

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

This paper supersedes a previous version titled “Vertical Integration between Hospitals and Insurers”. We thank our discussants, Christian Cox, Martin Gaynor, and Kate Ho, for their valuable suggestions. We also thank Joe Doyle, David Dranove, Liran Einav, Gautam Gowrisankaran, Igal Hendel, Ali Hortaçsu, Gastón Illanes, Robin Lee, Elena Prager, William Rogerson, Pietro Tebaldi, 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.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2024 by José Ignacio Cuesta, Carlos E. Noton, and Benjamin Vatter. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

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

ABSTRACT

We measure the impacts of vertical integration between insurers and hospitals. In the Chilean market, where half of private hospital capacity is vertically integrated, integration increases inpatient care spending by 6 percent and decreases consumer surplus and total welfare. Integrated insurers offer generous coverage at integrated hospitals, limited access to rival hospitals, and lower premiums. Competition for enrollees forces non-integrated insurers to provide additional coverage to high-quality non-integrated hospitals, resulting in plan networks that limit hospital competition. Whereas vertical integration reduces double marginalization, skewed cost-sharing structures—and their effect on hospital competition—more than compensate, leading to an overall negative welfare impact.

José Ignacio Cuesta
Stanford University
Department of Economics
579 Jane Stanford Way
Stanford, CA 94305
and NBER
jicuesta@stanford.edu

Benjamin Vatter
Massachusetts Institute of Technology
Sloan School of Management
Applied Economics
100 Main Street
Cambridge, MA 02142
bvatter@mit.edu

Carlos E. Noton
Escuela de Administración
Pontificia Universidad Católica de Chile
Avda Vicuña Mackenna 4860
Santiago
Chile
carlos.noton@uc.cl

An 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 policy-makers 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 more than half of all inpatient medical care is provided at hospitals that also sell insurance (AHRQ, 2023). Despite the prevalence of VI in healthcare markets, there is limited evidence about its impacts (Handel and Ho, 2021). In this paper, we study the equilibrium effects of VI between hospitals and insurers.

Hospitals and insurers interact in a vertical market. Downstream, insurers sell plans to consumers, offering access to upstream hospitals under a menu of prices and cost-sharing. VI in this market has ambiguous effects. The main arguments favoring VI are that it aligns insurer and hospital incentives to limit wasteful spending and eliminates the double marginalization problem (Spengler, 1950; Williamson, 1971). The counterargument is it grants market power to VI firms. It creates incentives for VI hospitals to use prices to increase rival insurers' costs or foreclose them and for VI insurers to limit patient access to rival hospitals (Salop and Scheffman, 1983; Hart and Tirole, 1990; Ordovery *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 integration and rich administrative data on claims, enrollment, and ownership. We combine this data with a bilateral oligopoly model with endogenous hospital prices, plan premiums, and cost-sharing. We find that VI firms steer demand to their hospitals by attracting consumers with low premiums and low hospital prices but limited access to rival hospitals. The generous coverage offered at VI hospitals induces non-VI rival insurers to overprovide coverage at high-quality, non-VI hospitals. This reduces hospital competition and increases prices, decreasing consumer surplus and total welfare.

The Chilean private healthcare market provides a well-suited setting for studying the effects of VI. It comprises a handful of hospitals and five insurers competing for individual enrollees in a regulated market. We study the 2013–2016 period, which features a well-established VI segment accounting for half of all hospital capacity and no horizontal mergers, entries, or exits to confound the effects of VI. In addition, insurance plans feature tiered hospital networks, with a base cost-sharing tier covering all inpatient care and a preferential tier covering a select set of hospitals. Tiered networks and access regulation imply that VI hospitals admit patients from all insurers, and access to care is mediated only

through cost-sharing and prices. Finally, there are no out-of-pocket maximums, exposing consumers to hospital prices at the margin and incentivizing them to shop for care.

The data show 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, relative to rival plans. Moreover, VI hospitals charge 13 percent lower prices to their integrated insurer's enrollees, conditional on diagnosis, demographics, and complexity. In contrast, we find no evidence that VI firms limit hospital spending, affect treatment decisions, or impact quality. We find suggestive evidence of demand steering: All enrollees can access all hospitals, but those who switch to VI insurers' plans spend nearly 30 percent more at integrated hospitals. 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 findings 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 coverage levels and hospital tiers. 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 an insurance plan. Finally, health risk is realized in the fourth stage, and consumers choose hospitals influenced by their plan's coverage and prices. This model bridges those in the literature on insurer-hospital competition (Ho, 2009; Gowrisankaran *et al.*, 2015; Ho and Lee, 2017, 2019) and VI (Lee, 2013; Crawford *et al.*, 2018). Its main novelty is incorporating endogenous 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—partially foreclosing them—may lead enrollees to switch to their integrated insurer. Third, VI insurers internalize that higher prices or lower coverage at rival hospitals might steer demand toward their hospitals. Finally, VI insurers internalize the value of their plans in attracting demand to their hospitals. These forces incentivize VI firms to charge lower prices internally, higher prices to rival insurers, skew plan coverage in favor of integrated hospitals, and set lower premiums on plans that steer demand to integrated partners. Non-VI rival insurers and hospitals do not face these incentives directly but compete in their presence.

We use our data to identify the primitives governing the impacts of VI. We identify

consumer hospital preferences by examining how they 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 that address price endogeneity in demand estimation. Moreover, we leverage the dissolution of a VI firm to identify how forces beyond price and quality (e.g., marketing) steer demand. Enrollment data informs of willingness to pay for heterogeneous plans, and risk variation identifies preferences for insurance. We use data on hospital average costs to identify the supply-side primitives, proving that cost information is crucial for identification. Finally, we use optimality conditions associated with plan tiering choices to formulate moment inequalities and identify plan design costs.

We estimate that consumers are significantly more elastic to premiums than out-of-pocket hospital prices. This limits VI hospitals' ability to harm rival insurers through higher prices, as lower premiums easily mitigate demand losses. Consumers, however, are sufficiently elastic to out-of-pocket prices for coverage to be an effective demand-steering instrument. This leads to the misallocation of nearly half of all medical spending relative to first-best, largely to the benefit of VI hospitals. Finally, consumers have heterogeneous network values driven by the match between their medical risk and hospital diagnosis-specific quality. This reduces the attractiveness of narrow networks and incentivizes VI firms to include other hospitals in their preferential tiers.

Using our model, we quantify the effects of banning VI. In the short run, when prices and premiums adjust, but plan design remains fixed, the impact on average prices reflects standard theoretical predictions. Double marginalization is reintroduced, raising prices among formerly integrated partners. Incentives to increase rivals' costs disappear, lowering prices between formerly VI hospitals and rival insurers. Enrollees of formerly VI insurers seek more care at non-VI hospitals, shifting countervailing power and increasing those hospitals' prices. Facing higher prices and no incentive to attract demand to their integrated hospitals, formerly VI insurers raise premiums. Non-VI insurers' ability to profit from lower prices at VI hospitals is limited since their plans' coverage is skewed away from them. Overall, consumer surplus and total welfare fall. Thus, in the short run, VI increases welfare.

It is only when insurers redesign their plans that the full effect of VI unfolds. Formerly VI insurers nearly halve preferential access to former partner hospitals, partially replacing it with access to high-quality non-VI hospitals. Network changes also reallocate consumers among formerly VI hospitals. Demand is shifted to lower-priced alternatives, reducing the average price paid by enrollees of formerly VI insurers for services at former partner hospitals. Non-VI insurers respond by slashing preferential access across the

board, but less so at high-quality hospitals. Access to formerly-VI and non-VI high-quality hospitals increases, stimulating competition among them. Overall, hospital spending falls by 6.4 percent, premiums by 2 percent, and total household spending by 2 percent. Consumer surplus increases by 62.7 million dollars and total welfare by 41.7 million per year. These gains accrue primarily to non-VI insurers and their enrollees, with consumers gaining an equivalent of 5 monthly premiums in surplus and insurers a 19 percent in profits. Thus, in the medium run, VI decreases welfare.

Overall, VI significantly changes the competitive landscape. Beyond the elimination of double marginalization and the incentive to increase rivals' costs, our results show VI meaningfully affects plan design. In the Chilean setting, we find that VI harms market efficiency by siloing VI hospitals from competitors through preferential networks and distorting coverage generosity. To inform the broader discussion on insurer-hospital integration, we explore the factors driving the welfare effects of VI. First, we show that cost efficiencies or quality improvements would make VI less detrimental—though the magnitudes of those impacts would need to be much larger than what the current evidence from the U.S. finds to overturn the results. Second, we show VI harms more consumers when high-quality hospitals are integrated, which limits valuable access to quality care. Finally, we show VI is more likely to harm welfare when consumers are more sensitive to premiums than to out-of-pocket hospital prices.

Our paper contributes to three strands of the literature. First, we contribute to the work 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; Ho and Lee, 2017; Chorniy *et al.*, 2020). 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 research fills a gap in the literature on the growing role of insurer-provider integration. We contribute a framework to study the equilibrium impacts of VI and empirical evidence from a compelling setting. Our work complements previous evidence comparing VI and non-VI plans' costs and quality (Johnson *et al.*, 2017; Parekh *et al.*, 2018; Park *et al.*, 2023).¹

Second, we contribute to the empirical literature on vertical contracting (Lee *et al.*, 2021). We bridge the literature on VI (Hortaçsu and Syverson, 2007; Atalay *et al.*, 2014; Crawford *et al.*, 2018; Luco and Marshall, 2020; Chen *et al.*, 2024) and healthcare com-

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

petition. We bring to bear features of healthcare markets, including cost-sharing, moral hazard, and upstream choice. Closely related, Crawford *et al.* (2018) finds that VI in television markets generates foreclosure incentives. Consumers in our market are more elastic to premiums than hospital prices, and insurers face strong competition from a public option, making foreclosure less profitable. Instead, the environment induces firms to steer demand using upstream prices and plan design. This second feature is a contribution of our framework, which connects to the 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).

