI hospitals receive from their integrated insurer is 38 percentage points larger than from rival insurers. Column (8) shows that this demand pattern implies that, despite lower prices, VI hospitals receive around five times as much revenue from their integrated insurer as from rivals. Table 1-B shows a similar pattern of results when focusing on the de-integration event.

4.2 Evidence from Switchers to VI Insurers

To further quantify how effective VI insurers are at steering demand toward integrated partners, we study hospital choices and spending of enrollees who switch to VI insurers. Whereas switchers' hospital choice sets are constant over time, whether a hospital is VI with their insurer may change upon switching. Thus, switchers experience changes in out-of-pocket prices but not in access.

We estimate the following event study regression:

$$y_{iht} = \sum_{\tau} \beta_\tau D_i \mathrm{VI}_{hm(i\tau)} + \sum_{\tau} \gamma_\tau D_i (1 - \mathrm{VI}_{hm(i\tau)}) + \alpha_i + \delta_{ht} + \varepsilon_{iht} \tag{4}$$

require readmission as part of their treatment (CMS, 2024). This increases the likelihood of readmissions capturing lower hospital quality. We include circulatory, infections, pregnancy, and respiratory admissions.

[13] We find a similar pattern of results when examining discretionary procedures in Appendix A.4.

Figure 3: Vertical integration, hospital choices, and expenditure

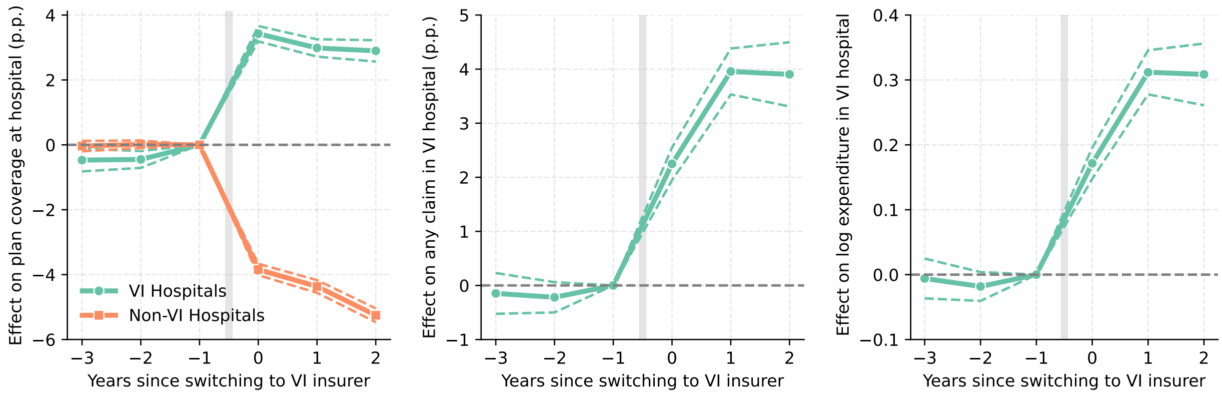

(a) Plan coverage **(b)** Any claims in VI hospital **(c)** Expenditure in VI hospital

Notes: This figure displays estimates from equation (4). 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 percent confidence intervals. The dependent variable in Figure 3c is $\log(1 + y)$ to accommodate zeros, but the results are similar when using expenditure in levels.

where y_{iht} is an outcome for patient i in hospital h at year t. The main explanatory variables are interactions between the indicator D_i for individual i ever switching to a VI insurer and the indicators $\text{VI}_{hm(i\tau)}$ for hospital h being VI with the insurer in which i is enrolled in τ years after switching. The coefficients β_τ measure the effect of changing the integrated status of a hospital relative to non-switchers, while γ_τ captures potential effects on non-VI hospitals. We include individual fixed effects α_i to control for persistent differences across patients (e.g., permanent differences in health) and hospital-time fixed effects δ_{ht} to control for differences in outcomes across hospitals and time that are constant across patients (e.g., seasonality in health shocks, quality differences). For ease of interpretation, we restrict the sample to enrollees who either never switch or switch to a VI insurer.

Enrollees switching to a VI insurer experience a change in plan coverage across hospitals. Figure 3a shows that switchers face 3 percentage points higher coverage at hospitals integrated with their new insurer and 4 percentage points lower coverage at rival hospitals. Likely as a result of the appeal of higher coverage, switchers are more likely to choose hospitals integrated with their insurer. Figure 3b shows that when an enrollee switches insurers, the probability of choosing a hospital VI with the new insurer increases by 4 percentage points relative to before switching, more than doubling the baseline rate of 3.5 percent. Moreover, Figure 3c shows that expenditure in hospitals integrated with the new insurer increases by around 30 percent after the enrollee switches.[14]

[14] A concern is that patients may switch to a VI insurer to gain better access to integrated hospitals. Our estimates show well-behaved trends leading to the switch and sharp impacts upon switching, suggesting that pre-existing health conditions or relationships with the hospital do not drive the results. However, we

4.3 Discussion

Overall, these results suggest VI firms steer enrollees to their hospitals using a combination of plan design and hospital prices. While these findings point to a strong effect of VI on market outcomes, they are insufficient to determine its equilibrium impact. For instance, price gaps could arise from either reduced double marginalization or cost efficiencies within VI firms or from market power and foreclosure of rivals. Our model builds on this descriptive evidence to parse the relative effectiveness of plan design and hospital prices as steering instruments and quantify their implications.

5 Model

We model the market as the subgame perfect equilibrium of a four-stage game. First, insurers design their plans, choosing coverage and networks. Second, hospitals and insurers negotiate over hospital prices, and insurers set premiums simultaneously. Third, households enroll in insurance plans. Fourth, health risks are realized, and enrollees choose which hospital to seek care at. We describe the model in reverse order. At the end of the section, we present a discussion of the model's assumptions and properties. We leave the specialization of some aspects to the Chilean setting for the next section.

5.1 Demand for Healthcare

When consumer $i \in \bar{\mathcal{I}}$, enrolled in a plan $j \in \mathcal{J}$ of insurer $m(j)$, falls ill with a diagnosis $d \in \mathcal{D}$ in year t, they choose a hospital $h \in \mathcal{H}$ to maximize an indirect utility given by:

$$u^H_{ihdt|j} = \alpha^H_i (1 - c_{jht}) p_{ijhdt} + \beta^H \text{distance}_{ih} + \chi^H_{hdt} + \xi^H_{m(j)hdt} + \epsilon^H_{ihdt|j} \tag{5}$$

where the first term is the disutility from out-of-pocket spending, determined by coverage c_{hjt} and price p_{ijhdt}, the second term is the disutility from travel, χ^H_{hdt} captures hospital quality for diagnosis d (e.g., due to the presence of specialists or specialized technology), and $\xi^H_{m(j)hdt}$ captures heterogeneity in access to quality across insurers. By allowing $\xi^H_{m(j)hdt}$ to vary across insurers and time, we allow VI to affect demand beyond systematic differences in prices, coverage, or quality—we isolate this channel for impacts of VI when estimating the model below. This accounts for features unobserved to the econometrician that might steer demand to integrated hospitals (e.g., marketing efforts, ease-of-pay benefits, or

cannot rule out that contemporaneous health shocks induce switching. For robustness, in Appendix Figure A.2, we repeat the analysis in a subsample of enrollees who move across neighborhoods. Moves may induce insurance switching, but are unlikely to be correlated with health shocks. The results are similar.

appointment assistance). Finally, $\epsilon^H_{ihdt|j}$ is an iid idiosyncratic shock distributed T1EV.

Patients may seek care outside the main set of hospitals. This care is offered at a price p_{0dt} and at the base plan coverage, $\underline{c}_{jt}$. We normalize the distance to hospitals for each consumer relative to the distance to their closest outside option provider, such that the indirect utility from the choice of the outside option is:

$$u^H_{i0dt|j} = \alpha^H_i(1 - \underline{c}_{jt})p_{0dt} + \eta^H_{m(j)dt} + \epsilon^H_{i0dt|j}$$

where the first term captures the disutility from out-of-pocket spending, and $\eta^H_{m(j)dt}$ relative preferences for the outside option. Thus, consumer i's probability of choosing hospital h for diagnosis d prior to the realization of $\boldsymbol{\epsilon}^H$ is $D^H_{ihdt|j} = \exp(\delta^H_{ihdt|j})/\sum_{h'\in\mathcal{H}\cup\{0\}}\exp(\delta^H_{ih'dt|j})$, where $\delta^H_{ihdt|j} = u^H_{ihdt|j} - \epsilon^H_{ihdt|j}$.

5.2 Demand for Insurance

In the third stage, consumers purchase insurance. Each policyholder i represents a household $\mathcal{F}_i \subset \bar{\mathcal{I}}$, and must choose an insurance plan before household members' medical needs and idiosyncratic preferences (ϵ^H) are realized.[15] Policyholder i chooses plan j if it maximizes the household's indirect utility:

$$u^M_{ijt} = \alpha^M_i|\mathcal{F}_i|\phi_{jt} + \beta^M \sum_{i'\in\mathcal{F}_i} WTP^H_{i'jt}(\boldsymbol{c}_{jt}, \boldsymbol{p}_{jt}) + \chi^M_{im(j)} + \xi^M_{m(j)t} + \epsilon^M_{ijt} \tag{6}$$

where the first term captures the disutility from premiums ϕ_{jt}, which scales with household size $|\mathcal{F}_i|$, and the second is the utility derived from the expected network surplus of plan j, often called network willingness-to-pay in the literature (Capps *et al.*, 2003; Ho and Lee, 2017). In particular, letting r_{id} be the probability that consumer i falls ill with condition d, the expected network surplus the consumer derives from plan j is $WTP^H_{ijt}(\boldsymbol{c}_{jt}, \boldsymbol{p}_{jt}) = \frac{1}{|\alpha^H_i|}\sum_{d\in\mathcal{D}} r_{id} \ln \sum_{h\in\mathcal{H}\cup\{0\}} \exp(\delta^H_{ihdt|j}(\boldsymbol{c}_{jt}, \boldsymbol{p}_{jt}))$. In addition, $\chi^M_{im(j)}$ and $\xi_{m(j)t}$ capture persistent and time-varying unobserved preference shocks for insurers, respectively. Finally, ϵ^M_{ijt} is an iid idiosyncratic preference shock distributed T1EV.

Households can opt for public coverage instead of private insurance. This outside option is offered at a premium ϕ_{0it} and grants access to a set of hospitals with a network surplus independent from the prices and coverages of private insurers. We normalize the network surplus of the outside option to zero, which sets the level for the unobserved preference shocks for private plans, $\chi^M_{im} + \xi^M_{mt}$. The indirect utility of the outside option is

[15]Formally, letting $\mathcal{I} \subset \bar{\mathcal{I}}$ denote policyholders, $\{\mathcal{F}_i\}_{i\in\mathcal{I}}$ is an exhaustive partition of $\bar{\mathcal{I}}$ such that $i \in \mathcal{F}_i$.

thus $u_{i0t} = \alpha_i^M \phi_{0it} + \epsilon_{i0t}^M$.

The choice set for each policyholder consists of the subset of plans available in their market segment and the public outside option. Thus, each policyholder i that belongs to a segment $s \in \mathcal{S}$ has access to plans $\mathcal{J}_{st} \cup \{0\}$. The choice probability of plan $j \in \mathcal{J}_{st}$ by i is $D_{ijt}^M = \exp(\delta_{ijt}^M) / \sum_{j' \in \mathcal{J}_{st} \cup \{0\}} \exp(\delta_{ij't}^M)$, where $\delta_{ijt}^M = u_{ijt}^M - \epsilon_{ijt}^M$.

5.3 Insurance Premium Competition

In the second stage, each insurer $m \in \mathcal{M}$ sets premiums to maximize profits, taking expectation over consumers' unrealized health risk and preference shocks $(\boldsymbol{\epsilon}^M, \boldsymbol{\epsilon}^H)$. Given coverages $\boldsymbol{c}_t$, premiums $\boldsymbol{\phi}_t$, and prices $\boldsymbol{p}_t$, insurer m's expected profits are:

$$\pi_{mt}^M(\boldsymbol{c}_t, \boldsymbol{\phi}_t, \boldsymbol{p}_t) = \sum_{j \in \mathcal{J}_{mt}} \sum_{i \in \mathcal{I}} D_{ijt}^M(\boldsymbol{c}_t, \boldsymbol{\phi}_t, \boldsymbol{p}_t)(|\mathcal{F}_i| \phi_{jt} - \sum_{i' \in \mathcal{F}_i} \sum_{d \in \mathcal{D}} \sum_{h \in \mathcal{H} \cup \{0\}} r_{i'd} D_{i'hdt|j}^H(\boldsymbol{c}_{jt}, \boldsymbol{p}_{m(j)t}) c_{jht} p_{i'jhdt} - \eta_{jt}) \quad (7)$$

where $\mathcal{J}_{mt}$ is the set of plans that insurer m offers in year t and $\mathcal{I}$ is the set of policyholders. For a given household, insurer m earns revenue from collected premiums and faces costs equal to its expected share of payments plus an administrative burden η_{jt}.