Finally, our work speaks to the literature on insurance contract design (Chade *et al.*, 2022; Marone and Sabety, 2022; Tilipman, 2022; Ho and Lee, 2023; Ghili *et al.*, 2024). We build on this literature, which focuses primarily on single-agent perspectives of optimal design, be it a regulator, monopolist, or employer. The literature highlights risk protection, moral hazard, and adverse selection as key determinants. In contrast, we focus on plan design under oligopolistic competition. The interaction between VI, coverage regulation, and affordable public care in our setting makes demand steering and competition over network value the key determinants of plan design. Finally, tiering decisions play a central role in this competition, constituting a static network design problem (Lee and Fong, 2013; Shepard, 2022; Serna, 2023). We see these approaches as complementary, shedding light on different aspects of how contracts are designed.

The remainder of the paper is organized as follows. Section 2 describes our setting and data. Section 3 provides descriptive evidence for how VI shapes market outcomes. Section 4 describes our model of competition in healthcare. Section 5 discusses the identification and estimation of the model and the main results from the estimates. Section 6 discusses our counterfactual analysis of the impacts of banning VI. Finally, Section 7 concludes.

2 Institutions and Data

2.1 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 taxable income by workers and retirees.² Private insurance is provided by a small number of firms (*Instituciones de Salud Previsional*,

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

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 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, gender, and the number of dependents. They are required to offer plans with tiered networks: Enrollees either face the same cost-sharing at all hospitals or have a subset of preferential hospitals where coverage is more generous.³ Insurers must disclose plan cost-sharing and tiers to potential enrollees. Regardless of plan tiers, insurers must afford access to all large general acute care hospitals, even if at varying cost-sharing levels. There are no out-of-pocket maximums, implying enrollees remain exposed to hospital prices regardless of spending. This, along with the negligible presence of deductibles, means cost-sharing is determined almost exclusively by plan coverage (i.e., one minus the coinsurance rate). Finally, guaranteed renewability allows enrollees to re-enroll in their previous plans under certain conditions (Atal, 2019; Figueroa, 2023). For more details on the market regulation, see 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 .⁴ The remainder of the private market is served by seven *closed* insurers associated with a few large companies and who 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 Chile, 2012). Public hospitals provide care at highly subsidized and regulated prices, while private hospitals offer higher quality at a 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.

³Single-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.

⁴One of these insurers operates using two brands. We account for this feature by allowing for heterogeneity in preferences over the services of these brands but otherwise treat them as a single firm.

Figure 1: 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.

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).⁵ As illustrated in Figure 1, 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 mergers, entries, or exits among them during it. Importantly, VI hospitals accept patients from all insurers. Finally, note that star hospitals are not VI.

2.2 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

⁵Chilean 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).

construction, and its connection with the regulatory environment in Appendix A.⁶

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 consumers' out-of-pocket share. We observe detailed diagnoses, services 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.

2.2.1 Descriptive Statistics. We display summary statistics in Appendix Table A.1. The average private plan policyholder 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. Plan enrollment is skewed: while insurers offer 1,431 plans in the market, 123 plans account for 90 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 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

⁶We 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.

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.⁷ Appendix A.3 further describes the market structure and the interaction between insurers and hospitals.

3 Descriptive Evidence

To describe how VI shapes market outcomes, we begin by exploring how insurers might affect the market through their plan networks. In particular, VI insurers may steer hospital demand by offering higher coverage to their enrollees at integrated hospitals. To study this behavior, we describe the coverage offered by plans at each hospital by estimating:

$$y_{jht} = \beta VI_{m(j)ht} + \eta_{m(j)} + \zeta_h + \rho_t + \varepsilon_{jht} \quad (1)$$

where y_{jht} is a network attribute of plan j at hospital h in year t , and $VI_{m(j)ht}$ indicates whether $m(j)$, the insurer offering plan j , is VI with hospital h . The regression includes insurer, hospital, and year fixed effects. The results in columns (1) and (2) in Table 1 suggest 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. To the extent patients are responsive to hospital prices, these patterns suggest VI insurers skew coverage toward their hospitals to steer demand to them.

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:

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

where y_{idjht} is an outcome for admission i for diagnosis 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, \mathbf{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.

⁷Appendix 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)
	A - Plan design		B - Admission outcomes				C - Hospital outcomes	
	Preferential hospital	Coverage rate	log cost proxy	log # services	Re-admission	log price	Share of admissions	log revenue
VI	0.336 (0.049)	5.843 (1.009)	0.029 (0.026)	-0.034 (0.022)	0.013 (0.006)	-0.137 (0.022)	0.382 (0.025)	1.875 (0.172)
<i>N</i>	15,741	15,741	567,752	567,752	204,223	567,752	264	264
<i>R</i> ²	0.150	0.240	0.208	0.612	0.050	0.692	0.555	0.788
Mean non-VI	0.137	63.101	2.534	16.039	0.081	5.269	0.128	8,865.659
Hospital FE	Y	Y	N	N	N	N	Y	Y
Insurer FE	Y	Y	N	N	N	N	Y	Y
Year FE	Y	Y	N	N	N	N	Y	Y
Interacted FE	N	N	Y	Y	Y	Y	N	N
Plan FE	N	N	Y	Y	Y	Y	N	N
Cost proxy	N	N	N	Y	Y	Y	N	N
Controls	N	N	Y	Y	Y	Y	N	N
Observation	plan-hospital-year		admission				insurer-hospital-year	

Notes: This table shows estimates from equations (1), (2) and (3). The unit of observation is reported in the bottom row. 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.

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 comparisons of admission outcomes within hospital-diagnosis-year across patients from an insurer that is VI with the hospital and those who are not.⁸

A proposed advantage of VI is an increased scope for insurers to control hospital spending. We explore this by estimating equation (2) using the cost proxy and number of services provided within an admission as dependent variables. Our 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 VI induces cost-control or treatment changes. Columns (3) and (4) of Table 1 show no significant differences in these outcomes. In addition, we explore whether VI is associated with treatment quality by estimating the same regression using an indicator for 30-day readmission as the dependent variable.⁹ Column (5) shows no

⁸We study whether patients from VI and non-VI insurers are balanced in observables in Appendix A.3.5.

⁹We 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

evidence of quality improvements. Taken together, these results suggest that cost efficiencies, treatment choices, and quality improvements associated with VI might not be present in our setting.^{10,11} 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. This pattern is consistent with the elimination of double marginalization lowering prices within VI firms, or VI hospitals' incentive to negotiate higher prices with rival insurers.

The results above show VI insurers skew coverage in favor of their hospitals, and that these charge lower prices for their enrollees. To quantify whether these behaviors are effective at steering demand toward their hospitals, we estimate:

$$y_{mht} = \beta VI_{mht} + \eta_m + \zeta_h + \rho_t + \varepsilon_{mht} \quad (3)$$

using admission flows from insurer m to hospital h as the dependent variable. Column (7) in Table 1 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 than from rivals.

These results suggest VI insurers successfully steer their enrollees to their hospitals using coverage and prices. We complement these results with an analysis that exploits variation in VI within insurer-hospital over time. This variation is produced by insurer m_c selling its stake in hospital h_{10} in 2014. Appendix Table A.4 shows results from estimating equations (1), (2) and (3) with insurer-hospital fixed effects. The results are qualitatively similar and reinforce our findings.

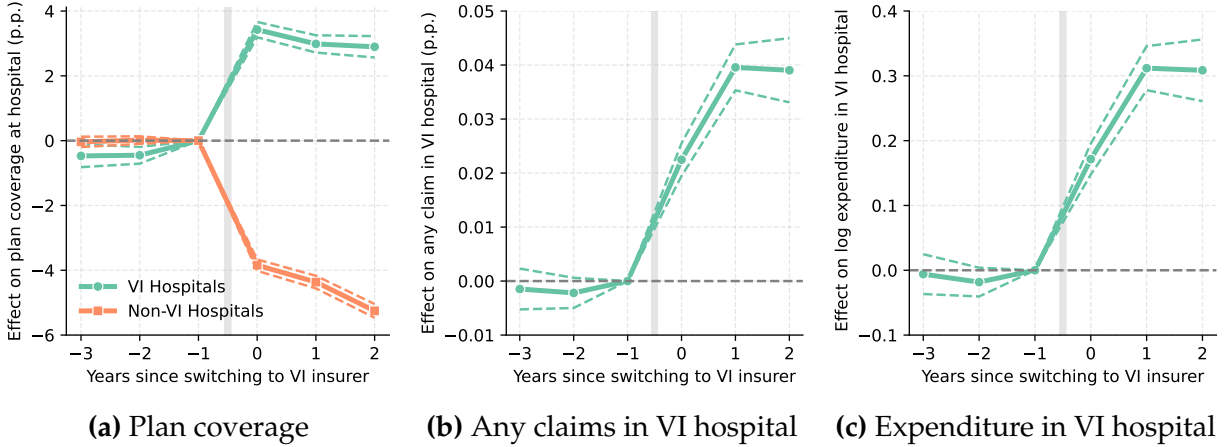
To further illustrate how effective VI insurers are at steering demand toward integrated partners, we study hospital choices and spending of enrollees who switch to VI insurers. Whereas switchers' hospital choice sets are constant over time, whether a hospital is VI with their insurer may change upon switching. Thus, switchers experience changes in

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.

¹⁰We find a similar pattern of results when examining discretionary procedures in Appendix A.4.

¹¹Our empirical strategy prevents us from ruling out that VI increases quality for all patients within a hospital. However, a common argument for quality gains from VI is an improvement in care coordination, which could not be delivered for patients of non-VI insurers.

Figure 2: Vertical integration, hospital choices, and expenditure



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 2c is $\log(1 + y)$ to accommodate zeros, but the results are similar when using expenditure in levels.

out-of-pocket prices but not in access. We estimate the following event study regression:

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

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 2a 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 2b 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 2c shows expenditure in hospitals integrated with the

new insurer increases by around 30 percent after the enrollee switches to a VI insurer.¹²

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.

4 Model

We model the market as the subgame perfect equilibrium of a four-stage game. First, insurers design plans by choosing preferential and base coverage and a set of preferential hospitals for each plan. Second, non-VI insurers and hospitals negotiate hospital prices, while VI firms set prices internally to maximize profits. Insurers set plan premiums simultaneously to maximize profits. Third, consumers choose a plan according to their preferences and risk. Fourth, health risks are realized, and consumers choose hospitals according to their preferences and coverage. We describe each stage in reverse order. We discuss our model assumptions and their connection to the setting in Appendix B.1.

4.1 Demand for Healthcare

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

$$u_{ihdt|j}^H = \alpha_i^H (1 - c_{jht}) \omega_{id} p_{m(j)ht} + \beta^H \text{distance}_{ih} + \gamma^H \text{VI}_{m(j)ht} + \chi_{hdt}^H + \xi_{m(j)hdt}^H + \epsilon_{ihdt|j}^H \quad (5)$$

where the first term is the disutility from out-of-pocket spending, which depends on price-sensitivity α_i^H , plan coverage c_{jht} , the resource intensity weight for condition d for a patient with i 's characteristics, ω_{id} , and the price negotiated between the insurer and the hospital, $p_{m(j)ht}$. The second term is the disutility from travel, and the third captures the

¹²A potential concern is that patients may switch to a VI insurer to gain better access to its 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 cannot rule out contemporaneous health shocks inducing enrollees to switch insurers. For robustness, we repeat the analysis in a subsample of enrollees that move across neighborhoods. These moves may induce enrollees to switch insurers, though they are unlikely to be driven by health shocks that drive hospital demand. The results are in Appendix Figure A.2 and are similar to those in the full sample.

impacts VI may have in steering demand to integrated hospitals (beyond coverage and prices). Moreover, χ_{hdt}^H captures hospital quality for condition d , and $\xi_{m(j)hdt}^H$ is a systematic preference shock. Finally, ϵ_{ihdtj}^H is an iid idiosyncratic shock distributed T1EV.

The outside option is to visit a public hospital.¹³ Given the wide availability of public providers, we assume consumers can always find one in their neighborhood. Public hospitals are in the non-preferential tier of private plans and are paid public list prices, p_{0dt} . Therefore, the indirect utility from this option is $u_{i0dtj} = \alpha_i^H(1 - \underline{c}_{jt})p_{0dt} + \eta_{m(j)dt}^H + \epsilon_{i0dtj}^H$, where \underline{c}_{jt} is plan j 's non-preferential coverage and $\eta_{m(j)dt}^H$ captures systematic preferences for the outside option relative to private hospitals. Consumer i 's probability of visiting hospital h for diagnoses d is $D_{ihdtj}^H = \exp(\delta_{ihdtj}^H) / \sum_{h' \in H} \exp(\delta_{ih'dtj}^H)$, where $\delta_{ihdtj}^H = u_{ihdtj}^H - \epsilon_{ihdtj}^H$.¹⁴

It is worth connecting this model to our broader analysis. First, consumer heterogeneous distaste for out-of-pocket spending impacts the variety of coverage offered in the market, and its correlation with medical risk affects insurers' ability to engage in selection via contract design. Second, perceived hospital quality governs how households of different risks value access to distinct networks. For example, households with children might value generous coverage for high-quality pediatric care, while older families might prefer plans with generous access to cardiovascular care. Therefore, quality shapes insurer plan tiering and network variety. Finally, preferences for VI hospitals capture VI firms' efforts to steer demand towards their hospitals using marketing. We treat this term as welfare-irrelevant, as the evidence does not suggest differential quality at VI hospitals for integrated enrollees. We discuss the implications of this assumption in Section 6.4.

4.2 Demand for Insurance

Consumers buy insurance at the household level, represented by a policyholder $i \in F$, where F is the household or family unit.¹⁵ Each household comprises one or more individuals and belongs to a market segment.¹⁶ For example, single males aged 25–40 are a separate segment from women with dependents aged 40–55. Each year, insurers

¹³The outside option groups public and small private hospitals. We identify the outside option with the public system and its prices, as the excluded small private hospitals are less relevant for complex inpatient stays, which constitute the bulk of spending in our analysis. Excluded private hospitals are not VI.

¹⁴We assume away hospital capacity constraints. To our knowledge, private capacity in Santiago is not binding. Public capacity constraints lead to wait times captured in relative preferences for public care.

¹⁵Formally, F is an element of a partition \mathcal{F} of \bar{I} . A singleton F captures consumers without dependents.

¹⁶Insurers effectively offer plans and prices by market segments defined by the policyholder's age group, gender, and whether they have dependents or not. We define age groups as 25 to 40, 40 to 55, and 55 to 65, and follow this market segmentation in our analysis. For additional details, see Appendix B.1.

provide a collection of insurance plans to each segment, characterized by a premium ϕ_{jt} and a coverage structure \mathbf{c}_{jt} that describes cost-sharing at each hospital. Policyholders choose a plan to maximize the indirect utility of their household, given by:

$$u_{ijt}^M = \alpha_i^M |F_i| \phi_{jt} + \beta^M \sum_{i' \in F_i} WTP_{i'jt}(\mathbf{p}_{m(j)t}, \mathbf{c}_{jt}) + \chi_{im(j)}^M + \xi_{jt}^M + \epsilon_{ijt}^M \quad (6)$$

where the first term captures households' disutility from premiums, governed by premium-sensitivity α_i^M , household size $|F_i|$, and per-enrollee premiums ϕ_{jt} . The second term is the household's expected network surplus from plan j , $WTP_{ijt} = \sum_{d \in D} r_{id} \ln \sum_{h \in H} \delta_{ihdt|j}^H$, which sums across household members and where r_{id} is consumer i 's annual risk of condition d (Capps *et al.*, 2003). The third term, $\chi_{im(j)}^M$, is an insurer-age fixed effect, capturing heterogeneous preferences for insurers. Finally, ξ_{jt}^M is a systematic unobserved preference for plan j in year t , and ϵ_{ijt}^M is an iid idiosyncratic preference shock distributed T1EV.

Households can opt for public coverage instead of private insurance. Under this program, they pay an income-dependent premium ϕ_{0it} and get access exclusively to the public hospital system under public prices. We normalize the household's expected utility from the public plan's network to zero, which sets the level for private plan preferences. Thus, the indirect utility of public insurance is $u_{i0t} = \alpha_i \phi_{0it} + \epsilon_{i0t}^M$ and the choice probability of plan j by consumer i 's household is $D_{ijt}^M = \exp(\delta_{ijt}^M) / \sum_{j' \in J_{st} \cup \{0\}} \exp(\delta_{ij't}^M)$, where $\delta_{ijt}^M = u_{ijt}^M - \epsilon_{ijt}^M$ and J_{st} is the set of private plans for segment s in year t .

Preferences over prices, premiums, and networks are central to our analysis. A high sensitivity to premiums relative to hospital prices limits VI firms' ability to attract enrollees by worsening rivals' networks. Non-VI insurers could react to higher prices at VI hospitals by lowering premiums, mitigating their losses, hence capping VI firms' benefits from foreclosure practices. Thus, α_i^H , β^M , and α_i^M are key in shaping the welfare effects of VI.

4.3 Insurance Premium Competition

Insurers simultaneously set plan premiums to maximize expected profits. Given premiums ϕ_t , prices \mathbf{p}_t , and coverages \mathbf{c}_t , insurer m 's expected profits are:

$$\pi_{mt}^M(\phi_t, \mathbf{p}_t, \mathbf{c}_t) = \sum_{j \in J_{mt}} \sum_{i \in I} D_{ijt}^M(\phi_t, \mathbf{p}_t, \mathbf{c}_t) (|F_i| \phi_{jt} - \sum_{i' \in F_i} \sum_{d \in D} r_{i'd} \sum_{h \in H} D_{i'hdt|j}^H(\mathbf{p}_{mt}, \mathbf{c}_{jt}) c_{jht} \omega_{i'd} p_{mht} - \eta_{jt}) \quad (7)$$

where J_{mt} is the set of plans insurer m offers in year t , and I denotes the policyholders who determine household enrollment decisions. Per household, insurer m earns revenue

equal to collected premiums and faces a cost equal to the expected share of payments plus an administrative burden η_{jt} . There is no risk adjustment, and household willingness to pay for insurance depends on their risk, creating a potential adverse selection problem. In our results below, this feature results in the underprovision of private insurance.

Non-VI insurers set premiums to maximize profits, π_{mt}^M . In contrast, VI insurers also consider their premiums' impacts on integrated hospitals. Hospital profits are given by:

$$\pi_{ht}^H(\phi_t, \mathbf{p}_t, \mathbf{c}_t) = \sum_{j \in I_t} \sum_{i \in I} D_{ijt}^M(\phi_t, \mathbf{p}_t, \mathbf{c}_t) \sum_{i' \in F_i} \sum_{d \in D} r_{i'd} D_{i'hd}^H(\mathbf{p}_{m(j)t}, \mathbf{c}_{jt}) \omega_{i'd} (p_{m(j)ht} - k_{hm(j)t}) \quad (8)$$

which combines the demand from all plans with the likelihood that their enrollees end up at hospital h . Each diagnosis has resource intensity ω_{id} , and the hospital obtains a payment of $p_{m(j)ht}$ and incurs a cost $k_{hm(j)t}$, per unit of resources. Hospital costs may vary across insurers depending on how easy it is to submit claims and get reimbursements.