VI insurers also consider the impact of their premiums on their integrated hospitals' profits. The expected profit of hospital h is:

$$\pi_{ht}^H(\boldsymbol{c}_t, \boldsymbol{\phi}_t, \boldsymbol{p}_t) = \sum_{m \in \mathcal{M}} \sum_{j \in \mathcal{J}_{mt}} \sum_{i \in \mathcal{I}} D_{ijt}^M(\boldsymbol{c}_t, \boldsymbol{\phi}_t, \boldsymbol{p}_t) \sum_{i' \in \mathcal{F}_i} \sum_{d \in \mathcal{D}} r_{i'd} D_{i'hdt|j}^H(\boldsymbol{c}_{jt}, \boldsymbol{p}_{mt})(p_{i'jhdt} - \omega_{i'd} k_{hmt}) \quad (8)$$

which accounts for h's expected demand from all plans and the profit from each admission. The cost of each admission combines the cost per unit of resources k_{hmt} with the expected resources required to treat condition d given the characteristics of consumer i', captured by weights $\omega_{i'd}$. Unit costs may vary across insurers within a hospital, capturing level differences in resource utilization or claim processing across insurers.

Given the definitions of insurer and hospital profits, we define the general objective of any firm f using a vector of weights $\boldsymbol{\theta}^f \in \Delta^{|\mathcal{M}|+|\mathcal{H}|}$ and the payoff function:

$$\tilde{\pi}_t(\boldsymbol{c}_t, \boldsymbol{\phi}_t, \boldsymbol{p}_t; \boldsymbol{\theta}_t^f) = \sum_{m \in \mathcal{M}} \theta_{mt}^f \pi_{mt}^M(\boldsymbol{c}_t, \boldsymbol{\phi}_t, \boldsymbol{p}_t) + \sum_{h \in \mathcal{H}} \theta_{ht}^f \pi_{ht}^H(\boldsymbol{c}_t, \boldsymbol{\phi}_t, \boldsymbol{p}_t) \quad (9)$$

Using this structure, the objective of a VI insurer m is defined by a vector $\boldsymbol{\theta}_t^m$, which is zero for all insurers different from m and all hospitals not integrated with m. Thus, m's premiums maximize $\tilde{\pi}_t(\boldsymbol{c}_t, \boldsymbol{\phi}_t, \boldsymbol{p}_t; \boldsymbol{\theta}_t^m)$. This formulation of the objective accommodates potential imbalances within the firm caused by differential regulatory scrutiny between the hospital and insurance market, tax credits, or partial internalization of integrated profits, as in Crawford *et al.* (2018). For example, if m and h are integrated, and $\theta^m_{ht} = \theta^m_{mt} = 1/2$, then m's premiums will maximize their joint profits. If instead, $\theta^m_{mt} = 2/3$ and $\theta^m_{ht} = 1/3$, then the insurer will imperfectly internalize profits at its integrated hospitals, placing half as much weight on upstream profits as on downstream profits.

We use equation (9) to describe the objectives of all firms. For a non-VI insurer m', it is $\tilde{\pi}_t(\boldsymbol{c}_t, \boldsymbol{\phi}_t, \boldsymbol{p}_t; \boldsymbol{\iota}^{m'})$, where $\boldsymbol{\iota}^{m'}$ is a vector that has zeros in all components except for the m'-th. For a VI hospital h, it is $\tilde{\pi}_t(\boldsymbol{c}_t, \boldsymbol{\phi}_t, \boldsymbol{p}_t; \boldsymbol{\theta}^h_t)$ where θ^h_t is a vector of weights that is zero for any insurer not integrated with h and any hospital that is not part of h's system. The objective of a non-VI hospital is defined analogously, with equal weight placed among horizontally integrated hospitals. When the integration status of a hospital or insurer is evident from the context, we use the shorthand $\tilde{\pi}_{mt}$ or $\tilde{\pi}_{ht}$ to denote their objective function.

5.4 Hospital Price Negotiations

Hospital prices are determined simultaneously to premiums by bilateral negotiations between insurers and hospitals following a Nash-in-Nash protocol (Crawford and Yurukoglu, 2012; Collard-Wexler *et al.*, 2019). Firms negotiate over an index price per unit of resources, p_{mht}, such that the payment received by hospital h for the treatment of condition d in year t for a consumer i enrolled in plan j from insurer m is $p_{ijhdt} = \omega_{id} p_{hmt}$.

Hospital prices are determined as:

$$p_{hmt} = \arg\max_{p_{hmt}} \quad (\tilde{\pi}_{mt}(p_{hmt}) - \tilde{\pi}_{mt\setminus h})^{\tau_{hmt}} (\tilde{\pi}_{ht}(p_{hmt}) - \tilde{\pi}_{ht\setminus m})^{1-\tau_{hmt}} \tag{10}$$

such that they maximize the weighted geometric mean of gains from trade, where τ_{hmt} is the insurer's relative bargaining weight. If an insurer and hospital agree on a price p_{hmt}, they obtain profits $\tilde{\pi}_{mt}(p_{hmt})$ and $\tilde{\pi}_{ht}(p_{hmt})$, respectively. If the firms do not reach an agreement, m removes h (and any of its integrated partners) from its networks, yielding the insurer a payoff $\tilde{\pi}_{mt\setminus h}$ and the hospital $\tilde{\pi}_{ht\setminus m}$. Assuming rational expectations and passive beliefs about rival prices and premiums (Horn and Wolinsky, 1988), disagreement payoffs $\tilde{\pi}_{mt\setminus h}$ and $\tilde{\pi}_{ht\setminus m}$ can be computed from equations (7) and (8) by removing hospital h (and its partners) from the choice sets of all of m's enrollees, recomputing hospital choice probabilities, networks surplus values, and plan enrollment probabilities. We assume that any hospital-insurer pair that negotiates a price does so under gains from trade.[16]

[16]Formally, for every pair (h, m) that negotiates a price in period t we assume that $\Delta_{hm}\tilde{\pi}_{mt} \equiv \tilde{\pi}_{mt} - \tilde{\pi}_{mt\setminus h} > 0$, $\Delta_{hm}\tilde{\pi}_{ht} \equiv \tilde{\pi}_{ht} - \tilde{\pi}_{ht\setminus m} > 0$. If these gains from trade conditions are not met, then bargaining is not individually rational for at least one of the players.

We assume that incentives in negotiations between VI firms are aligned, such that both m and h evaluate their objective using common weights $\theta^m = \theta^h$. In this case, the Nash-in-Nash protocol coincides with joint profit maximization.

5.5 Plan Design

In the first stage, insurers design their plan menus. Each insurer m is endowed with plans $\mathcal{J}_{mt}$ for which they must decide the set of preferential hospitals $\bar{\mathcal{H}}_{jt} \subset \mathcal{H}$, the base coverage $\underline{c}_{jt} \in [0,1]$, and the preferential coverage $\bar{c}_{jt} \in [\underline{c}_{jt}, 1]$.[17] Denoting the set of all tiered plan designs as $\mathcal{C} = \{\boldsymbol{c} \in [0,1]^{|\mathcal{H}|} \mid \sum_h (\max_{h'} c_{h'} - c_h)(c_h - \min_{h'} c_{h'}) = 0\}$, each insurer m solves:

$$\max_{\{\boldsymbol{c}_j \in \mathcal{C}\}_{j \in \mathcal{J}_{mt}}} \tilde{\pi}_{mt}(\boldsymbol{c}, \boldsymbol{\phi}^*(\boldsymbol{c}), \boldsymbol{p}^*(\boldsymbol{c}); \boldsymbol{\theta}^m) - \sum_{j \in \mathcal{J}_{mt}} N_{s(j)t} K_{mt}(\boldsymbol{c}_j) \tag{11}$$

where the first term is the firm's equilibrium profit objective and the second is its design cost, which is additive across plans and scales with the size of the market segment in which j is offered $N_{s(j)t}$. Design costs might include true costs, such as reinsurance premiums, or virtual costs, such as those associated with the Lagrangian of unobserved regulatory constraints. We note that insurers design plans under uncertainty as consumers' medical needs and preference shocks $(\boldsymbol{\epsilon}^H, \boldsymbol{\epsilon}^M)$ have not been realized at this stage.

5.5.1 Solving the Plan Design Problem. The plan design problem is a large combinatorial problem that requires further development to be solved at scale.[18] We start by noting that any tiered design $\boldsymbol{c}_j$ can be decomposed into its base and preferential coverages $(\underline{c}_j, \bar{c}_j)$ and a vector of weights $w_j \in [0,1]^{|\mathcal{H}|}$ indicating the relative position of hospitals in its tiers, such that $c_{jh} = \bar{c}_j w_{jh} + (1 - w_{jh})\underline{c}_j$. In this form, the set of tiered plans can be written as $\mathcal{C} = \{(\underline{c}, \bar{c}, w) \in [0,1]^{2+|\mathcal{H}|} \mid \underline{c} \leq \bar{c} \ \wedge w_h(1 - w_h) = 0 \ \forall h\}$. Note that, conditional on w_j, the problem of setting coverages is akin to that of setting prices, as coverage and prices appear jointly in hospital demand and insurance profits.[19] Thus, the complexity of the problem stems from the allocation of hospitals to tiers and not the choice of coverages.

[17]The outside option hospital is intentionally omitted from the potential set of preferential hospitals. The set of preferential hospitals is allowed to be empty.

[18]We model a single geographic market and a single price for each insurer-hospital-year. Thus, the design of plans in one segment affects all prices, forcing insurers to consider the impact of plan design jointly.

[19]Holding $\boldsymbol{w}$ and $\boldsymbol{p}$ fixed, the design problem is analogous to the equilibrium of a game in which insurers pay a predetermined fee per hospital admission and resource unit, and where insurers first determine hospitals prices at the plan-hospital level subject to linear constraints and then compete in premiums. The linear constraints correspond to prices satisfying the constraints imposed by the coverage tiers and being bounded above by the hospital fee. This would be the case, for example, if all hospitals were public.

Therefore, it is natural to define a proxy problem that relaxes tiering constraints. Letting $C^* = \{(\underline{c}, \bar{c}, w) \in [0,1]^{2+|\mathcal{H}|} \mid \underline{c} \leq \bar{c}\}$ be the set of all designs, λ a positive constant, and G a positive and increasing function, we define the regularized problem as:

$$\max_{\{\underline{c}_j, \bar{c}_j, \boldsymbol{w}_j \in C^*\}_{j \in \mathcal{J}_{mt}}} \tilde{\pi}_{mt}(\boldsymbol{c}, \boldsymbol{\phi}^*(\boldsymbol{c}), \boldsymbol{p}^*(\boldsymbol{c})) - \sum_{j \in \mathcal{J}_{mt}} N_{s(j)t} \hat{K}_{mt}(\boldsymbol{c}_j) - \lambda \sum_{j \in \mathcal{J}_{mt}} \sum_{h \in \mathcal{H}} G(w_{jh}(1 - w_{jh})) \tag{12}$$

Note that $\tilde{\pi}_{mt}$ is well defined on C^*, as demand and profits accept any coverage, regardless of tiering structure. In contrast, design costs might only be defined on tiered designs. Thus, we denote by $\hat{K}_{mt}$ a continuous design cost on C^*. We say that the regularized problem is a relaxation of the original problem if $\hat{K}_{mt}$ is a continuous extension of K_{mt} to C^*. The following result, proven in Appendix B.3, is key to our study of plan design.

Proposition 1. *A continuous extension of K_{mt} always exists and is bounded whenever K_{mt} is bounded. Moreover, let $\mathfrak{c}^*$ be the set of solutions to the design problem and $\tilde{\mathfrak{c}}(\lambda, G)$ be the set of solutions of its relaxation, then $\mathfrak{c}^* = \lim_{\lambda \to \infty} \tilde{\mathfrak{c}}(\lambda, G)$.*

The intuition behind this result is that the set of plan designs that emerges under strict tiering rules is the same as those under sufficiently heavy penalties for lack of tiering. As we can choose the structure of virtual penalties G arbitrarily, we can use it to "concavify" the problem by making G convex. Moreover, if G is differentiable, the regularized objective is differentiable whenever $\hat{K}$ is differentiable. Thus, each regularized problem is as complex as the problem of choosing coverage levels conditional on tiers.

To our knowledge, this approach is novel in the study of contract design. However, it builds on well-established results in combinatorial optimization. Proposition 1 is a continuous extension result for the plan design problem (Lucidi and Rinaldi, 2010). Murray and Ng (2010) provide algorithms for the sequence problem defined in the proposition, showing its remarkable performance in an array of canonical large-scale optimization problems. In Appendix C.5, we provide additional evidence showing how it benchmarks against solving the plan design problem by grid search. We note that the impact on time and precision is similar to the difference between solving a multiproduct pricing problem by grid search rather than relying on first-order optimality conditions.

We compute exact solutions to the plan design problem using proposition 1. Relative to alternative approaches, we expect this one to be well-suited to the plan design problem. First, the firm's objective $\tilde{\pi}$ is well-defined on the unconstrained set of designs C^*, thus requiring no manual extension. Second, when K_{mt} is unknown to the econometrician, it is indistinguishable from its continuous extension as it generates the same set of choices and payoffs on all admissible designs. Third, the regularized approach taps into the insurer's mechanism design problem. Intuitively, insurers use coverage differences across hospitals to steer consumers by making it incentive-compatible to follow the insurer's preferred demand allocation. To steer enrollees away from a subset of hospitals (e.g., non-VI hospitals), insurers will naturally assign them low coverage. Gains from multiple coverage levels might be present among desirable hospitals, but as coverage is bounded above by one and below by the incentive compatibility constraint, the value of such allocations will likely be small. Thus, a small penalty for lack of tiering should suffice for insurers to offer tiered networks. We confirm this intuition in our simulations, where the regularized solution converges to the optimum within a few increments of λ.

5.6 Equilibrium

An equilibrium of the game consists of plan coverages $\mathbf{c}$, prices $\mathbf{p}$, premiums ϕ, enrollment and hospital choices, and beliefs about risk and preference shocks such that: (i) the coverage of each insurer m maximizes the objective in equation (11); (ii) the premium vector of each insurer m maximizes $\tilde{\pi}_{mt}$; (iii) the price negotiated between each h and m, p_{hm} solves the Nash bargaining condition in equation (10); (iv) households enrollment choices maximize their indirect utility of choice in equation (6); (v) hospital admission choices maximize individual indirect utilities conditional on enrollment in equation (5); and (vi) all agents beliefs about population risk and preference shocks are rational.

There might be several such equilibria. For example, if consumers and hospitals perceive some insurers to be ex-ante (i.e., prior to plan design) identical in preferences and costs, then a permutation of one equilibrium plan design configuration will likely yield a different equivalent equilibrium. To address this, we adopt a strict refinement, focusing on the unique equilibrium achievable by best response iteration starting from the status quo. We see this equilibrium as the natural transition point for the market in counterfactuals. The refinement is strict, as such equilibrium requires the best response mapping to be locally contracting, which is not guaranteed. In exchange for this loss, the approach delivers a unique and coherent prediction.

5.7 Discussion

The model builds on the frameworks in Gowrisankaran *et al.* (2015) and Ho and Lee (2017) for healthcare competition and in Crawford *et al.* (2018) for vertical integration. By bridging these models and adding endogenous plan design, our model captures the mechanisms described in Section 2: self-preferencing and narrow networks through strategic plan design, changes in double marginalization, and raising rivals' cost incentives.

The incentives for skewing hospital networks emerge from the plan design objective in equation (11). If the weight placed by an insurer m on the profits of a hospital h increases, then, all else equal, insurer m will shift coverage to increase h's profits. This is achieved by making h preferential in more plans (i.e., self-preferencing), making h's rivals preferential in fewer plans (i.e, network narrowing), and increasing the wedge between preferential and base coverage in those plans. In turn, double marginalization arises from the Nash bargaining in equation (10) and its effect on premiums. As m and h's objective weights ($\boldsymbol{\theta}$) converge to be equal due to VI, the price they agree on will converge to their joint profit maximization value, which will eliminate double markup distortions on downstream premiums. This integration also affects their gains from trade with rivals. When h negotiates with a non-VI insurer $\hat{m}$ over a price $p_{h\hat{m}t}$, a marginal increase in price (or a disagreement) will tend to benefit m by reducing the network surplus of plans owned by $\hat{m}$, inducing enrollees to substitute towards m. Thus, VI increases the derivative of $\tilde{\pi}_{ht}$ with respect to $p_{h\hat{m}t}$ and the value of $\tilde{\pi}_{ht\setminus\hat{m}}$, all else equal. This generates the raising rivals' cost incentive and the potential for foreclosure. Finally, note that non-VI insurers can respond to such incentives by skewing their networks away from VI hospitals.

The model has some simplifying assumptions. First, following the literature on insurer-hospital competition (Handel and Ho, 2021), we abstract from hospital capacity, sales intermediaries, and inertia. To the extent capacity constraints translate to wait times and factor into χ^H_{hdt}, we assume those remain stable in counterfactuals. To our knowledge, private hospital capacity is rarely binding in our setting. Similarly, if insurers compete through sales agents, affecting plan demand through χ^M_{im}, we assume this influence remains unchanged in counterfactuals. Analogously, we do not seek to separate inertia from systematic preference heterogeneity. In addition, we focus on how VI shapes plan design, prices, and premiums, rather than risk protection. Thus, we do not model risk preferences. While VI might impact equilibrium risk protection, we consider these effects secondary to those on network structure, which may occur while holding risk protection constant. Finally, network tiering is decided unilaterally by the insurer in our setting. VI might have different competitive implications if tiering decisions can be negotiated.

6 Identification and Estimation

We now present the specialization of our model to the Chilean setting, its identification, and our estimation strategy. We discuss these elements following the four-stage structure of the model. Appendix C provides supporting evidence and proofs.

6.1 Preferences over Hospitals

The first step is to estimate hospital demand. The main parameters of interest are the coefficients on hospital out-of-pocket prices α_i^H, the parameter on distance to hospitals β^H, the hospital-diagnosis heterogeneity χ_{hdt}^H, and the hospital-insurer heterogeneity $\xi_{m(j)hdt}^H$. To specify the model for our setting, we define the outside option to private hospitals as a composite of many public or small private hospitals. Given the wide availability of public providers, we use the public system's list prices as the relevant outside option price. Since all neighborhoods have public providers, distance to the outside option is uniformly zero.

In terms of identification, the primary challenge stems from potential price endogeneity.[20] We address this with an instrument based on the prices negotiated by closed insurers with hospitals in our sample. As discussed in Section 3, closed insurers are isolated employer groups that have formed insurance companies to provide coverage to their employees. Consumers in our sample are not eligible for closed plan coverage. Moreover, closed insurer enrollees often live in different areas, are employed in different industries, and have sociodemographic profiles different from those of our population of interest. Thus, the prices negotiated by closed insurers reflect the preferences of a separate group of consumers but capture variation in hospital costs over time. These instruments are in the spirit of Hausman (1994), albeit for different firms in a segmented market. We use data on closed insurer claims to predict out-of-pocket prices for each admission in our hospital choice data. We use the coverage of the plan we are instrumenting, as the market timing implies insurers design plans before plan-specific hospital preferences are realized.

We estimate preferences for hospitals via maximum likelihood, leveraging individual-level hospital choice data. We use a control function approach to incorporate the instruments (Petrin and Train, 2010) and the inversion approach of Berry (1994) to recover our model's rich structure of unobserved preferences. In particular, we define the aggregate insurer-hospital-diagnosis-year effect $\bar{\xi}_{m(j)hdt}^H \equiv \chi_{hdt}^H + \xi_{m(j)hdt}^H - \eta_{m(j)dt}^H$, and $\bar{\boldsymbol{\xi}}^H$ as the vector of those effects, and maximize the data likelihood:

$$\max_{\{\alpha_i^H\}_{i\in\mathcal{I}},\beta^H,\bar{\boldsymbol{\xi}}^H} \sum_{i,h,d,t} y_{ihdt} \ln D_{ihdt|j}^H(\alpha_i^H, \beta_i^H, \bar{\boldsymbol{\xi}}^H) \quad \text{s.t} \quad D_{hdt|m}^{H*} = \sum_{i|j(i)\in\mathcal{J}_{mt}} D_{ihdt|j}^H(\alpha_i^H, \beta_i^H, \bar{\boldsymbol{\xi}}^H) \quad \forall h, d, t$$

where y_{ihdt} is a choice indicator, $D_{ihdt|j}^H$ the model implied individual choice probability, and

[20] Given the insurer-hospital-diagnosis-year fixed effect, price preferences are identified from variation in the likelihood of visiting a hospital for the same condition-insurer-hospital combination under different coverage levels. The source of endogeneity can be viewed as either an expected component of consumer idiosyncratic preferences or an unmodeled plan-specific residual preference.

$D^{H*}_{hdt|m}$ the observed market share of hospital h among enrollees of insurer m for condition d in year t.[21] We project the estimates of $\bar{\boldsymbol{\xi}}^H$ on its components, as:

$$\bar{\xi}^H_{m(j)dht} = \chi^H_{hdt} + \gamma^H \mathrm{VI}_{m(j)ht} - \eta^H_{m(j)dt} + \tilde{\xi}^H_{m(j)hdt}$$

where we have split $\xi^H_{m(j)hdt}$ into a VI-specific component and its residual variation, $\tilde{\xi}^H_{m(j)hdt}$. The de-integration event in our sample allows the separate identification of γ^H from the set of fixed effects included in the regression.

Consumers who are older, have higher incomes, or have dependents are less price elastic, and single women are less elastic than single men, as shown by Table 2-A. The model implies a median own-price demand elasticity of -0.79, which is larger than similar estimates from the U.S.[22] Price sensitivity is likely stronger in Chile since consumers can obtain price quotes from providers or insurers for planned inpatient care, and are always exposed to hospital prices on the margin due the structure of insurance plans. In our analysis, price elasticities primarily serve as a channel for network steering. If in other settings, consumers were instead responsive to variation in coverage across hospitals (e.g., moving a hospital out of network or from fixed copay to coinsurance), then demand steering effects similar to those we document for Chile are likely to emerge.

The estimates also show a substantial distaste for travel and meaningful heterogeneity across hospitals. A 25-year-old single male would be willing to pay nearly \$60 to reduce his travel distance by a mile, which implies significant preference heterogeneity for accessing different hospitals. Similarly, hospitals are meaningfully differentiated in diagnosis-specific quality.[23]

The estimates imply that VI induces additional demand between integrated partners equivalent to a \$1,510 reduction in out-of-pocket payments. For comparison, this is 77 and 95 percent of the quality gap between the best non-VI provider for pregnancy care and the best hospitals integrated with insurers m_a and m_b, respectively. Hence, VI insurers effectively steer consumers towards integrated hospitals using mechanisms beyond prices, plan coverage, or access to common quality. Given that the evidence presented in Section 4

[21] We use a random 30 percent sample of admissions to fit individual-level preferences and the full sample for share matching condition and estimating $\bar{\xi}^H$. The results are robust to changes in the sample size.

[22] For example, Aron-Dine *et al.* (2013) estimate an elasticity of -0.03, while Prager (2020) estimates -0.16.

[23] Appendix Figure A.3 shows that, conditional on price and distance, consumers prefer star hospitals for almost all conditions. Insurer m_a's integrated hospitals are of moderately low quality, while two of the three hospitals integrated with m_b are above the median in almost every condition and comparable to star hospitals in managing circulatory, infectious, and blood diseases. The best oncological care is provided at non-VI hospitals, and the public system has the highest quality for perinatal conditions.

Table 2: Estimates of consumer preferences

	(1)	(2)	(3)	(4)
	A. Healthcare		**B. Insurance**	
	Coef.	S.E.	Coef.	S.E.
A: Price (α_i^H) / B: Premium (α_i^M)				
× Age ∈ $[25, 40)$	-1.536	(0.010)	-29.382	(0.053)
× Age ∈ $[40, 55)$	-1.355	(0.011)	-26.999	(0.052)
× Age ∈ $[55, 65]$	-1.335	(0.011)	-27.789	(0.054)
× Female × Single	0.263	(0.010)	10.940	(0.051)
× Has dependents	0.223	(0.009)	16.097	(0.049)
× High income	0.247	(0.005)	7.711	(0.027)
Distance to hospital (β^H)	-0.089	(0.001)		
VI demand shifter (γ^H)	2.323	(0.008)		
Network surplus (β^M)			1.297	(0.006)
Control function				
A: Price / B: Premium	0.699	(0.010)	7.987	(0.029)
Network surplus			-0.582	(0.008)
Median elasticity	-0.79		-2.25	
N	261,857		163,034,142	

Notes: Panel A presents estimates of preferences for hospitals. The sample is a 30 percent random sample of inpatient admissions. The model includes insurer-hospital-diagnosis-year fixed effects. Panel B presents estimates of preferences for plans. The model includes an insurer-age fixed effect. Heterogeneity in price and premium sensitivity depends on policyholder attributes, where high income indicates those above 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 control function parameter is the coefficient on the first-stage residual. The reported elasticities are the median own-price in Panel A and own-premium in Panel B. Appendix Table A.6 shows estimates without the control function approach.

shows no effect of VI on treatment or quality—at least to the extent observable in inpatient data—we are left to conjecture what drives this effect. One plausible mechanism is direct steering through referrals, marketing, and preferential scheduling. Alternatively, patients might perceive significant benefits from preferential treatment at VI hospitals, such as better customer care or more convenient payment facilities. We remark that regardless of the mechanism, it must be consistent with the descriptive evidence and thus cannot meaningfully affect the intensity or quality of care.

6.2 Preferences over Plans

The second step is to estimate plan demand. The main parameters of interest are the coefficients on plan premiums, α_i^M, and on network surplus, β^M. To take the model to the data, we define policyholders' choice sets based on the market segments described in Section 3, grouping consumers by gender, household composition, and age groups (25–

40, 40–55, 55–65).[24] In terms of the outside option, we note that the public plan network consists primarily of public hospitals and relies on list prices and fixed coverage design. Hence, the assumption that the network surplus of the outside option is independent of the private market equilibrium is consistent with the setting.

The main identification challenge is that unobserved plan preferences might correlate with premiums and network surplus. We address premium endogeneity by using public hospital list prices to compute an actuarially fair premium for each household-plan, given their medical risk and the plan's cost-sharing structure. We use the average of these simulated premiums among rivals by market segment and year as an instrument. This captures household expected costs and the competitive pressure on premiums. The instrument is excluded under the assumption that the public insurer does not condition on future realizations of private enrollees' preferences when choosing list prices. Similarly, we instrument network surplus—which depends on hospital prices—using the average of rival plans' network surplus, the average share of rivals offering preferential access to the same hospitals as the plan, and the fraction of plans offered by the same insurer that have the same preferential network structure.[25] From the game's timing, plan coverage is not a function of unobserved plan preferences.