VI firms set premiums by maximizing a weighted sum of insurer profits π_{mt}^M and any integrated hospital's profits π_{ht}^H . Without loss, we normalize the weight placed on integrated insurer profits to one and express the objective of an integrated insurer m as:

$$\tilde{\pi}_{mt}(\phi_t, \mathbf{p}_t, \mathbf{c}_t) = \pi_{mt}^M(\phi_t, \mathbf{p}_t, \mathbf{c}_t) + \sum_{h \in H} \theta_{hmt} \pi_{ht}^H(\phi_t, \mathbf{p}_t, \mathbf{c}_t) \quad (9)$$

where, for example, $\theta_{hmt} = 1$ if m and h are VI, and they value a dollar in hospital profits equally to a dollar in insurance profits. If, instead, profits are worth 50 percent more at the hospital than at the insurer, then $\theta_{hmt} = 1.5$. Such imbalance can be caused by differential regulatory scrutiny or unequal power within the VI firm. Finally, if m and h operate as distinct firms despite being co-owned, then $\theta_{hmt} = 0$. We interpret $\tilde{\pi}_m$ as the firm's objective but not necessarily its profit. This way of modeling the internalization of profits across integrated partners is akin to that in Crawford *et al.* (2018).

We use equation (9) to capture the objective of VI and non-VI firms. For a VI hospital h integrated with an insurer m , we use the shorthand $\tilde{\pi}_h$ to denote $\tilde{\pi}_{mt}$. Analogously, if m is not VI, $\tilde{\pi}_{mt} = \pi_{mt}^M$ by definition, and if h is not VI, then $\tilde{\pi}_{ht} = \pi_{ht}^H$. If a hospital is not integrated and is part of a system of hospitals (as in our counterfactuals), then $\tilde{\pi}_{ht}$ includes the profits of all hospitals in the system. Equilibrium premiums ϕ^* thus solve:

$$\max_{\phi_{mt}} \tilde{\pi}_{mt}(\phi_t, \mathbf{p}_t, \mathbf{c}_t) \quad \forall m \in M, t \quad (10)$$

VI shapes equilibrium premiums. In particular, VI insurers offer plans providing pref-

erential access to integrated hospitals and internalize profits earned by steering demand to them. This creates an incentive to reduce premiums, and as premiums are strategic complements, this places downward competitive pressure on rival insurer premiums. The magnitude of this effect depends on consumers' relative sensitivity to out-of-pocket hospital spending and premiums. These incentives also alter plan design choices and hence access to care altogether, as we discuss in Section 6.

4.4 Hospital Price Negotiations

Insurers and hospitals determine prices simultaneously with premiums. VI firms set prices internally to maximize their weighted objective, $\tilde{\tau}_{mt}$. Non-VI pairs, instead, engage in Nash bargaining (Crawford and Yurukoglu, 2012; Collard-Wexler *et al.*, 2019) by solving:¹⁷

$$\max_{p_{hmt}} (\Delta_{mh}\tilde{\tau}_{mt}(p_{hmt}) + l_{mht}\Delta_{mh}WTP_{mt})^{\tau_h} (\Delta_{mh}\tilde{\tau}_{ht}(p_{hmt}))^{1-\tau_h} \quad (11)$$

where τ_h is insurer m 's bargaining weight when negotiating with hospital h , and $\Delta_{mh}\tilde{\tau}_{mt}$ is the incremental value it derives when hospital h is available for its enrollees. Analogously, $\Delta_{mh}\tilde{\tau}_{ht}$ is the incremental value hospital h obtains if enrollees of insurer m can access its services. If h is part of a system, it threatens to deny access to m 's enrollees to all of its system's hospitals upon disagreement. This approach builds on related work on hospital-insurer bargaining (Gowrisankaran *et al.*, 2015; Ho and Lee, 2017).

The legal penalty $l_{mht}\Delta_{mh}WTP_{mt}$ is a fine levied on insurers in the event of disagreement. It captures enrollees' right to maintain access to providers, an integral part of the Chilean regulatory environment.¹⁸ We model the penalty as proportional to the expected loss in network surplus across all the insurer's plans, $\Delta_{mh}WTP_{mt}$, where l_{mht} is a random variable capturing expectations about the severity of the legal case drawn independently across negotiations from a distribution L .¹⁹ Appendix B.2 provides further details.

The distinction between the pricing protocol of VI and non-VI pairs captures the essence of VI. When setting prices, VI hospitals internalize the value of attracting enrollees to their integrated insurer, even if they then seek care at other hospitals. This eliminates double marginalization by aligning incentives (Spengler, 1950) and distorts how firms

¹⁷Hospitals and insurers hold rational expectations and passive beliefs about all other ongoing negotiations as in Horn and Wolinsky (1988).

¹⁸These access rights form part of consumers' right to guaranteed renewability (Atal, 2019).

¹⁹This can be thought of as each negotiating pair observing a common iid signal about the severity of the legal case if they disagree. The multiplier, l_{mht} , captures their conditional expectations.

negotiate prices with rivals. VI hospitals internalize that a higher price to a rival insurer reduces rival plans' network surplus, shifting demand to their integrated insurer. This induces VI hospitals to foreclose or worsen their rivals' networks (Hart and Tirole, 1990). Analogously, VI insurers acknowledge that higher prices from rival hospitals shift demand toward integrated hospitals, decreasing their incentive to negotiate prices down.

4.5 Plan Design

In the first stage, insurers design plan menus. For each plan j , insurer m must select a base coverage level $\underline{c}_j \in [0, 1]$, a preferential coverage level $\bar{c}_j \in [0, 1]$, and a set of preferential hospitals $\bar{H}_j \subset \{H \cup \emptyset\}$. Denoting the set of all tiered designs as \mathcal{C} , each insurer solves:²⁰

$$\max_{\{c_j \in \mathcal{C}\}_{j \in J_{mt}}} \tilde{\pi}_{mt}(\phi^*(c), p^*(c), c) - \sum_{j \in J_{mt}} M_{jt} (K_m^r(c_j) + K_m^o(c_j)) \quad (12)$$

where the first term is the firm's profit objective—inclusive of the impact on equilibrium premiums and prices—and the second term is the firm's underwriting cost, which scales with market size M_{jt} and comprises regulatory costs $K_m^r(c_j)$ and organizational costs $K_m^o(c_j)$.

Regulatory costs $K_m^r(c_j)$ capture the shadow cost of regulatory scrutiny. As this might entail costs beyond plan denial, we model it as a flexible component to be estimated.²¹ Organizational costs $K_m^o(c_j)$ capture tiering costs. These include marketing costs associated with offering preferential plans and organizational pressures to tier integrated hospitals beyond their effect on profits.²² Overall, underwriting costs represent the value of relaxing implicit regulatory constraints and the cost associated with transfers among players in our model. They cancel out in our welfare analysis but help rationalize minor deviations from profit-maximizing plan design and capture a complex regulatory environment.

Given the game's timing, consumer unobserved plan preferences are unknown at the plan design stage. Insurers know ξ_{jt}^M only up to consumer systematic preferences for each insurer-year, with the residual variation realizing once plans are designed and approved. We assume insurers evaluate this uncertainty using rational expectations, setting them to their expected mean of zero. In contrast, insurers are aware of consumer hospital preferences, which are not plan-specific. Insurers are also informed of the realizations

²⁰Formally, $\mathcal{C} = \{c \in [0, 1]^{|H|} \mid \exists 0 < \underline{c} \leq \bar{c} < 1, \bar{H} \subset H \cup \emptyset \text{ s.t. } c_h = \bar{c} \mathbb{1}\{h \in \bar{H}\} + \underline{c} \mathbb{1}\{h \notin \bar{H}\}\}$.

²¹The regulator overlooks plans as described in Appendix A.2. The threat of scrutiny is one mechanism through which the regulator might alleviate the effects of adverse selection, given the lack of risk adjustment.

²²Organizational pressures to tier hospitals include efforts to build brand loyalty or to signal cooperation between the managers of integrated insurers and hospitals.

of regulatory penalty multipliers l_{mht} and hospital costs. They observe rival insurers' approved plans and can forecast their subsequent pricing and premium choices.

Endowed with their information and facing the pricing subgame, insurers design plans to optimize plan value, steer demand, and control the allocative efficiency of spending. Consumers prefer plans with higher coverage. However, increasing coverage lowers hospital demand elasticity, leading to higher negotiated prices and insurer costs. Higher coverage at a hospital also steers demand towards it. In particular, insurers are concerned with spending valued below the cost of coverage, as it entails payments that cannot be fully recovered through premiums. To control this loss, insurers use coverage and prices to limit consumer moral hazard spending. Importantly, when considering spending at integrated hospitals, VI insurers share the regulator's perspective on spending efficiency, benchmarking it against costs rather than prices. However, they value spending at rival hospitals less than the regulator. Therefore, VI insurers overprovide coverage to their hospitals, rationalizing why almost all preferential plans of VI firms include at least one of their hospitals as preferential. For non-VI firms, competing against VI plans that overprovide coverage has more nuanced consequences, which depend on whether coverage is a strategic complement or substitute on the margin.

4.5.1 Solving the Plan Design Problem. Plan design involves a large-scale combinatorial problem, as all plans interact in hospital price negotiations. It is analogous to forming an extensive collection of interdependent networks, as tiering decisions are akin to selecting which hospitals are in the network. The following result, proven in Appendix B.3, is key to our study of equilibrium networks.

Proposition 1. *For $\lambda \geq 0$, and a positive and increasing $G(\cdot)$, let $\tilde{c}^*(\lambda)$ be a solution to the problem:*

$$\max_{\substack{\underline{c}_m \leq \tilde{c}_m \in [0,1]^{|J_{mt}|} \\ \mathbf{w}_m \in [0,1]^{|J_{mt}| \times |H|}}} \tilde{\pi}_{mt}(\phi^*(\mathbf{c}), \mathbf{p}^*(\mathbf{c}), \mathbf{c}) - \sum_{j \in J_{mt}} M_{jt}(K_m^r(\mathbf{c}_j) + K_m^o(\mathbf{c}_j)) - \lambda \sum_{j \in J_{mt}} \sum_{h \in H} G(w_{jh}(1 - w_{jh})) \quad (13)$$

where $c_{jh} = \bar{c}_j w_{jh} + (1 - w_{jh}) \underline{c}_j$. Then $\lim_{\lambda \rightarrow \infty} \tilde{c}^*(\lambda)$ is a solution to equation (12).