We estimate plan demand using individual-level enrollment data, which we complement in two ways. First, we use our hospital demand estimates to compute the network surplus (WTP^H_{ijt}) of all plans. Second, we use data on the share of publicly insured individuals by demographic group to account for the outside share. We construct a representative publicly insured household for each market and compute its network surplus from private plans. We then estimate preferences by maximizing the weighted likelihood of observed enrollment, where weights account for the share of publicly insured individuals. As with hospitals, we incorporate our instruments using a control function approach. We specify preference heterogeneity for insurers as varying by policyholder age.

Table 2-B shows the estimates. Consumers who are older, wealthier, or have dependents are less elastic to premiums. Single women are substantially less responsive to premiums than single men. The median own-premium elasticity is -2.3.[26] Heterogeneity

[24]The market segmentation by gender and household structure follows the regulatory framework. In principle, insurers can price plans at narrower age groups, but a review of marketing material and insurer websites suggests that most plans are grouped within these broader age tiers. This somewhat simplifies the analysis by reducing the dimensionality of the pricing problems, but otherwise should not meaningfully affect our results. see Appendix B.1 for further discussion.

[25]The likelihood that prices correlate with unobserved preferences is limited since negotiations aggregate all plans, reducing the effect the demand for any given plan has on prices.

[26]For comparison, Ho and Lee (2017) find a premium elasticity between -1.2 and -3, Curto *et al.* (2021)

in premiums implies significant differences across consumers in their relative valuation of network surplus. For example, a dollar increase in annual network surplus is worth about half a dollar in premiums for a low-income single man aged 25 to 40. In contrast, a high-income family whose policyholder is 55 to 65 years old values the same dollar increase in network surplus at nearly three dollars in premiums. These households, however, face higher premiums and greater exposure to hospital prices. Overall, households are more sensitive to changes in premiums than out-of-pocket hospital prices.

These findings suggest that, in equilibrium, VI hospitals' benefit from raising rivals' costs might be weaker than in a standard vertical market without cost-sharing or one with equal elasticities. To see this, note that an equilibrium price increase has a quality and a cost effect on rivals' plans. The quality effect stems from the negative impact of hospital prices on network surplus. When consumers are more responsive to premiums than to out-of-pocket prices, the loss in network surplus due to a dollar increase in out-of-pocket prices can be offset by less than a dollar decrease in premiums. The cost effect stems from the direct impact of prices on insurer payments. When consumers are more responsive to premiums than out-of-pocket prices, the insurer can limit the enrollment impact of the pass-through of costs to premiums by adjusting coverage and allocating some of the cost increase to out-of-pocket prices. Hence, the unequal elasticities and the fact that the insurer has two levers—premiums and coverage—will tend to mitigate demand losses from price increases by rival hospitals.

6.3 Price and Premium Setting Parameters

Pricing and premium setting depend on five primitives: hospital costs k_{hmt}, bargaining weights $\boldsymbol{\tau}$, insurers' administrative costs $\boldsymbol{\eta}$, and VI objective weights $\boldsymbol{\theta}$. We begin by presenting general results on the cross-sectional identification of these parameters in a large class of Nash-in-Nash bilateral oligopoly models (Gowrisankaran *et al.*, 2015; Ho and Lee, 2017).[27] We denote $\boldsymbol{X} = (\boldsymbol{c}, \boldsymbol{r}, \boldsymbol{\omega}, \boldsymbol{D}^H, \boldsymbol{D}^M, \Omega)$ the collection of observed (or identified) coverage, risk, resource weights, demand functions, and ownership matrices, and $\boldsymbol{Y} = (\boldsymbol{p}, \boldsymbol{\phi})$ the observed prices and premiums. We denote $F_{\boldsymbol{X},\boldsymbol{Y}}$ the implied data-generating process. For any hospital h, we denote K_h and T_h its identified sets of costs and

estimate -1.1, and Tebaldi (2024) estimates it to be between -1.3 and -2.

[27]Our pricing model includes cost-sharing like Gowrisankaran *et al.* (2015), downstream recapture as Ho and Lee (2017), and vertical integration as Crawford *et al.* (2018). The results do not require the number of hospitals, insurers, or periods to grow with the sample size. In fact, conditional on demand being identified, these are finite-sample identification (i.e., uniqueness) results. The proof does not rely on special properties of the logit model used for demand and can be extended to a larger class of additive random utility models.

bargaining weights. The following result is proven in Appendix C.2 and discussed below.

Proposition 2. *Given the model assumptions and $F_{\boldsymbol{X},\boldsymbol{Y}}$, then:*

1. ***Identification of** $(\boldsymbol{\eta},\boldsymbol{\theta})$: $(\boldsymbol{\eta}_{jt})_{j\in\mathcal{J}_m}$ is identified for all non-VI insurers $m \in \mathcal{M}$. Let $\boldsymbol{k}_t$ be identified and $c_{jht} \in [0,1)$ for all j,h such that $m(j)$ is VI with h at t, then $(\boldsymbol{\eta}_{jt})_{j\in\mathcal{J}_m}$ and $\boldsymbol{\theta}^m_t$ are identified for all m VI up to θ^m_{mt}.*

2. ***General identification of** $(\boldsymbol{k},\boldsymbol{\tau})$: $(\boldsymbol{k},\boldsymbol{\tau})$ are not separately identified. Moreover, for every hospital $h \in \mathcal{H}$ there exists a smooth bijection $g_{ht} : K_h \to T_h$, such that $(\boldsymbol{k}_{ht}, \boldsymbol{\tau}_{ht})$ are separately identified if and only if there exists additional $\mu^K_{ht} : \mathbb{R}^{|\mathcal{M}|} \to \mathbb{R}^{m_k}$ and $\mu^T_{ht} : \mathbb{R}^{|\mathcal{M}|} \to \mathbb{R}^{m_T}$ differentiable constraints with full row-rank Jacobians, such that $\{\boldsymbol{k}_{ht} \in K_h \mid \mu^K_{ht}(\boldsymbol{k}) = 0 \wedge \mu^T_{ht}(g_{ht}(\boldsymbol{k})) = 0\}$ is singleton.*

3. ***Specialized identification of** $(\boldsymbol{k},\boldsymbol{\tau})$: For hospital $h \in \mathcal{H}$, let $\bar{k}_{ht}$ denote its mean cost and $\bar{\boldsymbol{X}}$ the extended set of observables that includes $\bar{k}_{ht}$ for all h and t, then $(\boldsymbol{k},\boldsymbol{\tau})$ are identified from $F_{\bar{\boldsymbol{X}},\boldsymbol{Y}}$ up to heterogeneity within hospital-year in bargaining weights.*

Proposition 2.1 resolves the identification of the premium-setting and VI-specific parameters. First, for non-VI insurers, insurance admin costs are identified from premium optimality conditions, as is standard in Nash-Bertrand pricing games. For VI insurers, payoff objectives depend on the costs of their integrated hospitals. However, conditional on those costs being known, insurance administrative costs are point-identified from premium optimality conditions. Similarly, the relative weight placed on the profits of each member of a VI firm is point-identified from the optimality of prices within a VI firm. Hence, $\boldsymbol{\theta}^m_t$ is identified only up to some normalization (e.g., $\boldsymbol{\theta}^m_{mt} = 1$). This normalization, however, is without loss as only the relative weights are relevant for predicting counterfactual behavior of the integrated firm. Hence, Part 1 states that, conditional on identifying hospital costs, VI introduces no additional complexity to the identification challenge. Importantly, this is the only part of the proposition that requires positive cost-sharing ($c_{jht} < 1$), and it is only needed for identifying $\boldsymbol{\theta}^m$ from prices. For an alternative strategy that identifies $\boldsymbol{\theta}^m$ from upstream input choices, see Crawford *et al.* (2018).[28]

Proposition 2.2 shows that bargaining weights and hospital costs are not separately identified without additional constraints and offers a solution. The first part formalizes an intuition dating back to Grennan (2013), noting that for every hospital h and year t, there are $|\mathcal{M}|$ bargaining optimality conditions but $2|\mathcal{M}|$ unknowns ($|\mathcal{M}|$ bargaining

[28] We do not follow the strategy of identifying $\boldsymbol{\theta}$ from input choices given that in our model, input choices (plan design) are subject to unknown and potentially discontinuous costs, making that strategy impractical.

weights and $|\mathcal{M}|$ costs), preventing separate identification. The second part formalizes the usual solution to this problem: imposing modeling or data constraints. For example, Ho and Lee (2017) imposes homogeneity constraints on bargaining weights (μ^T) and orthogonality of residual cost heterogeneity to instruments given observed cost data (μ^K). Constraints are often imposed simultaneously on bargaining weights and hospital costs, and the proposition states the condition for such constraints to deliver identification.[29] Proposition 2.3 provides one common set of constraints that always satisfy the conditions for identification. It shows that data on the total or mean cost of a hospital is sufficient to separately identify costs and bargaining weights, as long as the latter are homogeneous across insurers.

Having established the general identification properties of the model, we now specialize the framework to our setting. In particular, the Chilean market features potential disagreement penalties for insurers. Hence, we specify insurer m's gains from trade with hospital h as $\Delta^H_{hm}\tilde{\pi}_h \equiv \tilde{\pi}_h - \tilde{\pi}_{h\backslash m} + \zeta_{hmt}$, where ζ_{hmt} is the disagreement penalty. We model $\zeta_{hmt} = z_{hmt} l_{hmt}$, where z_{hmt} is a known quantity and l_{hmt} is an iid random shock distributed L, where L is absolutely continuous and has full support. We assume that ζ_{hmt} is public information at the beginning of year t. We now denote $\bar{\boldsymbol{X}}$ the collection of observables, including both average hospital costs for each period and z_{hmt}, which we specify below. The following proposition establishes the identification properties of the extended model:

Proposition 3. *Without loss, write $k_{hmt} = \bar{k}_{ht} + \boldsymbol{s}_{hmt}\boldsymbol{\beta}_t$ with $\boldsymbol{s}_{hmt} \in \mathbb{R}^d$ with $d \leq |\mathcal{H}|(|\mathcal{M}| - 1)$ and stacked matrix $\boldsymbol{S}_h$.*[30] *Assume $\mathbb{E}[l_{hmt}|\bar{\boldsymbol{X}}, \boldsymbol{S}_h] = 0$ for all h, m, t, and for any t, either (i) $c_{jht} = 1$ for all j, h, or; (ii) $c_{jht} \in (\underline{c}, \bar{c}) \subset (0, 1)$ for all j, h, and for every $i \in \mathcal{I}$, $D^M_{i0t} > \frac{\bar{c}-\underline{c}}{2-(\bar{c}+\underline{c})}$. Then, $(\bar{\boldsymbol{k}}, \boldsymbol{\tau}, \boldsymbol{\beta}, L)$ are identified from $F_{\bar{\boldsymbol{X}},\boldsymbol{Y}}$ up to heterogeneity within hospital-year in $\boldsymbol{\tau}_{ht}$.*

Starting with the assumptions, we have a standard exclusion restriction on the disagreement penalty shock and a restriction on cost sharing. The latter takes on one of two forms: either there is no cost-sharing (as in fully vertical markets), or cost sharing is partial, but the insurance outside option is sufficiently strong. This second variant is the one useful for our application, where $\bar{c} \approx 0.8$ and $\underline{c} \approx 0.5$, implying a bound $D^M_{i0t} > 3/7$; well below the market share of public insurance. This assumption is sufficient to establish

[29]In Appendix C.2, we show in a graphical example that the sufficient and necessary conditions in the proposition are not trivial: In a setting with only two insurers, pairing a constraint that bargaining weights are homogeneous within hospitals with an identical constraint on costs fails to achieve identification, while pairing it with information on total costs attains identification. Nevertheless, either restriction reduces the problem to one with two equations and two unknowns.

[30]This parametrization is without loss as we can always take $\boldsymbol{s}_{hmt}$ to be insurer-year fixed-effects, with one normalized insurer per hospital. This recovers the setup of the Proposition 2.

a law of demand for hospitals when coverage is partial and heterogeneous.

This final result is the most practical. Conditional on simplifying the bargaining weight structure—a consequence of Proposition 2.3—it provides boundaries on our ability to shift statistical uncertainty between hospitals' gains from trade (through $\boldsymbol{k}_{ht}$) and insurers' gains from trade (through ζ_{hmt}). These different models are observationally equivalent for any pair of hospital-insurer negotiations for which gains from trade are positive; hence, where to assign the heterogeneity should depend on the context. We have institutional reasons to expect a non-negligible ζ_{hmt} term in our setting. Moreover, a model without such uncertainty is rejected by the data, since a few insurers have minimally negative gains from trade with some high-priced non-VI hospitals, yet agreement occurs in the data.

Having established identification, we proceed to estimation. Given we have a limited number of periods, we adopt a specification that is parsimonious in time heterogeneity. First, following Proposition 3, we assume that bargaining weights are homogeneous within a hospital, across insurers, and over time. Second, since our data on hospital average costs shows little variation over time, we assume that average hospital costs are constant and specify $\boldsymbol{s}_{hmt}$ as insurer-year fixed effects. Third, we assume the penalty that insurer m faces upon a disagreement with hospital h is set by the courts as proportional to the loss of surplus for m's enrollees. Thus, we specify z_{hmt} as a measure of surplus lost, which is independent of the price p_{hmt}.[31] Finally, we assume L is a normal distribution with unknown mean and variance. We adopt a two-step iterative approach to jointly recover the pricing and premium-setting parameters. We start from a guess of hospital costs and use the optimality conditions for premiums and VI hospital prices to recover $\boldsymbol{\eta}$ and $\boldsymbol{\theta}$. Then, we use costs and VI weights to maximize the likelihood of hospital prices for non-VI pairs to estimate $\boldsymbol{k}$, τ_h, and the parameters in $L(\mu_l, \sigma_l^2)$. We update the guess of hospital costs and iterate until convergence. We present the details in Appendix C.3.