The proposition states firms are in the same position whether they face a strict requirement to provide tiered networks or increasingly large penalties for not doing so. In equation (13), insurers freely design their plans' coverage which, without loss, consists of a base coverage \underline{c}_j , a preferential coverage \bar{c}_j , and the position of each hospital between the two, $w_{jh} \in [0, 1]$. The term $w_{jh}(1 - w_{jh})$ is the degree to which a design violates the tiering requirement, $G(\cdot)$ the regulatory strictness, and λ the dollar value of the penalty.

The proof proceeds in two steps. First, it shows the combinatorial problem of equation (12) is equivalent to a continuous problem with tiering constraints. Second, it applies the maximum theorem to transform the equivalent problem into a sequence of simpler problems. The latter is crucial for computational applicability as the penalty λ controls the trade-off between exploration and feasibility: a small value permits exploration of non-tiered profitable designs while a large value effectively restricts attention to locally feasible solutions. An increasing penalty sequence allows exploration while guaranteeing feasibility at the limit. The regularizing function $G(\cdot)$ is arbitrary and chosen to improve the objective's concavity, limiting the incidence of local maxima. The approach is not heuristic, guaranteeing a feasible solution to the original problem upon convergence.

To our knowledge, this approach is novel in the study of contract design. However, it builds on well-established methods in combinatorial optimization. Equivalence results like those in the first step of our proof are well-known (Lucidi and Rinaldi, 2010; Yakovlev, 2017), and our regularized convergent approach is a specialization of the method of Murray and Ng (2010), who also document its performance in solving large-scale combinatorial optimization problems. In Appendix C.4, we provide additional evidence showing how our approach benchmarks against solving the problem by brute force.

4.6 Equilibrium

A subgame perfect Nash equilibrium consists of plan coverages \mathbf{c} , prices \mathbf{p} , and premiums ϕ , such that firms maximize their objectives according to equations (12), (11), and (9), and consumers optimally choose plans and hospitals according to equations (6) and (5).

There might be several subgame perfect Nash equilibria. To address this, we adopt a strict refinement, focusing on the unique subgame perfect 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. This refinement requires solving all game stages as best response iterations. In Appendix C, we reformulate the price- and premium-setting subgames as solutions to fixed-point problems that operationalize this strategy.

5 Identification and Estimation

We estimate the model in four steps. First, we leverage hospital choices to recover preferences over hospitals. Second, we use enrollment choices to recover preferences over plans. Third, we use the estimated demand model, data on hospital average costs, and optimality conditions for premium setting and price negotiations to recover the relative weights in VI firms' objectives, hospital bargaining weights and costs, the distribution of legal penalties to insurers in case of bargaining breakdowns, and insurer administrative costs. Fourth, we use the optimality of observed plan designs to recover underwriting costs. The following sections describe the estimation approach, identification arguments, and estimates. Appendix C provides supporting proofs and details.

5.1 Preferences over Hospitals

We estimate preferences for hospitals via maximum likelihood, leveraging individual-level hospital choice data. We use the inversion approach of Berry (1994) to recover our model's rich structure of systematic preferences. In particular, we absorb consumer responsiveness to VI marketing, $\gamma^H VI_{m(j)ht}$, and unobserved preferences, $\lambda_{hdt}^H + \xi_{m(j)hdt}^H - \eta_{m(j)dt}^H$ in an insurer-hospital-diagnosis-year fixed effect, $\bar{\xi}_{mhd}^H$. Our estimator is:

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

where y_{ihdt} is a choice indicator, $D_{ihdt|j}^H$ the model implied individual choice probability, and $D_{hdt|m}^{H*}$ the observed market share of hospital h among enrollees of insurer m for condition d in year t .²³ We recover consumer responsiveness to VI marketing and their unobserved preferences by projecting the estimates of $\bar{\xi}^H$ on its components.

The primary identification challenge stems from potential price endogeneity.²⁴ We address this issue with an instrument based on the prices negotiated by closed insurers with hospitals in our sample. As discussed in Section 2, 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-insurance coverage.

²³We 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.

²⁴Given 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.

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, making the existence of common demand shocks unlikely. Thus, the prices negotiated by closed insurers reflect the preferences of a separate group of consumers but still 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. Given our estimator is non-linear and out-of-pocket prices are individual-specific, we adopt a control function approach (Petrin and Train, 2010). We provide details in Appendix C.1.

Consumers who are older, have higher incomes, or have dependents are less price elastic, as shown by Table 2-A. Single women are also less elastic than single men. The model implies that the median own-price demand elasticity for inpatient care is -0.79 , which is larger than estimated for consumers in the U.S.²⁵ There are three reasons for this higher elasticity in the Chilean market. First, hospitals are clearly differentiated in quality and prices. Second, consumers can obtain price estimates from providers or insurers for planned inpatient care. Finally, consumers always pay for care at the margin since there are no caps on out-of-pocket payments. Therefore, shopping for care is easier in this setting, and consumers have higher stakes in doing so.

The estimates show consumers have a substantial distaste for travel. A 25-year-old single male enrollee would be willing to pay nearly \$60 to reduce his travel distance by a mile. This disutility from travel creates significant preference heterogeneity across consumers for access to different hospitals, stimulating horizontal differentiation in networks. Similarly, consumers perceive providers to be meaningfully differentiated in diagnosis-specific quality. Appendix Figure A.3 shows consumers prefer the two highest-priced non-VI hospitals for almost every medical need. Insurer m_a 's integrated hospitals are of moderately low quality, with only one having high-quality pregnancy, respiratory, and endocrine condition care. In contrast, two of the three hospitals integrated with insurer m_b are above median quality in almost every condition. They are particularly good at 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.

Consumer responsiveness to VI marketing is identified from the termination of the

²⁵For example, Aron-Dine *et al.* (2013) estimate an elasticity of -0.03 , while Prager (2020) estimates -0.16 .

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: Monthly Premium (α_i^M)				
× Age \in [25, 40)	-1.536	(0.010)	-29.382	(0.053)
× Age \in [40, 55)	-1.355	(0.011)	-26.999	(0.052)
× Age \in [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 Marketing (γ^H)	2.323	(0.008)		
Network surplus (β^M)			1.297	(0.006)
<hr/>				
Control function				
A: Price / B: Premium	0.699	(0.010)	7.987	(0.029)
Network surplus			-0.582	(0.008)
<hr/>				
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.7 shows estimates without the control function approach.

VI agreement between insurer m_c and hospital h_{10} in 2014. The estimates in Table 2-A imply that the strength of steering toward VI hospitals is 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 marketing effectively steers consumers towards integrated hospitals despite large quality differentials.²⁶

5.2 Preferences over Plans

We estimate plan demand using individual-level enrollment data, which we complement in two ways. First, we use our hospital demand estimates to compute network surplus

²⁶In principle, we cannot distinguish between steering through welfare-irrelevant or welfare-relevant means. We find support for our assumption in two observations. First, the steering effect seems exceedingly large to be rationalized by quality differentials, more so given we find no meaningful treatment changes in Section 3. Second, if steering were welfare-relevant, it would be in the interest of both m_c and h_{10} to maintain it after their dissolution—it substantially contributes to m_c 's plans network surplus and profits to h_{10} .

($WTP_{ijt}(\mathbf{p}_{m(j)t}, \mathbf{c}_{jt})$) for all plans offered in their market segment. Second, we use our aggregate 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 their network surplus from private plans. We then estimate preferences over plans by maximizing the weighted likelihood of the data, where weights account for the share of publicly insured individuals.

The main identification challenge is the correlation between premiums and unobserved preference shocks. We address this 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. From the game's timing, plan coverage is not a function of unobserved plan preferences, as these have not been realized by the time insurers design plans. The instrument is excluded from unobserved preference shocks since the regulator is unlikely to consider those when setting public hospital prices.

To a lesser degree, unobserved plan preferences might also influence hospital prices.²⁷ 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. These instruments are in the spirit of (Berry *et al.*, 1995). Given the model timing, network choices reflect variation in regulatory and hospital costs, not unobserved preference shocks. As with hospital demand, we adopt a control function approach. Appendix C.1 provides additional details.

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 -2.3.²⁸ Heterogeneity 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,

²⁷The likelihood that prices correlate with unobserved preferences is limited since negotiations aggregate all plans, reducing the effect the demand for any given plan plays in shaping prices.

²⁸For comparison, Ho and Lee (2017) find a premium elasticity between -1.2 and -3, Curto *et al.* (2021) estimate -1.1, and Tebaldi (2024) estimates it to be between -1.3 and -2.

face higher premiums and greater exposure to hospital prices. Overall, households are more sensitive to changes in premiums than hospital prices. Therefore, insurers can mitigate losses from higher negotiated prices—for example, from a VI hospital—by slightly lowering premiums. This curtails VI firms’ benefit from increasing prices to rivals and incentivizes them to attract consumers with lower premiums instead.

5.3 Price and Premium Setting Parameters

Pricing and premium setting depends on five primitives: hospital costs k_{hmt} , bargaining weights τ_h , the relative weight VI firms place on hospital profits θ_{mht} , the distribution of the regulatory penalty for lost access L , and insurer administrative cost η_{jt} . The following proposition establishes the identification of these objects, derived in Appendix B.4.