Table 3 shows the estimates.[32] Insurers hold slightly more bargaining power than hospitals. However, the highest-quality non-VI hospitals— star hospitals h_1 and h_6—have the most bargaining power and the highest costs. All low-cost hospitals are VI, but not all VI hospitals are low-cost. Moderate differences in bargaining weights across hospitals suggest most of the price variation is rationalized by differences in costs and gains from trade. The latter includes consumers' preferences and the ability of VI firms to recapture profit

[31] We assume the courts will use the average among m's rivals' prices at hospital h at time t to impute the unobserved p_{hmt} that would have resulted under agreement. We assume rivals are not strategic regarding their influence on these penalties. Using past or public prices for the imputation yields similar estimates.

[32] We restrict attention to each insurer's 70 percent most popular plans, excluding a tail of tiny plans. We capture over 90 percent of the enrollment of most insurers.

Table 3: Estimated price and premium setting parameters

	(1)	(2)	(3)	(4)	(5)	(6)
	A. Bargaining weight (τ)		**B. Hospital cost** (k^H)		**C. VI weight** (θ)	
Hospital	Coef.	S.E.	Coef.	S.E.	Coef.	IQR
h_1	0.210	(0.015)	4.692	(0.020)		
h_2 (m_a)	0.518	(0.011)	1.622	(0.043)	1.519	[1.471, 1.528]
h_3 (m_a)	0.588	(0.012)	1.765	(0.051)	1.756	[1.708, 1.826]
h_4 (m_b)	0.680	(0.017)	2.100	(0.100)	1.264	[1.217, 1.306]
h_5	0.534	(0.032)	3.016	(0.081)		
h_6	0.229	(0.022)	6.042	(0.030)		
h_7 (m_b)	0.648	(0.015)	2.964	(0.051)	1.319	[1.316, 1.369]
h_8 (m_a)	0.514	(0.012)	3.231	(0.041)	1.333	[1.166, 1.357]
h_9	0.167	(0.011)	3.499	(0.017)		
h_{10} (m_c)	0.493	(0.016)	3.456	(0.030)	2.367	-
h_{11} (m_b)	0.658	(0.017)	1.678	(0.085)	1.874	[1.753, 1.978]
Penalty mean (μ_l)	1.331	(0.022)				
Penalty S.D. (σ_l)	1.498	(0.006)				
Median administrative cost overhead (η_j/ϕ_j)			18.17%			
Median insurance markup			58.27%			
Median hospital markup			26.74%			
N			220			

Notes: The first column lists hospitals and, in parentheses, their integrated insurer. In Panels A and B, we report standard errors in parentheses. In Panel B, hospital costs are in thousands of dollars, normalized to the resource intensity of an average delivery for a woman aged between 25 and 40. In Panel C, the interquartile range of VI weights (θ) estimates across years are shown in brackets. Hospital h_{10} is only integrated for one year and hence does not have a range of estimates. Median administrative cost and insurance markup are at the plan level and are weighted by enrollment.

losses from disagreements. The estimated distribution of disagreement penalties suggests that, on average, insurers expect a penalty equivalent to 133 percent of the consumer surplus loss from disagreement. The estimates indicate that VI firms place more weight on their integrated hospitals' profits, which is consistent with the uneven regulatory environment, as insurers face more scrutiny than hospitals, creating an incentive to tunnel profits (Gandhi and Olenski, 2024). Taken together, these estimates imply that the median insurer spends 18 percent of its revenue paying for non-inpatient care costs—including outpatient services, limited prescription drug coverage, and administration—and the median plan and hospital markups are 58 and 27 percent, respectively.[33,34]

[33]Insurers in Chile also serve a role akin to short-term disability insurers in the U.S., paying for lost wages for short terms. This introduces additional costs to insurers that are not captured in our data or analysis. It is independent of market structure or plan design and, therefore, secondary to our analysis.

[34]We quantify the extent of allocative inefficiencies related to moral hazard and adverse selection in Appendix D.1. The notion of allocative moral hazard we consider is that spending would be efficient if consumers chose care based on its costs and quality, rather than their out-of-pocket prices, which are distorted by cost-sharing and markups. Our results show that pricing and coverage distortions from VI

6.4 Plan Design Costs

Plan design involves a mix of continuous and discrete choices associated with coverage and tiering. To account for this, we divide the cost structure into a smooth part, K^r_{mt}, and a discontinuous part, K^o_{mt}, writing $K_{mt}(\boldsymbol{c}_j) = K^r_{mt}(\boldsymbol{c}_j) + K^o_{mt}(\boldsymbol{c}_j)$. We refer to these as the *regulatory* and *tiering* costs of plan design, respectively, although the names are merely suggestive of their potential origins. For example, regulatory costs might capture the shadow cost of regulatory scrutiny, related to minimum coverage and solvency requirements in our setting.[35] Tiering costs might include marketing costs associated with offering preferential plans and organizational pressures to place integrated hospitals in the preferential tier of plans, beyond their effect on profit. We identify the smooth regulatory cost from the first-order optimality conditions of coverage rates and the tiering costs from profit inequalities implied by the optimality of observed tiers. Importantly, we allow VI to affect both cost structures flexibly.

We specify regulatory costs 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 differentiable function, $\underline{c}_{jt}$ and $\bar{c}_{jt}$ are the base and preferential plan coverage rates, and $\underline{\zeta}_{jt}$ and $\bar{\zeta}_{jt}$ are mean-zero iid normal shocks with unknown variance that capture idiosyncratic differences in regulatory scrutiny across plans. We compute the marginal profit of each plan's base and preferential coverage, accounting for their impacts on prices, premiums, and demand. We then estimate the regulatory cost of coverage by maximum likelihood, flexibly estimating $c^K(\cdot)$ using a polynomial series. Appendix C.4 provides details.

Regarding tiering costs, we specify $K^o_m(\boldsymbol{c}_{jt}) = \sum_{h\in\mathcal{H}} w_{hjt}(\vartheta_{hmt} + \varsigma_{hjt})$, where w_{hjt} indicates whether hospital h is preferential for plan j, ϑ_{hmt} is the cost of tiering, and ς_{hjt} is an unobserved cost shock. As firms observe these cost shocks when designing plans, tiering decisions are subject to an unobserved selection problem (Pakes *et al.*, 2015). Following Canay *et al.* (2023), we assume that for a vector of instruments Z^K_{hjt}, the unobserved cost shocks satisfy $\mathbb{E}[\varsigma_{hjt}|Z^k_{hjt}] = 0$ and $\mathbb{E}[\varsigma_{hjt}|Z^k_{hjt}, w_{hjt}] \leq \bar{\varsigma}$ for some known positive $\bar{\varsigma}$.[36] Using these assumptions, the optimality of observed tiering implies lower and upper bounds

incentives and market power lead to 47 percent higher spending than efficient. Misallocation is stronger at VI hospitals, since VI insurers provide higher coverage at them. In addition, we find evidence for adverse selection as measured by a positive correlation between consumers' willingness to pay (WTP) and expected inpatient cost for each plan (Einav *et al.*, 2010). This correlation operates through two channels: Riskier consumers benefit more from higher plan generosity and also from greater access to high-quality hospitals.

[35]Given the evidence of partial enforcement in the data and previous literature discussed in Appendix A.2, we model regulation as an unknown cost rather than a strict constraint.

[36]As noted by Canay *et al.* (2023), assuming an upper bound on the conditional expectation of cost shocks is weaker than assuming conditional independence. In our setting, conditional independence leads to point identification of tiering costs, as is commonly documented in the literature.

Table 4: Estimates of plan design cost parameters

	Coef.	S.E.		Coef.	S.E.		Coef.	S.E.
A. Coverage tier levels			**B. Mean coverage x hospital**			**C. Spread x mean coverage x insurer**		
Base	-2.802	(0.392)	h_1	-3.368	(0.666)	m_a	-2.221	(0.903)
Preferential	0.714	(0.247)	h_2	-0.327	(0.644)	m_b	0.761	(0.750)
Base^2	4.871	(2.344)	h_3	0.623	(0.349)	m_d	0.153	(0.924)
Preferential^2	-2.973	(0.472)	h_4	0.039	(0.525)	m_e	-1.590	(0.874)
Base x Preferential	6.389	(1.254)	h_5	1.423	(0.398)	**D. Regulatory cost shock variance**		
VI coverage × m_c	16.148	(3.403)	h_6	-4.520	(1.097)	ln(ζ)	1.222	(0.018)
VI coverage × m_b	9.994	(0.786)	h_7	-0.169	(0.406)	ln($\bar{\bar{\zeta}}$)	2.500	(0.018)
VI coverage × m_a	6.527	(0.869)	h_8	-0.855	(0.529)	**E. Tiering costs bounds (ϑ)**		
			h_9	0.674	(1.578)	Non-VI	[0.120, 0.139]	
			h_{10}	1.949	(1.113)	VI × m_a	[-0.066, 0.083]	
			h_{11}	-3.885	(0.543)	VI × m_b	[-0.386, 0.009]	
						VI × m_c	[-0.061, 0.495]	
N	3,256	Moments	30					

Notes: Panels A, B, and C show elements of $c^K(\cdot)$ that enter the regulatory cost (K^r_m) exponentiated. Spread is the difference between preferential and base coverage. Mean coverage is averaged within a plan across hospitals. The functional form of $c^K(\cdot)$ includes segment fixed effects interacted with each plan's coverage spread, omitted for brevity. Estimates in Panel C are relative to the normalized effect of m_c. Panel D shows estimates of the variance of the regulatory costs of plan design. We present them in logs as our estimation strategy recovers the distribution as log-normals. Panel E shows the estimated set of tiering parameters. All costs are in millions of dollars per hundred thousand market segment enrollees. The upper bound on tiering cost shocks $\bar{\xi}$ is set to 2/3 of a million dollars; approximately the interquartile range of the estimated distribution of profit differences from changing tiering decisions, $\Delta\tilde{V}_{hmt}$. Identified sets are estimated using the bootstrap method in Chernozhukov *et al.* (2019).

on tiering costs: if insurer m decided to leave hospital h in the base tier of plan j, then tiering costs must have been sufficiently large; if not, then tiering costs must have been sufficiently small. These optimality conditions form the basis of our estimator, which we implement using the test of Chernozhukov *et al.* (2019). Appendix C.4 provides details.

We allow tiering costs to vary depending on whether the hospital-insurer pair is VI and differ across VI firms. As instruments, we use the average non-price component of healthcare indirect utility of each plan-hospital ($\chi^H_{hdt} + \xi^H_{m(j)hdt}$), weighted by medical risk and aggregated across consumers. This instrument captures the relative value a plan offers when accessing a hospital due to factors other than coverage generosity and prices. It is unlikely to be correlated with changes in tiering costs, as it is driven by consumers' geographic location and the match quality of hospitals and insurers with specific diagnoses. These instruments are relevant, as the design-invariant value of a plan's access to a hospital affects the relative value of providing preferential coverage at that hospital. We also include the average of this instrument across other plans of the same insurer and separately across rivals.

Table 4 shows the estimates. The estimated regulatory costs indicate that the average marginal cost associated with a percentage point increase in the base or preferential coverage of a plan is equal to −\$1,549 and \$928, respectively. The negative marginal cost of base coverage suggests that insurers face meaningful lower bounds on base coverage: Insurers would rather offer lower coverage, but regulation constrains them to offer coverage at least as generous as the public plan's at private hospitals. The positive marginal cost of preferential coverage suggests that reinsurance and liquidity requirements curtail preferential plan generosity. Among the tiering costs parameters, the only one relevant for our main counterfactual is the one associated with tiering by non-VI insurers, which is estimated to cost between 0.12 and 0.14 million per one hundred thousand market segment enrollees. Given the tight bounds, we use the median point for counterfactuals.

7 Equilibrium Effects of Vertical Integration

To quantify the impact of VI, we simulate a counterfactual VI ban and compare it to the baseline market structure. We then decompose the impact of VI on premiums and prices from its effect on plan design by simulating a second counterfactual VI ban that holds plan design fixed at the status quo. In both counterfactuals, the VI profit weights are eliminated ($\theta_{ht}^{m} = 0$), forcing formerly integrated partners to negotiate prices, and the VI hospital demand shifter is removed ($\gamma^{H} = 0$), eliminating VI insurers' ability to steer patients other than through price and coverage. The first counterfactual equilibria involve redesigned plans, and both include reoptimized premiums, renegotiated hospital prices, and new enrollment and hospital choices by consumers. To focus on the impacts of vertical linkages, we preserve the hospital systems formed by VI firms. We present results for 2016. Throughout, we refer to formerly VI firms in the counterfactual simply as VI to reduce redundancy. We provide computational details in Appendix C.5.

7.1 Effects on Market Outcomes and Welfare

Breaking up VI firms has stark effects on them, as well as equilibrium effects on the rest of the market. Table 5-A reports results for changes in market outcomes, and Table 6-A reports impacts on market efficiency and welfare.

Plan design. Insurers respond to a VI ban by drastically changing their plan supply. The coverage for hospitals in the base coverage tier falls substantially at both VI and non-VI insurers by 6.5 and 21.8 points, respectively. In contrast, the coverage in the preferential tier increases slightly for both. The combination of these changes leads to a larger gap in generosity across tiers, intensifying demand steering through provider networks. The structure of these networks also changes markedly in the counterfactual. Since VI insurers no longer internalize the profits of their former partners ($\theta_{ht}^{m} = 0$), they drastically shrink the presence of VI hospital in the preferential tier of their plans: the likelihood of a VI hospital being preferential in its formerly integrated insurer's plan falls by around a third, from 67.5 to 45 percent. This change in plan design by VI insurers induces competitive responses by rivals in the same direction: the likelihood of a VI hospital being preferential in a non-VI insurer plan drops by half, from 22.1 to 11.9 percent. Importantly, while access to VI hospitals worsens, both VI and non-VI insurers shift their plan design towards increased preferential coverage at non-VI star hospitals. In particular, star hospitals are 16 times more likely to be preferential in a VI insurer's plan than at baseline.[37] Taken together, these results suggest insurers offer a more diverse set of plans in the absence of VI, mostly driven by the shift of VI insurers away from their hospitals. The similarities between VI and non-VI plan coverage responses to the new environment suggest that plan generosity, much like prices, is a strategic complement.

Hospital prices and plan premiums. The average price of VI hospitals to their formerly VI insurers increases by 23.2 percent, consistent with higher hospital markups and double marginalization. VI insurers increase premiums by 9 percent due to higher costs and the elimination of the incentive to attract enrollees with lower premiums to steer them toward their integrated hospitals.

Changes in plan design impact preferential access to hospitals and stimulate competition on both sides of the market. VI hospitals are no longer siloed by preferential access from their partners, and VI insurers are no longer differentiated from their competitors by their preferential access to VI hospitals. The literature on countervailing power notes that the net effect of these forces is ambiguous (Ho and Lee, 2017). Intensified hospital competition lowers their bargaining leverage, reducing prices. Similarly, increased insurer competition makes them more easily substitutable from the hospitals' perspective, raising prices. Column (3) of Table 5 shows that for the average hospital, insurer competition dominates, resulting in negotiated hospital prices increasing by 10.4 percent.

While the average hospital price increases, column (4) in Table 5 shows that the average price paid by consumers decreases due to demand steering. Low-value care is reallocated to the outside option due to price increases, particularly so at the city's periphery, where distance to low-cost private hospitals limits access. High-value care, instead, is directed more towards star hospitals and high-quality VI hospitals. While prices at star hospitals increase moderately—due to increased demand from VI insurers—prices at high-quality VI hospitals drop significantly for two reasons. First, high-quality VI hospitals face

[37]Appendix Figures A.4-A and A.4-B show a more detailed breakdown of changes in coverage across all insurer-hospital pairs. The results indicate that insurer competition intensifies on two extremes. First, nearly all insurers increase their preferential access to the highest-quality star hospital (h_6). Second, they increase access to the cheapest, most centrally located hospital (h_{11}).

Table 5: Effects of vertical integration on hospital prices and plan design

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
			A. Full effect		B. Decomposition			
	Baseline				Endogenous prices and premiums		+Endogenous plan design	
	Raw	Weighted	Raw change	Weighted change	Raw change	Weighted change	Raw change	Weighted change
Plan coverage rate (p.p.)								
VI insurer base coverage	52.54	48.46	-6.26	-6.56	–	–	-6.26	-6.57
VI insurer preferential coverage	78.15	76.23	2.10	1.81	–	–	2.10	2.53
Non-VI insurer base coverage	58.88	58.67	-21.76	-19.64	–	–	-21.76	-21.14
Non-VI insurer preferential coverage	85.75	84.94	0.11	0.93	–	–	0.11	0.49
Plan preferential tiering (p.p.)								
VI insurer self-preferencing	67.49	65.43	-22.68	-26.79	–	–	-22.68	-28.76
VI insurer to other-VI hospital	13.56	12.77	-7.39	-6.60	–	–	-7.39	-7.64
VI insurer to non-VI hospital	4.69	4.97	1.15	2.39	–	–	1.15	2.86
VI insurer to star hospital	0.48	0.45	7.42	8.17	–	–	7.42	8.36
Non-VI insurer to other-VI hospital	22.12	21.40	-10.21	-9.34	–	–	-10.21	-8.90
Non-VI insurer to non-VI hospital	16.91	17.98	-7.46	-9.26	–	–	-7.46	-9.01
Non-VI insurer to star hospital	13.97	18.22	2.02	-1.91	–	–	2.02	-0.79
Hospital prices ($M)								
Within VI firm	3.61	4.04	0.84	-0.13	0.08	-0.06	0.76	-0.08
VI hospital to other-VI insurer	3.78	3.91	-0.11	-0.29	-0.28	-0.28	0.17	-0.01
VI hospital to non-VI insurer	3.69	3.84	0.28	-0.34	-0.10	-0.23	0.38	-0.11
Non-VI hospital to VI insurer	5.86	6.31	0.24	0.45	0.54	0.02	-0.30	0.43
Non-VI hospital to non-VI insurer	5.44	6.10	1.05	0.27	1.74	1.23	-0.69	-0.97
All hospitals	4.32	3.89	0.45	-0.44	0.46	-0.10	-0.01	-0.33
Plan premiums ($M)								
VI insurer	1.22	1.15	0.11	0.04	0.05	0.03	0.06	0.00
Non-VI insurer	1.38	1.29	-0.05	-0.07	0.02	0.02	-0.07	-0.09
All insurers	1.31	1.23	0.02	-0.02	0.03	0.03	-0.02	-0.05
Hospital market shares (p.p.)								
VI hospital	38.66	–	-16.69	–	-15.74	–	-0.94	–
Non-VI hospital	30.95	–	0.27	–	2.83	–	-2.55	–
Insurer market shares (p.p.)								
VI insurer	10.84	–	-2.84	–	-2.39	–	-0.45	–
Non-VI insurer	13.18	–	2.17	–	0.01	–	2.16	–
Admission shares (p.p.)								
Within VI firm	72.00	–	-50.67	–	-43.96	–	-6.72	–
VI hospital from other-VI insurer	3.76	–	7.33	–	9.10	–	-1.77	–
VI hospital from non-VI insurer	24.24	–	43.34	–	34.85	–	8.49	–
Non-VI hospital from VI insurer	22.95	–	10.64	–	12.55	–	-1.91	–
Non-VI hospital from non-VI insurer	77.05	–	-10.64	–	-12.55	–	1.91	–

Notes: Plan coverage and tiering in percentage points, prices in thousands of dollars per unit of resources, premiums in thousands per year, and shares in percentage terms. VI insurer self-preferencing is the likelihood that a VI hospital is preferential in a VI plan. Other-VI and non-VI 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 hospital resources; for premiums and coverage, it is by plan demand. Panel B displays partial changes: Columns (5) and (6) account for changes in hospital prices and premiums, keeping coverage fixed; columns (7) and (8) show the additional impact of plan design adjustments. Their sum is the Full effect. All changes reported in columns (3)–(8) are in levels.

intensified competition from star hospitals, as the vertical segmentation produced by VI networks is eliminated. Second, they face reduced bargaining leverage as the incentive to raise rivals' costs and the demand steering by formerly VI insurers are eliminated under the VI ban. Thus, more high-value care is delivered at high-quality providers, resulting in lower quality VI hospitals seeing the greatest loss in demand and sharpest increases in prices, mainly with former partners. Accordingly, average prices paid by consumers decrease by 11 percent.[38]

Demand. Changes in plan design, premiums, and hospital prices reallocate enrollment and admissions. In the hospital market, VI hospitals lose almost half their market share as former partners no longer steer demand their way. Accordingly, admissions shares equilibrate between VI and non-VI insurers. In particular, before the ban, VI insurers accounted for 72 percent of admissions at their integrated hospitals, whereas after the ban, they account for 21 percent, similar to other insurers. Similarly, the share of admissions from VI insurers at non-VI hospitals increases from 23 to 34 percent. These results largely reflect the shift in VI insurers' plan coverage towards rival hospitals and the removal of the VI demand shifter. In the insurance market, VI plans' network surplus decreases despite improvements in access to star hospitals. This leads to a 26 percent loss in market share, most of which is recaptured by non-VI rivals.

Spending. The equilibrium outcomes described above lead to a decrease in inpatient spending of 6.7 percent by private plan enrollees and of 6 percent by all enrollees, as shown in Table 6-A. Even though plan preferential tiers include only about half the number of providers than at baseline, there is more variety in preferential structure and slightly more generous coverage at preferential providers. Thus, households of heterogeneous risk find better-matching networks, and price-sensitive consumers have greater access to cheaper medium- to low-quality hospitals. Higher prices at star hospitals are partially offset by lower average premiums, stimulated by increased insurance competition and greater access to high-quality formerly VI hospitals at lower prices. As a result, total household spending on healthcare and insurance decreases by 2 percent, despite plan actuarial value falling by 11 percentage points on average.

Market efficiency. Removing the distortions introduced by VI improves market efficiency. On the hospital side, we evaluate the allocative efficiency of demand by comparing realized demand patterns with those that would emerge from a competitive, frictionless

[38]Appendix Figures A.4-C and A.4-D show changes in demand and prices across all pairs of insurers and hospitals. Importantly, the highest quality VI hospital, h_7, sees its price decrease by nearly 20 percent to almost all insurers, including its former VI insurer.

Table 6: Effects of vertical integration on spending, efficiency and welfare

	(1)	(2)	(3)	(4)
		A. Full effect	**B. Decomposition**	
			Endogenous prices and premiums	+Endogenous plan design
	Baseline	Change	Change	Change
Healthcare spending				
Inpatient spending | private plan ($M)	1.04	-0.07	-0.12	0.06
Inpatient spending ($M)	0.50	-0.03	-0.03	0.01
Total household spending ($M)	1.10	-0.02	-0.01	-0.01
Actuarial value (%)	0.66	-0.11	-0.02	-0.08
Insurance market efficiency				
Spending allocative inefficiency (%)	44.10	-16.43	-14.45	-1.98
Median marginal value of coverage ($M)	15.08	-8.99	-6.88	-2.11
Spread marginal value of coverage ($M)	56.97	-37.37	-31.28	-6.09
Consumer surplus				
VI enrollees (per member, $M)	–	-0.08	-0.66	0.59
Non-VI enrollees (per member, $M)	–	0.52	-0.56	1.08
Total consumer surplus ($MM)	–	62.65	-86.26	148.91
Share better off	–	0.72	0.01	0.71
Profits ($MM)				
VI hospitals	105.04	-54.83	-50.50	-4.33
Non-VI hospitals	109.56	4.46	57.81	-53.35
VI insurers	554.68	-106.77	-85.25	-21.53
Non-VI insurers	713.83	136.18	35.06	101.12
VI profit objective	698.77	-181.14	-156.73	-24.41
Total welfare ($MM)	–	41.69	-129.13	170.82

Notes: Healthcare spending in thousands per household. Actuarial value is the share of expected payments covered by insurers. Spending allocative inefficiency is relative to the first-best inpatient spending. The marginal value of coverage is the derivative of total welfare with respect to a plan's coverage at a hospital, accounting for equilibrium effects. Consumer surplus for VI enrollees is the average surplus among VI plan enrollees, unweighted by demand. Non-VI enrollees' consumer surplus is defined analogously. Profits and total consumer surplus are measured in millions of dollars per year. Panel A displays the Full effect of banning VI. Panel B displays partial changes: Column (3) accounts for changes in hospital prices and premiums keeping coverage fixed; column (4) shows the additional impact of plan design adjustments. Their sum is the Full effect. All changes reported in columns (3)–(8) are in levels.

market. In this first-best, hospital care is priced at costs, and demand is not influenced by cost-sharing, networks, or steering by VI insurers. Relative to this benchmark, we find that 44 percent of inpatient spending is inefficient at baseline and that banning VI eliminates 37 percent of this allocative inefficiency.

On the insurance side, we use our plan demand model to compute the marginal value of increasing coverage to assess the efficiency of insurance provision. At baseline, coverage is inefficiently low on average: a marginal increase in coverage to the median patient increases welfare by thousands of dollars. Importantly, there is vast dispersion in this inefficiency, largely driven by VI insurers' incentive to skew coverage in favor of their hospitals and away from rivals. Banning VI eliminates incentives for self-preferencing and increases competition, leading to a more efficient coverage provision. Both the median and dispersion of the marginal value of coverage decrease drastically after banning VI.[39]

Welfare. Evaluating the impacts of banning VI on consumer welfare requires taking a stance on whether the VI demand shifter, γ^H, is welfare-relevant. Motivated by the descriptive evidence on the impacts of VI on hospital quality in Section 4, we develop our baseline welfare analysis assuming it is not.[40] Table 6 shows this quantification, indicating that consumer surplus increases by $62.7 million. Consumer gains are heterogeneous, with 72 percent of individuals being better off, and non-VI plan enrollees benefiting the most, gaining an equivalent of 5 average monthly premiums.[41]

Banning VI leads to an equilibrium with industry profits $21 million lower than at baseline. In total, VI firms lose $161.6 million in profits, 87 percent of which is captured by rivals. More intense hospital competition leads to a decrease in hospital profits of $50.4 million. In turn, lower inpatient costs lead to an increase in insurer profits of $29.4 million, most of which is captured by non-VI insurers, whose profits increase by 19 percent.[42]

Combining these results, total welfare increases by $41.7 million, since gains to consumers more than compensate for lower industry profits. This result suggests that any benefits of VI in terms of reducing double marginalization are offset by distortions in hospital prices, premiums, and plan design. These findings, however, depend on whether γ^H is actually welfare-relevant. Below, we present two pieces of evidence related to the impact of this assumption on the welfare estimate. First, we show that the surplus gains from banning VI do not stem mechanically from the elimination of the purely-distortionary γ^H. In fact, banning VI while holding plan design at the status quo level, eliminates γ^H but

---

[39]Coverage is underprovided due to adverse selection, which is worsened by intensified insurer competition. In particular, we estimate a correlation between residualized willingness to pay for insurance and expected cost of 0.37, which increases to 0.54 once VI is banned. However, improvements in access, prices, and spending efficiency curtail its adverse effects on welfare.