Proposition 2. *If, for any negotiating pair (h, m) , bargaining is individually rational, $\frac{\partial \pi_h^H}{\partial p_{hm}} \geq 0$ at equilibrium prices, hospital costs k_{hmt} are non-negative and can be decomposed as $k_{hmt} = k_{ht}^H + \tilde{k}_{mt}^H$ and the disagreement penalty multiplier l_{hmt} is drawn iid from L with finite variance, then under rank sufficiency conditions established in Appendix B.4, $(\theta_{hmt}, k_{hmt}, \tau_h, \eta_{jt}, L)$ are identified.*

The proof proceeds in three steps. First, it shows the optimality conditions for premiums and prices of VI firms uniquely resolve the added complexity of VI (i.e., the introduction of unobserved θ_{hmt}). Hence, identifying the model with and without VI is essentially identical. Second, it shows the model without disagreement penalties or restrictions on hospital costs is identified from the optimality of negotiated prices and data on average hospital costs. The latter is necessary; the passthrough of costs to prices reflects hospitals’ bargaining power, providing additional identifying information. Finally, it shows that adding disagreement penalties and restrictions on hospital costs transforms the problem into two simultaneous systems of equations, allowing us to leverage standard results on the identification of non-linear regression models (Amemiya, 1983).

The proof provides valuable insight into the problem. It does not rely on our demand model, requiring minimal conditions on substitution patterns to hold. It also outlines the model’s identification without disagreement penalties or cost restrictions, speaking to the identification of Nash bargaining models broadly. Particularly, it shows that jointly identifying hospital costs and bargaining weights requires external cost data.²⁹ Fundamentally, the result leverages average cost data, the identification of the demand model,

²⁹The proof shows that alternative information on hospital costs could also resolve the identification problem. For example, the cost of certain standard services, labor, or of serving specific insurers could all have been used in lieu of total cost data.

the constant marginal cost assumption, and the general premium- and price-setting protocols. Most other structures can be relaxed. For example, bargaining weights τ_h could vary across time or include interactions with insurer identities. This, however, requires leveraging time series variation for identification, which is limited in our sample.

Having established identification, we proceed to estimation. We impose two simplifying assumptions to close the gap between what is feasible to identify and what is practical in finite samples. First, we assume hospital costs admit a block structure $k_{hmt} = k_h^H + \tilde{k}_{mt}^H$. Second, we assume L is a normal distribution with unknown mean and variance.³⁰ We aggregate consumers by the demographic variables determining their preference heterogeneity and restrict attention to each insurer’s 70 percent most popular plans.³¹ We estimate the model in two iterative steps. Starting from an initial guess of hospital costs, we use the optimality conditions for premiums and VI firm hospital prices to recover administrative costs η_{jt} and VI weights θ_{hmt} . Second, we use costs and weights to maximize the constrained likelihood of hospital prices for non-VI pairs to recover estimates of hospital costs k_{hmt} , bargaining weights τ_h , and the distribution of penalties $L(\mu_l, \sigma_l^2)$. The constraints match the assumptions of Proposition 2. We update the guess of hospital costs and iterate until convergence. We present additional details in Appendix C.2.

Table 3 shows the estimates. Insurers hold slightly more bargaining power than hospitals broadly. However, the highest-quality non-VI hospitals—the 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. The moderate differences in bargaining weights across hospitals suggest most of the variation in negotiated prices is rationalized by differences in costs and gains from trade. The latter include consumer hospital preferences and VI firms’ ability to recapture profit losses from disagreement. Moreover, the estimated distribution of disagreement penalties implies that, on average, insurers expect a penalty equivalent to 133 percent of the consumer surplus loss from disagreement. The estimates also suggest VI firms place substantial weight on their integrated hospitals’ profits. This is consistent with the uneven regulatory environment, as insurers face more scrutiny than hospitals, creating an incentive to tunnel profits to providers as in 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, lim-

³⁰ A negative penalty multiplier may capture that the legal blame upon disagreement falls on the hospital.

³¹ The aggregation of consumers is largely without loss, given the coarse preference structure. The only source of potential imprecision is the slight undersampling of large households. Our results are robust to increasing the sample size. The plans this restriction excludes consist of a long tail of tiny plans; hence, including them does not alter our results. For most insurers, we capture over 90 percent of their enrollment.

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.

ited prescription drug coverage, and administration—and the median plan and hospital markups are 58 and 27 percent, respectively.³²

5.3.1 Selection and Moral Hazard. Two frictions challenge market efficiency: adverse selection and moral hazard. Figure 3a documents adverse selection by showing the correlation between consumers’ willingness to pay (WTP) and expected inpatient cost for each plan (Einav *et al.*, 2010). The correlation is 0.62 overall and 0.44 conditional on market segment and year. Within each market, consumers at the 75th percentile of WTP cost 33 percent more to insure than those at the 25th percentile, on average. Adverse selection operates through two channels: Riskier consumers benefit more from higher plan generosity and also from greater access to high-quality hospitals. Making plan networks homogeneous eliminates the second channel, revealing that 27 percent of

³²Insurers 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, not captured in our data or analysis. It is independent of market structure or plan design and, therefore, secondary to our analysis.

Figure 3: Adverse selection and moral hazard



Notes: Figure 3a shows the correlation between consumers' expected network surplus and their inpatient cost at each plan. Drawn circles bin the horizontal axis into 30 quantiles, plotting the mean of cost within each. Lines show the associated regression on the entire sample. *Raw* (green) is computed without controls, *Residualized* (orange) controls for a market-year fixed effect, and *Homogeneous networks* (blue) recomputes the previous by setting each plan's coverage across hospitals and each hospital's price across insurers equal to their mean. This eliminates heterogeneity across plans in access while preserving heterogeneity in coverage generosity. Figure 3b shows the share of additional spending produced by allocative moral hazard relative to a counterfactual scenario in which consumers choose hospitals according to cost and quality, hence removing the effect of coverage, price negotiations, and VI marketing. *Full* bars show the additional spending in the status quo, while the *Marketing Only* bars include only the effect of VI marketing.

adverse selection is due to selection on networks (Shepard, 2022). Hence, selection likely affects how insurers choose plan generosity and preferential tiers.

Figure 3b shows the effect of allocative moral hazard on total spending. Holding plan choices fixed, allocative efficiency is achieved when consumers choose hospitals based on cost and quality. Status quo allocations are distorted by price negotiations, coverage choices, and VI marketing. This increases total spending by 47 percent, 8.7 percent of which is attributable to VI marketing. This distortion is roughly equal across VI and non-VI insurers. However, VI marketing nearly doubles inefficient spending at VI hospitals and decreases it at non-VI hospitals. Coverage distortions increase spending valued below cost at all private hospitals. Despite these distortions being relatively larger at VI hospitals, non-VI hospitals are more expensive on average and contribute more to moral hazard spending. Nevertheless, VI plays an important role in spending efficiency, which shapes the welfare impacts of VI we examine in Section 6.

5.4 Plan Design Costs

We identify and estimate plan design costs in two steps. First, we identify the regulatory cost, K_m^r , from the coverage rates' optimality conditions. We specify $K_m^r(c_{jt}) = \exp(c^K(c_{jt})) + \underline{c}_{jt}\underline{\zeta}_{jt} + \bar{c}_{jt}\bar{\zeta}_{jt}$, where $c^K(\cdot)$ is a flexible polynomial of the coverage rate of each plan, \underline{c}_{jt} and

\bar{c}_{jt} are the base and preferential plan coverage rates, and ζ_{jt} and $\bar{\zeta}_{jt}$ are mean-zero iid normal shocks with unknown variance. This flexible cost function captures the regulator's concern for plan generosity and its impact on firm coverage choices. The unobserved cost shocks 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 exploit these conditions to estimate the regulatory cost of coverage by maximum likelihood. Appendix C provides details.

We rely on the optimality of tiering choices to identify and estimate the organizational cost of plan design, K_m^o . We specify $K_m^o(c_{jt}) = \sum_{h \in H} w_{hjt}(\vartheta_{hmt} + \varsigma_{hjt})$, 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 set of instruments Z_{hjt}^K , the unobserved cost shocks satisfy $\mathbb{E}[\varsigma_{hjt}|Z_{hjt}^K] = 0$ and $\mathbb{E}[\varsigma_{hjt}|Z_{hjt}^K, w_{hjt}] \leq \bar{\zeta}$ for some known positive $\bar{\zeta}$.³³ Using these assumptions, the optimality of observed tiering implies:

$$[(\Delta_{jht} \tilde{V}_{mt} - M_{jt} \vartheta_{hmt})(1 - w_{hjt}) - \bar{\zeta} w_{hjt}] Z_{hjt}^K \leq 0 \quad (14)$$

$$[(\Delta_{jht} \tilde{V}_{mt} + M_{jt} \vartheta_{hmt}) w_{hjt} - \bar{\zeta} (1 - w_{hjt})] Z_{hjt}^K \leq 0 \quad (15)$$

where $\Delta_{jht} \tilde{V}_{mt}$ is the difference in insurance profit minus regulatory costs when the tiering decision of hospital h in plan j is inverted.³⁴ Equation (14) implies a lower bound on tiering costs ϑ_{hmt} , stating that if insurer m decided to leave hospital h in the base tier of plan j , then tiering costs must have been sufficiently large. Equation (15) states that if plan j chose to make hospital h preferential, tiering costs must have been sufficiently small, placing an upper bound on costs. These optimality conditions form a basis for estimation, which we implement using the test of Chernozhukov *et al.* (2019).

We impose that these moment conditions are satisfied in expectation for each insurer and year. We let the tiering cost vary depending on whether the hospital-insurer pair (h, m) are VI in year t and differ across VI firms. As instruments, we use the average estimated non-price component of healthcare indirect utility of each plan-hospital, aggregated across

³³As 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.

³⁴Formally, $\Delta_{jht} \tilde{V}_{mt} = \tilde{\pi}_{mt}(c_{jt}) - \tilde{\pi}_{mt}(c'_{jt}) - M_{jt}(K_m^r(c_{jt}) - K_m^r(c'_{jt}))$ where the coverage vector c'_{jt} equals c_{jt} for all components except at c'_{jht} , which is inverted relative to c_{jht} : $c'_{jht} = w_{hjt}c_{jt} + (1 - w_{hjt})\bar{c}_{jt}$, where w_{hjt} is the original preferential tier indicator.

Table 4: Estimates of underwriting 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_c	-1.590 (0.874)
Base x Preferential	6.389	(1.254)	h_5	1.423	(0.398)	D - Regulatory cost shock variance	
VI coverage $\times m_c$	16.148	(3.403)	h_6	-4.520	(1.097)	$\ln(\underline{\epsilon})$	1.222 (0.018)
VI coverage $\times m_b$	9.994	(0.786)	h_7	-0.169	(0.406)	$\ln(\bar{\epsilon})$	2.500 (0.018)
VI coverage $\times m_a$	6.527	(0.869)	h_8	-0.855	(0.529)	E - Tiering costs (ϑ)	
			h_9	0.674	(1.578)	Non-VI	[0.120, 0.139]
			h_{10}	1.949	(1.113)	VI $\times m_a$	[-0.066, 0.083]
			h_{11}	-3.885	(0.543)	VI $\times m_b$	[-0.386, 0.009]
						VI $\times m_c$	[-0.061, 0.495]
N	3,256						
Moments	30						

Notes: This table displays estimates of parameters governing underwriting costs. Panels A, B, and C show elements of $c^K(\cdot)$ and enter the regulatory cost (K_m^r) exponentiated. Spread stands for the difference between preferential and base coverage. Mean coverage is averaged within a plan across hospitals. The functional form of $c^K(\cdot)$ also includes market fixed effects, interacted with each plan's coverage spread, which are omitted for brevity. Estimates in Panel C are relative to the normalized effect of insurer m_c . 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 socks ($\bar{\epsilon}$) is set to 2/3 of a million dollars per 100,000 enrollees, which is approximately the interquartile range of $\Delta \tilde{V}_{hmt}$. Identified sets are estimated using the bootstrap method described in Chernozhukov *et al.* (2019).

consumers according to their medical risk. This instrument captures the relative value a plan offers when accessing a hospital due to factors other than coverage generosity, prices, and premiums. It is unlikely to be correlated with changes to the organizational cost of tiering, 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 estimated underwriting costs. Regulatory costs decrease in base coverage, consistent with the regulator's stated preferences for coverage generosity. Lacking risk adjustment, this might capture regulatory efforts to stop the market from unraveling due to cream skinning. Introducing a preferential tier requires paying a regulatory cost at any relevant level. Higher coverage at VI hospitals increases regulatory costs, indicating a regulatory effort to curtail VI's effect on plan design. Panel C shows the regulator does not treat insurers meaningfully differently, and Panel D indicates substantial non-systematic variance in regulatory costs across plans.³⁵ Organizational tiering costs in Panel E indi-

³⁵The estimated regulatory forces are consistent with the regulation described in Appendix A.2.

cate a reasonably precise estimate for the cost of making a non-VI hospital preferential in a plan. As our main counterfactual analysis eliminates VI altogether, this is the key coefficient of interest, which we set to the estimated median. For VI firms, we cannot reject a zero cost of tiering. We can reject that tiering costs are the same for VI and non-VI insurers for all but m_c , whose VI partnership with hospital h_{10} ends in 2014.

While marginal underwriting costs are small relative to insurance profits at an average of 0.6 percent, they capture regulatory and organizational pressures on plan design. Ignoring these costs yields a baseline prediction with lower coverage, narrower networks, and VI plans that are more skewed toward their hospitals. As our main results in Section 6 show VI distorts plan design at the expense of market welfare, these estimates suggest the regulator might be aware of these distortions and attempts to limit them.

6 Equilibrium Effects of Vertical Integration

Using the model, we evaluate the impact of VI by simulating a VI ban. In the counterfactual, the VI joint profit objectives are eliminated ($\theta_{mht} = 0$), forcing formerly integrated partners to negotiate prices, and VI marketing is removed ($\gamma^H = 0$), eliminating VI insurers' ability to steer patients through means other than price and coverage. The counterfactual equilibrium involves redesigned plans, reoptimized premiums, renegotiated hospital prices, and new enrollment and care 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 additional details in Appendix C.

6.1 Hospital Prices, Plan Premiums, and Plan Design

The elimination of VI has a stark effect on formerly integrated firms, as shown in Table 5-A. Consistent with the theory on double marginalization, the average price of formerly VI hospitals to their insurer increases by \$842 per unit of resources or 23.2 percent. VI insurers increase premiums by 9 percent due to higher costs and the elimination of the incentive to attract enrollees with lower premiums and steer them toward their partners. They also redesign their plans, reducing the likelihood that a VI hospital is preferential by a third and increasing the preferential likelihood of other hospitals. In particular, star hospitals are 16 times more likely to be preferential in VI plans post-ban. Changes to plan coverage structure occur primarily along the tiering margin: base coverage is only 6.3 percentage points lower, and preferential coverage is 2.1 percentage points higher.

Changes in VI plans' preferential access stimulate competition on both sides of the market. VI hospitals are no longer siloed by preferential access from their partners, and non-VI insurers are no longer dissuaded from providing preferential access to VI hospitals by the foreclosure incentives of the latter. As noted by the literature on countervailing power (Ho and Lee, 2017), the net effect of these forces is ambiguous. Intensified hospital competition reduces their bargaining leverage, lowering prices. Similarly, increased competition among insurers makes them more easily substitutable from hospitals' perspective, increasing prices. Column (3) of Table 5 shows that for the average hospital, insurer competition dominates, resulting in higher prices. Changes in plan preferential networks reported in the table and additional evidence shown in Appendix Figure A.4 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 centric hospital (h_{11}). Demand and prices for these providers increase. The similarities between VI and non-VI plan coverage responses to the new environment suggest plan generosity, much like prices, is a strategic complement.

While the average hospital price increases, column (4) in Table 5 shows that the average price consumers pay decreases. Moreover, the average paid price in the outside option falls despite list prices staying constant. The mechanics of demand steering underlie these changes: While more care is delivered at public hospitals, as shown in Table 6, it is of lower resource usage and, hence, of lower complexity. Increased prices at non-VI hospitals push some consumers to the outside option. However, improved access to expensive high-quality and centric cheap hospitals implies that primarily low-complexity care at the city periphery is reallocated to the outside option. Table 6 confirms this improves efficiency, eliminating 37 percent of moral hazard spending. Price increases at star hospitals are offset by this reallocation of low-value care to the public system and a significant decrease in prices at high-quality VI providers. Thus, while the average VI hospital increases prices, consumers are steered away from it. Instead, they are redirected to high-quality VI hospitals, which now face greater competition from star hospitals, reduced leverage by not receiving steered demand from their VI insurers, and no incentives to foreclose rivals. Accordingly, prices at those hospitals decrease.

In total, inpatient spending by private plans falls by 6.7 percent per enrollee, and total inpatient spending falls by 6 percent. Total household spending on healthcare and insurance decreases by 2 percent, despite plan actuarial value falling by 11 percentage points. Plans' preferential tiers include only about half the number of providers than at baseline, but there is more variety in preferential structure and slightly more generous

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

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Baseline		A - Full effect		B - Decomposition			
	Raw	Weighted	Raw change	Weighted change	Short run		Medium run	
					Raw change	Weighted change	Raw change	Weighted change
Hospital prices (by hospital / insurer)								
Within VI	3.61	4.04	0.84	-0.13	0.08	-0.06	0.76	-0.08
VI / VI	3.78	3.91	-0.11	-0.29	-0.28	-0.28	0.17	-0.01
VI / Non-VI	3.69	3.84	0.28	-0.34	-0.10	-0.23	0.38	-0.11
Non-VI / VI	5.86	6.31	0.24	0.45	0.54	0.02	-0.30	0.43
Non-VI / Non-VI	5.44	6.10	1.05	0.27	1.74	1.23	-0.69	-0.97
Public / VI	1.57	1.49	0.00	-0.05	0.00	-0.04	0.00	-0.01
Public / Non-VI	1.58	1.49	0.00	-0.03	0.00	-0.01	0.00	-0.02
All	4.32	3.89	0.45	-0.44	0.46	-0.10	-0.01	-0.33
Premiums (by insurer)								
VI	1.22	1.15	0.11	0.04	0.05	0.03	0.06	0.00
Non-VI	1.38	1.29	-0.05	-0.07	0.02	0.02	-0.07	-0.09
All	1.31	1.23	0.02	-0.02	0.03	0.03	-0.02	-0.05
Plan design (by insurer)								
VI - Base coverage	52.54	48.46	-6.26	-6.56	0.00	0.01	-6.26	-6.57
VI - Preferential coverage	78.15	76.23	2.10	1.81	0.00	-0.72	2.10	2.53
VI - Self-preferencing	67.49	65.43	-22.68	-26.79	0.00	1.98	-22.68	-28.76
VI - Other-VI-preferencing	13.56	12.77	-7.39	-6.60	0.00	1.04	-7.39	-7.64
VI - Non-VI-preferencing	4.69	4.97	1.15	2.39	0.00	-0.47	1.15	2.86
VI - Star-hospital-preferencing	0.48	0.45	7.42	8.17	0.00	-0.19	7.42	8.36
Non-VI - Base coverage	58.88	58.67	-21.76	-19.64	0.00	1.49	-21.76	-21.14
Non-VI - Preferential coverage	85.75	84.94	0.11	0.93	0.00	0.44	0.11	0.49
Non-VI - Other-VI-preferencing	22.12	21.40	-10.21	-9.34	0.00	-0.44	-10.21	-8.90
Non-VI - Non-VI-preferencing	16.91	17.98	-7.46	-9.26	0.00	-0.25	-7.46	-9.01
Non-VI - Star-hospital-preferencing	13.97	18.22	2.02	-1.91	0.00	-1.12	2.02	-0.79
Healthcare spending								
Actuarial value		0.66		-0.11		-0.02		-0.08
Inpatient spending private plan		1.04		-0.07		-0.12		0.06
Inpatient spending		0.50		-0.03		-0.03		0.01
Total household spending		1.10		-0.02		-0.01		-0.01

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

coverage at preferential providers. Thus, households of heterogeneous risk find better-matching networks, and elastic consumers have greater access to cheaper medium- to low-quality hospitals. Higher prices at high-quality 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.