[40]Recall that the VI demand shifter γ^H is identified from demand changes between h_{10} and m_c following their de-integration, not rationalized by changes in prices and coverage. The evidence in Section 4 suggests this event did not affect quality or treatment decisions, suggesting γ^H captures selective referrals, marketing, or similar steering mechanisms.

[41]Appendix Table A.7 shows results for heterogeneity of changes in consumer surplus across consumers. Most consumers benefit from banning VI, particularly those in the city periphery, those with dependents, and those who are older. Private plan enrollees benefit more than those in the public plan, given the low substitution between private and public plans.

[42]We do not incorporate costs associated with plan design in profits. These costs fall mechanically as designing VI plans is estimated to be more costly. However, these costs might reflect intensified regulatory scrutiny rather than true payments. Hence, our insurer profit changes should be considered a lower bound.

decreases surplus. Second, in Section 7.4.1 below, we examine how the welfare effects of VI would change if a fraction of γ^H was welfare-relevant.

7.2 Decomposing the Roles of Hospital Prices, Plan Premiums, and Plan Design

To better understand the impacts of banning VI, we disentangle the changes induced by prices and premiums from those caused by plan redesign. Consistent with the timing of our model, we first consider an equilibrium in which firms only adjust prices and premiums upon a VI ban. We compare it to the status quo to isolate the contribution of price and premium responses to the overall effect of VI, and then to the full impact of a VI ban to isolate the incremental contribution of plan design responses. Tables 5-B and 6-B show the results from this decomposition.

When only premiums and price negotiations adjust, banning VI leads to increases in average prices and premiums. Consistent with the introduction of double marginalization, hospital prices within VI firms and VI insurer premiums increase. Similarly, the elimination of the incentive to raise rivals' costs leads to a decrease in the prices charged by VI hospitals to rival insurers. Average paid hospital prices decrease, though the change is smaller than when plan design can adjust, consistent with the absence of changes in steering through plan design. A key driver of these results is that VI insurers are locked into plans that offer preferential coverage to hospitals they no longer own, creating two distortions. First, former partners hold them up by controlling the value of their preferential plans, charging them higher prices. Second, as these higher prices lead to higher premiums, some consumers drop VI insurers' plans, creating a service gap in the market: Non-VI insurers could benefit from competing for these enrollees by redesigning plans to attract them.[43] Nevertheless, insurers cannot fix either distortion without adjusting the design of their plans. Overall, consumers are pushed toward less generous plans, and the introduction of double marginalization increases demand from VI plan enrollees to non-VI hospitals. This reduces the dependence of non-VI hospitals on non-VI insurers, allowing them to increase prices by 20 percent. Average premiums increase by 2.4 percent, as a result of both higher hospital prices and VI insurers not having an incentive to lower premiums in order to attract demand to their hospital.

In this environment, VI increases welfare. Intuitively, with fixed plan design, the effects of VI depend on consumers' relative premium- and price-elasticities. Given that consumers are more premium- than price-elastic, VI hospitals' ability to harm rival insur-

[43]For example, they could offer a handful of plans similar to those from VI insurers, providing value to consumers while limiting the overall leverage VI hospitals have in negotiations.

ers with higher prices is limited since rivals can adjust premiums in response to mitigate demand losses.[44] Similarly, VI insurers are more effective at steering demand by lowering premiums and can capture profits more effectively through hospital prices. Thus, in this counterfactual with fixed coverage structure, the price and premium effects of VI on welfare are mostly positive. Once VI is banned, low price elasticities create large losses from double marginalization and increased prices at star hospitals. Overall, consumer surplus falls by 86 million, and only one percent of consumers are better off than at baseline.

It is only once insurers redesign their plans that the market improves. VI insurers redesign plans to steer consumers away from VI hospitals, but enrollment losses reduce their bargaining leverage, increasing average VI prices. Demand steering, however, works in their favor, translating into lower prices for consumers. Plan redesign strongly intensifies competition between high-quality VI hospitals and non-VI star hospitals. At baseline, star hospitals have outstanding power over non-VI insurers as the sole non-VI providers of quality care. Non-VI insurers are pushed to provide generous coverage at star hospitals and accept higher prices from them, as providing preferential coverage at high-quality VI hospitals is too expensive: VI hospitals would internalize this as an encroachment on their integrated plans' value, pushing them to negotiate higher prices. Banning VI eliminates this distortion, leading high-quality VI hospitals to further decrease prices to both rival and formerly VI insurers. Intensified competition among insurers for star hospitals translates to higher countervailing power and prices. However, higher enrollment at non-VI insurers increases their bargaining leverage, leading to a small net price decrease.

Non-VI insurers are the primary beneficiaries of plan design adjustments. While the key distortion in coverage is driven by VI insurers' inability to adjust when plan design is fixed, it is their rivals who reap the benefits of allowing star hospitals and high-quality VI hospitals to compete. By redesigning their plans, they attract more enrollees from VI insurers by offering similar network surpluses and lower premiums. Thus, VI insurers cannot maintain their dominant position once VI is banned, even with endogenous plan design. The end result is a more even distribution of profits among insurers.

Discussion. Our results describe how VI shapes market outcomes. When plan design is fixed, two forces govern the effects of VI: the reduction of double marginalization and the incentive to raise rivals' costs. The former dominates the overall impacts, as VI firms have little hope of recapturing enrollees from rival insurers due to insurer competition

[44] We explore this further in Appendix D.2, simulating a VI ban under different price and premium demand elasticities. We find that VI is welfare decreasing whenever consumers are more premium-than price-elastic, consistent with the intuition that this limits the impact of the raising rivals' cost threat.

and rivals' ability to adjust premiums. However, once plan design is allowed to adjust, equilibrium plan design responses overturn the effects of VI. VI insurers are incentivized to overprovide coverage at their hospitals, increasing rival non-VI insurers' dependence on non-VI star hospitals. As VI and non-VI insurers steer demand toward different hospitals, hospital competition weakens, increasing hospital bargaining leverage. This leads to higher prices, lower access, and less efficient spending, making VI detrimental.

This decomposition delivers two lessons about VI. First, the gains from banning VI are not mechanically due to the elimination of the potentially distortionary VI hospital demand shifter, γ^H. Those distortions are eliminated in the first step of the decomposition, yet banning VI in that context leads to a welfare loss. Second, accounting for endogenous plan design responses is crucial. VI firms can only retain profits if they can steer demand toward their hospitals when profitable, which is implemented primarily by plan design. As shown, this effect is strong enough to flip the sign of the welfare effect of VI.

7.3 The Role of the Relative Quality of VI Hospitals

Our analysis highlights the adverse effects of VI on hospital competition. In particular, the results suggest that siloing high-quality VI hospitals from their non-VI counterparts is a key driver of anticompetitive effects. To study this margin, we leverage that our market contains two VI firms, one with high-quality hospitals, m_b, and one with medium- to low-quality, m_a. We simulate counterfactual scenarios for the de-integration of each VI firm separately. The results, shown in Appendix Tables A.8 and A.9, indicate that banning only the high-quality VI firm leads to a similar total welfare change as the full ban. However, consumer surplus is nearly 2.5 times larger, indicating a significant reallocation of welfare gains. This ban increases competition among high-quality hospitals while maintaining pressure from a dominant low-priced VI firm. In turn, this competitive pressure pushes non-VI premiums downward, reallocating welfare gains to consumers.

Banning only the low-quality VI firm leads to greater but more unequal welfare gains. The results show that intensified insurance market competition leads all insurers to increase access to star hospitals. The pressure to provide additional access to high-quality care from lower-priced insurers pushes the remaining VI insurer, m_b, to increase access to star hospitals beyond what is attained under the previous counterfactual. A more competitive insurance market also results in lower incentives to raise rivals' costs for the remaining VI hospitals, leading to lower prices. Insurer m_b, however, still controls a dominant share of preferential access to its own high-quality VI hospitals. This results in uneven gains from the ban: While 80.5 percent of consumers would prefer a ban on the high-quality VI firm than remaining in the status quo, only 14 percent would rather have the low-quality VI firm banned. Together, these counterfactuals confirm that the siloing of high-quality providers plays a crucial role in determining the effects of VI.[45]

7.4 Cost Efficiencies and Quality Improvements

Proponents of VI in healthcare argue it might lead to reducing costs or increasing quality. The analysis above ignores such effects, since the evidence in Section 4 suggests that VI either does not impact cost or quality in our setting, or does so on aspects not captured well by inpatient claims data. In this section, we revisit this assumption and study the extent to which cost efficiencies or quality gains would affect the welfare impact of VI.

7.4.1 Cost Efficiencies. VI may eliminate wasteful spending by inducing hospitals to internalize their costs. While the evidence in Section 4 suggests VI does not impact treatment choices in ways related to cost, we cannot rule out reductions in administrative or information costs. To gauge the potential impact of such changes, we implement our counterfactual VI ban under varying degrees of cost efficiency. In the simulations, we assume banning VI would increase VI hospitals' cost of treating patients from their formerly integrated insurers. Appendix Figure A.5a shows that the welfare gains from banning VI decrease with cost efficiencies. Higher costs for formerly VI hospitals under the ban generate moderate gains for non-VI hospitals and losses for insurers. Consumers bear the brunt of the loss through higher prices and worse access. Welfare gains from banning VI are halved at a 7 percent cost efficiency, and VI becomes welfare-neutral at 18 percent. These magnitudes are substantial: recent evidence on merger efficiencies finds gains between 2 and 7 percent (Schmitt, 2017; Demirer and Karaduman, 2024).

7.4.2 Quality Improvements. VI might lead to quality improvements through better care coordination. We study this possibility by implementing our analysis for a range of quality improvements induced by VI. To do so, we adjust the fraction of the VI hospital demand shifter γ^H that is welfare-relevant. Appendix Figure A.5b shows the results. When quality effects are larger, the consumer surplus and total welfare gains from banning VI are smaller. A VI ban becomes welfare-neutral when 25 percent of the VI demand shifter is treated as quality. This magnitude would imply that VI hospitals can produce $473 worth of quality improvements per admission. Equivalently, it would require the average VI hospital to

[45] For comparison, in the U.S., some of the highest-quality hospitals are VI, including Mass General Brigham, Massachusetts General Hospital, Cleveland Clinic, and Mount Sinai Hospital, according to AHRQ (2023) and the U.S. News Ranking 2023-2024 (U.S. News, 2024).

close 23 percent of its quality gap with the highest-quality hospital at will. Given our descriptive evidence, this would need to happen without major changes to treatment decisions or readmission rates, which seems unlikely. Nevertheless, this is a meaningful bound for quality effects. For comparison, recent evidence on VI between primary care physicians and insurers in the U.S. finds quality gains from improved matching with specialists of about an order of magnitude smaller (Cho, 2025).[46]

8 Conclusion

This paper investigates the effects of VI between hospitals and insurers. Using comprehensive data from Chile, we show that enrollees of VI insurers are substantially more likely to seek care at VI hospitals, which charge lower out-of-pocket prices to their insurers' enrollees but treat them the same as other patients. Motivated by this evidence, we model and quantify how VI affects premiums, hospital prices, plan design, access to care, and welfare. We bridge the gap between the literature on healthcare competition, VI, and contract design by introducing profit sharing between integrated hospitals and insurers and endogenous plan design into a bilateral oligopoly model.

We find that VI decreases welfare, but only once its impact on plan design is accounted for. The benefits of eliminating double marginalization outweigh VI firms' incentives to foreclose rivals and the additional market power they gain. However, VI insurers overprovide coverage at their hospitals, creating an uneven competitive landscape for rivals. Skewed plan networks shield VI hospitals from strong non-VI competitors, leading to higher overall prices and demand misallocation. Breaking up VI firms equilibrates the insurance market, enhances efficiency, and reduces total spending.

Our findings contribute to the growing body of evidence on insurer-provider VI, a trend shaping healthcare markets in the U.S. and abroad. While we show that VI would need to drastically improve quality or reduce costs to counteract its impact on the Chilean market, the implications of VI on firm incentives to invest in these areas remain largely unknown and potentially significant. Additionally, VI may play a crucial role in selection, as integrated insurers might possess better information about population risk than non-integrated insurers and could design plans leveraging this information asymmetry. These issues open exciting avenues for future research on VI.

[46] Cho (2025) finds VI increases consumer utility by an equivalent of 8.3 miles traveled. A recent experimental study estimates willingness to pay to avoid a mile traveled at $11.45 (Schwartz *et al.*, 2021).

References

AHRQ (2023). Compendium of U.S. Health Systems. Last reviewed Dec, 2023. `https://www.ahrq.gov/chsp/data-resources/compendium.html`, accessed: 2024-06-12.

Aron-Dine, A., Einav, L. and Finkelstein, A. (2013). The RAND Health Insurance Experiment, Three Decades Later. *The Journal of Economic Perspectives*, **27**, 197–222.

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

Atalay, E., Hortaçsu, A. and Syverson, C. (2014). Vertical Integration and Input Flows. *American Economic Review*, **104** (4), 1120–48.

Baker, L. C., Bundorf, M. K. and Kessler, D. P. (2016). The Effect of Hospital/physician Integration on Hospital Choice. *Journal of Health Economics*, **50**, 1–8.

Berry, S. (1994). Estimating Discrete-Choice Models of Product Differentiation. *RAND Journal of Economics*, **25** (2), 242–262.

Bontemps, C., Gualdani, C. and Remmy, K. (2023). *Price Competition and Endogenous Product Choice in Networks: Evidence from the US Airline Industry*. TSE Working Paper 23-1415.

Brot-Goldberg, Z. C., Che, C. and Handel, B. R. (2022). *Pharmacy Benefit Managers and Vertical Relationships in Drug Supply: State of Current Research*. Working Paper 29959, National Bureau of Economic Research.

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.

Capps, C., Dranove, D. and Ody, C. (2018). The Effect of Hospital Acquisitions of Physician Practices on Prices and Spending. *Journal of Health Economics*, **59**, 139–152.

—, — and Satterthwaite, M. (2003). Competition and Market Power in Option Demand Markets. *The Rand Journal of Economics*, **34** (4), 737–763.

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

Chade, H., Marone, V. R., Starc, A. and Swinkels, J. (2022). *Multidimensional Screening and Menu Design in Health Insurance Markets*. Tech. rep., National Bureau of Economic Research.

Chen, L., Yi, L. X. and Yu, C. (2024). The Welfare Effects of Vertical Integration in China's Movie Industry. *American Economic Journal: Microeconomics*, **16** (2), 20435.

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.

Cho, H. (2025). Provider-Insurer Integration and Healthcare Delivery, Manuscript.

Chorniy, A., Miller, D. and Tang, T. (2020). Mergers in Medicare Part D: Assessing Market power, Cost Efficiencies, and Bargaining Power. *International Journal of Industrial Organization*, **68** (C).

Clínicas de Chile, A. (2012). *Dimensionamiento del Sector de Salud Privado de Chile*. Tech. rep., Clínicas de Chile A.G.

CMS (2024). Readmission Measures: Methodology. Accessed: 2024-06-21.

Collard-Wexler, A., Gowrisankaran, G. and Lee, R. S. (2019). Nash-in-Nash Bargaining: A Microfoundation for Applied Work. *Journal of Political Economy*, **127** (1), 163–195.

Cooper, Z., Craig, S. V., Gaynor, M. and Van Reenen, J. (2018). The Price Ain't Right? Hospital Prices and Health Spending on the Privately Insured. *The Quarterly Journal of Economics*, **134** (1), 51–107.

Copetta, M. (2013). Prestadores de Salud, Isapres y Holdings: ¿Relación Estrecha? *Depto Estudios, Superintendencia de Salud*.

Craig, S. V., Grennan, M. and Swanson, A. (2021). Mergers and Marginal Costs: New Evidence on Hospital Buyer Power. *RAND Journal of Economics*, **52** (1), 151–178.

Crawford, G. S., Lee, R. S., Whinston, M. D. and Yurukoglu, A. (2018). The Welfare Effects of Vertical Integration in Multichannel Television Markets. *Econometrica*, **86** (3), 891–954.

— and Yurukoglu, A. (2012). The Welfare Effects of Bundling in Multichannel Television Markets. *American Economic Review*, **102** (2), 643–85.

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

Cutler, D. M., Dafny, L., Grabowski, D. C., Lee, S. and Ody, C. (2020). *Vertical Integration of Healthcare Providers Increases Self-Referrals and Can Reduce Downstream Competition: The Case of Hospital-Owned Skilled Nursing Facilities*. Working Paper 28305, National Bureau of Economic Research.

Dafny, L. (2009). Estimation and Identification of Merger Effects: An Application to Hospital Mergers. *Journal of Law and Economics*, **52** (3), 523–550.

—, Duggan, M. and Ramanarayanan, S. (2012). Paying a Premium on Your Premium? Consolidation in the US Health Insurance Industry. *American Economic Review*, **102** (2), 1161–85.

Demirer, M. and Karaduman, Ö. (2024). *Do Mergers and Acquisitions Improve Efficiency? Evidence from Power Plants*. Tech. rep., National Bureau of Economic Research.

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

Dix, R., Moran, K. and Nguyen, T. M. A. (2024). Costs of Technological Frictions: Evidence from EHR (Non-) Interoperability, Manuscript.

DOJ (2024). New Health Care Task Force Will Tackle Competition Problems in Health Care Markets. `https://shorturl.at/Q4bD8`, accessed: 2024-06-12.

Einav, L., Finkelstein, A. and Cullen, M. R. (2010). Estimating Welfare in Insurance Markets using Variation in Prices. *The Quarterly Journal of Economics*, **125** (3), 877–921.

Eizenberg, A. (2014). Upstream Innovation and Product Variety in the U.S. Home PC Market. *The Review of Economic Studies*, **81** (3), 1003–1045.

Fan, Y. (2013). Ownership Consolidation and Product Characteristics: A Study of the US Daily Newspaper Market. *American Economic Review*, **103** (5), 1598–1628.

— and Yang, C. (2022). *Estimating Discrete Games with Many Firms and Many Decisions: An Application to Merger and Product Variety*. Working Paper 30146, National Bureau of Economic Research.

FTC (2024). *Docket No. ATR 102: Request for Information on Consolidation in Health Care Markets*. Tech. rep., Federal Trade Commission.

Gal-Or, E. (1999). Mergers and Exclusionary Practices in Health Care Markets. *Journal of Economics & Management Strategy*, **8** (3), 315–350.

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

Gandhi, A. and Olenski, A. (2024). *Tunneling and Hidden Profits in Health Care*. Working Paper 32258, National Bureau of Economic Research.

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

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.

Gray, C., Alpert, A. E. and Sood, N. (2023). *Disadvantaging Rivals: Vertical Integration in the Pharmaceutical Market*. Working Paper 31536, National Bureau of Economic Research.

Grennan, M. (2013). Price Discrimination and Bargaining: Empirical Evidence from Medical Devices. *American Economic Review*, **103** (1), 145–77.

Guardado, J. (2023). Policy Research Perspectives: Competition in Commercial PBM Markets and Vertical Integration of Health Insurers with PBMs. *American Medical Association*.

Handel, B. and Ho, K. (2021). Chapter 16 - The Industrial Organization of Health Care Markets. In K. Ho, A. Hortaçsu and A. Lizzeri (eds.), *Handbook of Industrial Organization, Volume 5*, *Handbook of Industrial Organization*, vol. 5, Elsevier, pp. 521–614.

Hart, O. and Tirole, J. (1990). Vertical Integration and Market Foreclosure. *Brookings Papers on Economic Activity, Microeconomics*, pp. 205–285.

Hausman, J. A. (1994). *Valuation of New Goods under Perfect and Imperfect Competition*. Working Paper 4970, National Bureau of Economic Research.

Ho, K. (2009). Insurer-Provider Networks in the Medical Care Market. *American Economic Review*, **99** (1), 393–430.

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

— and Lee, R. S. (2019). Equilibrium Provider Networks: Bargaining and Exclusion in Health Care Markets. *American Economic Review*, **109** (2), 473–522.

— and — (2023). Health Insurance Menu Design for Large Employers. *The RAND Journal of Economics*, **54** (4), 598–637.

Horn, H. and Wolinsky, A. (1988). Bilateral Monopolies and Incentives for Merger. *The RAND Journal of Economics*, **19** (3), 408–419.

Hortaçsu, A. and Syverson, C. (2007). Cementing Relationships: Vertical Integration, Foreclosure, Productivity, and Prices. *Journal of Political Economy*, **115** (2), 250–301.

Johnson, G., Lyon, Z. M. and Frakt, A. (2017). Provider-offered Medicare Advantage Plans: Recent Growth and Care Quality. *Health Affairs*, **36** (3), 539–547.

Koch, T. G., Wendling, B. W. and Wilson, N. E. (2021). The Effects of Physician and Hospital Integration on Medicare Beneficiaries' Health Outcomes. *Review of Economics and Statistics*, **103** (4), 725 739.

Lee, R. S. and Fong, K. (2013). Markov Perfect Network Formation: An Applied Framework for Bilateral Oligopoly and Bargaining in Buyer-seller Networks.

—, Whinston, M. D. and Yurukoglu, A. (2021). Chapter 9 - Structural Empirical Analysis of Contracting in Vertical Markets. In K. Ho, A. Hortaçsu and A. Lizzeri (eds.), *Handbook of Industrial Organization, Volume 4*, *Handbook of Industrial Organization*, vol. 4, Elsevier, pp. 673–742.

Liebman, E. (2022). Bargaining in Markets with Exclusion: An Analysis of Health Insurance Networks, manuscript.

Lucidi, S. and Rinaldi, F. (2010). Exact Penalty Functions for Nonlinear Integer Programming Problems. *Journal of Optimization Theory and Applications*, **145** (3), 479–488.

Luco, F. and Marshall, G. (2020). The Competitive Impact of Vertical Integration by Multiproduct Firms. *American Economic Review*, **110** (7), 2041–64.

Marone, V. R. and Sabety, A. (2022). When Should There Be Vertical Choice in Health Insurance Markets? *American Economic Review*, **112** (1), 30442.

Miller, K., Petrin, A., Robert, T. and Michael, C. (2023). *The Optimal Geographic Distribution of Managed Competition Subsidies*. Tech. rep., Technical Report.

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

Ordover, J. A., Saloner, G. and Salop, S. C. (1990). Equilibrium Vertical Foreclosure. *The American Economic Review*, **80** (1), 127–142.

PAI (2024). PAI-Avalere Study on Physician Employment-Practice Ownership Trends 2019-2023.

Pakes, A., Porter, J., Ho, K. and Ishii, J. (2015). Moment Inequalities and Their Application. *Econometrica*, **83** (1), 315–334.

Parekh, N., Hernandez, I., Radomski, T. R. and Shrank, W. H. (2018). Relationships between Provider-led Health Plans and Quality, Utilization, and Satisfaction. *Am J Manag Care*, **24** (12), 628–32.

Park, S., Langellier, B. A. and Meyers, D. J. (2023). Differences between Integrated and Nonintegrated Plans in Medicare Advantage. *Health Services Research*, **58** (3), 560568.

Petrin, A. and Train, K. (2010). A Control Function Approach to Endogeneity in Consumer Choice Models. *Journal of Marketing Research*, **47** (1), 3–13.

Prager, E. (2020). Healthcare Demand under Simple Prices: Evidence from Tiered Hospital Networks. *American Economic Journal: Applied Economics*, **12** (4), 196–223.

Salop, S. C. and Scheffman, D. T. (1983). Raising Rivals' Costs. *The American Economic Review P& P*, **73** (2), 267–271.

Schmitt, M. (2017). Do Hospital Mergers Reduce Costs? *Journal of Health Economics*, **52**, 74–94.

Schwartz, A. J., Yost, K. J., Bozic, K. J., Etzioni, D. A., Raghu, T. and Kanat, I. E. (2021). What is the Value of a Star when Choosing a Provider for Total Joint Replacement? A Discrete Choice Experiment. *Health Affairs*, **40** (1), 138–145.

Serna, N. (2024). Determinants of Provider Networks: Risk Selection vs. Fixed Costs, Manuscript.

Shepard, M. (2022). Hospital Network Competition and Adverse Selection: Evidence from the Massachusetts Health Insurance Exchange. *American Economic Review*, **112** (2), 578–615.

Song, Z., Johnson, W., Kennedy, K., Biniek, J. F. and Wallace, J. (2020). Out-Of-Network Spending Mostly Declined In Privately Insured Populations With A Few Notable Exceptions From 2008 To 2016. *Health Affairs*, **39** (6), 10321041.

Spengler, J. J. (1950). Vertical Integration and Antitrust Policy. *Journal of Political Economy*, **58** (4), 347–352.

Tebaldi, P. (2024). Estimating Equilibrium in Health Insurance Exchanges: Price Competition and Subsidy Design under the ACA. *The Review of Economic Studies*, **92** (1), 586–620.

Tilipman, N. (2022). Employer Incentives and Distortions in Health Insurance Design: Implications for Welfare and Costs. *American Economic Review*, **112** (3), 9981037.

U.S. News (2024). U.S. News Best Hospitals 2024-2025. `https://health.usnews.com/best-hospitals`, accessed: 2024-07-25.

Vatter, B. (2025). Quality Disclosure and Regulation: Scoring Design in Medicare Advantage. *Econometrica*.

Williamson, O. (1971). The Vertical Integration of Production: Market Failure Considerations. *American Economic Review*, **61** (2), 112–23.

Wollmann, T. G. (2018). Trucks without Bailouts: Equilibrium Product Characteristics for Commercial Vehicles. *American Economic Review*, **108** (6), 1364–1406.