6.2 Choices, Efficiency, and Welfare

Lower prices and redesigned plans reallocate enrollment and admissions, as shown in Table 6. VI hospitals lose almost half their market share as former partners no longer steer demand their way. Accordingly, admissions shares equilibrate: before the ban, VI insurers accounted for 72 percent of admissions at partner hospitals, whereas post-ban, they account for 21 percent. Despite improvements in VI plan access to star hospitals, the elimination of VI self-preferencing decreases their plans' network surplus, leading to a 26 percent loss in market share, mostly recaptured by non-VI rivals.

Banning VI increases efficiency in insurance provision. At baseline, coverage is inefficiently low on average: a marginal increase in coverage to the median patient increases welfare by tens of thousands of dollars. These distortions are consistent with VI firms' incentive to skew coverage in favor of their hospitals and the presence of adverse selection. Banning VI eliminates incentives for self-preferencing and increases competition, leading to a more efficient coverage provision. However, adverse selection still affects the market and coverage remains underprovided. Residual adverse selection is worsened by intensified insurer competition, yet improvements in access, prices, and spending efficiency curtail its effects.

In total, VI firms lose \$161.6 million in profits, 87 percent of which is captured by rivals. Lower hospital prices and more intense competition benefit non-VI insurers, whose profits increase by 19 percent. Consumer surplus increases by \$62.7 million. However, 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.³⁶

6.3 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. In the short run, network and coverage choices are held fixed, whereas premiums and price negotiations adjust. Table 5-B shows that the short-run price effects of VI are limited, with a lower impact of double

³⁶We provide average impacts on consumer surplus by policyholder characteristics in Appendix Table A.8. While most groups benefit on average from banning VI, consumers located further from the city center, those with dependents, and those who are older benefit the most. Moreover, private plan enrollees benefit more than those in the public plan, given the low substitution between private and public plans.

Table 6: Effects of vertical integration on choices and welfare

	(1)	(2)	(3)	(4)
		A - Full effect	B - Decomposition	
	Baseline	Change	Short run Change	Medium run Change
Efficiency				
Moral hazard spending	44.10%	-16.43	-14.45	-1.98
Median marginal value of coverage	15.08	-8.99	-6.88	-2.11
Spread marginal value of coverage	56.97	-37.37	-31.28	-6.09
Adverse selection	0.37	0.17	0.05	0.12
Market shares				
VI hospital	38.66%	-16.69	-15.74	-0.94
Non-VI hospital	30.95%	0.27	2.83	-2.55
VI insurer	10.84%	-2.84	-2.39	-0.45
Non-VI insurer	13.18%	2.17	0.01	2.16
Admission shares (by hospital / insurer)				
Within VI	72.00%	-50.67	-43.96	-6.72
VI / VI	3.76%	7.33	9.10	-1.77
VI / Non-VI	24.24%	43.34	34.85	8.49
Non-VI / VI	22.95%	10.64	12.55	-1.91
Non-VI / Non-VI	77.05%	-10.64	-12.55	1.91
Profits (in millions)				
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
Consumer surplus				
VI enrollees (per member)	-	-0.08	-0.66	0.59
Non-VI enrollees (per member)	-	0.52	-0.56	1.08
Total consumer surplus (in millions)	-	62.65	-86.26	148.91
Share better off	-	0.72	0.01	0.71
Total welfare (millions)	-	41.69	-129.13	170.82

Notes: Moral hazard spending 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 given hospital, accounting for equilibrium effects. Adverse selection is the residual correlation between WTP for insurance and cost, conditional on market segment. Profits and total consumer surplus are measured in millions of dollars per year. Consumer surplus for VI enrollees is the average surplus conditional on enrolling in a VI plan, unweighted by demand. Non-VI consumer surplus is defined analogously. All values are in thousands of dollars unless stated otherwise.

marginalization and a smaller decrease in average paid hospital prices. 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 insurer plans creating a service gap in the market: Non-VI insurers could benefit from competing for these enrollees by redesigning plans to attract them.³⁷ Nevertheless, insurers cannot fix either distortion in the short run. 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.

In the short run, VI increases welfare. Intuitively, with fixed networks, the effects of VI depend on consumer relative premium and price elasticities. As consumers are more premium- than price-elastic, VI hospitals' ability to harm rival insurers with higher prices is limited, as rivals can adjust premiums in response to mitigate demand losses. Similarly, VI insurers are more effective at steering demand with lower premiums and can capture profits more effectively through hospital prices. Thus, in the counterfactual with fixed coverage, 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 vastly 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 medium-run plan design adjustments. While the key distortion in coverage is driven by VI insurers' inability to adjust in the

³⁷For 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.

short run, 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.

This decomposition delivers two lessons about VI. First, the gains from banning VI are not mechanically due to the elimination of distortive VI marketing. Those distortions are eliminated in the short run, yet banning VI in that context leads to a welfare loss. Second, that accounting for endogenous plan design responses to changes in market structure is relevant. VI firms can only retain profits if they can steer demand toward their hospitals when profitable, and steering is largely implemented by plan design.

6.4 The Roles of Cost Efficiencies, Quality, and Preferences

The welfare effects of VI uncovered above are local to the Chilean healthcare market. In this section, we use our model to extrapolate beyond our setting by varying components that might play a role in other markets.³⁸

6.4.1 Cost Efficiencies. VI may eliminate wasteful spending by inducing hospitals to internalize their costs. While we find no support for this in our setting, we implement our counterfactual analysis under varying cost efficiency degrees to evaluate its impact. 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 the welfare gains from banning VI decrease with cost efficiencies. Higher costs for formerly VI hospitals under a VI 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 seem substantial: Recent research finds that VI plans in Medicare Advantage incur higher costs than their rivals (Park *et al.*, 2023).

6.4.2 Quality Effects. 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. We adjust the fraction of the VI effect in hospital demand—treated as marketing in the main analysis—that is welfare-relevant. Appendix Figure A.5b

³⁸Medicare Advantage is the natural comparison for the Chilean market. Both are regulated private insurance markets with a dominant public option and substantial insurer-hospital VI.

shows the results. We find that the higher the quality effects of VI, the lower the consumer surplus and welfare gains from banning VI. A VI ban becomes welfare-neutral when 25 percent of the VI effect on demand is due to quality. This magnitude implies VI hospitals can produce \$473 worth of quality improvements for each enrollee of their integrated insurer per year and that the average VI hospital can close 23 percent of its quality gap with the highest-quality hospital at will. For comparison, evidence from Medicare Advantage finds that VI insurers' plans have approximately 10 percent higher quality scores than non-VI insurers' plans (Parekh *et al.*, 2018).³⁹

6.4.3 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 distinct VI firms, one with high-quality hospitals, m_b , and one with medium- to low-quality, m_a . We simulate counterfactual scenarios for the dissolution of each VI firm independently. Appendix Tables A.10 and A.9 show the results. Banning only the high-quality VI firm leads to a similar total welfare change as the full ban. However, consumers' 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

³⁹Estimates from Vatter (2024) suggest this would be worth at most \$403.6 per enrollee-year. This calculation assumes the additional quality of VI is achieved by improved outcome quality obtained at VI hospitals and that 10 percent higher scores translate to 10 percent higher outcome quality. Evidence, however, suggests most quality improvements by VI firms come from patient satisfaction, which would vastly reduce their surplus value (Johnson *et al.*, 2017).

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.⁴⁰

6.4.4 Premium- and Price-sensitivity. Consumer responsiveness to prices and premiums is pivotal for the effects of VI. Theoretically, consumers must be sufficiently elastic to out-of-pocket hospital prices for plan design to steer demand among providers. Additionally, consumers must be sufficiently premium-elastic for VI insurers to attract consumers with low premiums to plans with limited access to rival hospitals. To quantify these forces, we simulate the impact of banning VI under alternative price and premium elasticities. Appendix Figure A.6 shows the results.

For moderate shifts from the baseline, VI enhances welfare when consumers are more price- than premium-elastic and reduces welfare when the opposite is true. Intuitively, when consumers are significantly more responsive to prices than premiums, steering consumers becomes costly. Consumers become very responsive to differences in hospital access, limiting VI insurers' profits from skewing coverage. In this scenario, VI firms lean on their price advantage—a product of the elimination of double marginalization—resulting in welfare improvements. In contrast, when consumers are more responsive to premiums, VI insurers can skew coverage and attract consumers with lower premiums while capitalizing on higher hospital prices. In both cases, the ability of VI firms to capture demand by increasing rivals' costs is limited by insurance competition and the ability of rival insurers to adjust premiums and coverage. Fundamentally, insurers can shift a marginal price increase to where it harms their enrollees less, lowering coverage and premiums when premiums are more important or increasing them when prices are.

When elasticities decrease by more than 50 percent from the status quo estimates, VI almost invariably enhances welfare. Intuitively, consumers must be price-sensitive for steering to be effective. Otherwise, VI firms struggle to direct demand within their networks, limiting their gains from skewing coverage. Consequently, they primarily benefit from eliminating double marginalization, which generally increases welfare.⁴¹

⁴⁰For 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).

⁴¹For the U.S., the evidence indicates that enrollees exhibit limited elasticity to out-of-pocket prices, placing them in the region of welfare-enhancing VI. In our model, however, this positive effect of VI occurs because of a breakdown of steering mechanisms. In contrast, consumers in the U.S. are significantly responsive to network structure, all but avoiding out-of-network care. This mechanism could likely substitute for steering through coinsurance rates, acting as if consumers were price elastic in our model.

6.5 Discussion

Our results describe how VI shapes market outcomes. In the short run, 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 downstream insurer competition and rivals' ability to adjust premiums. However, equilibrium plan design responses overturn the effects of VI in the medium run. VI insurers are incentivized to overprovide coverage at their hospitals, pushing rival non-VI insurers to increase their 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.

7 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 insurer's 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 and VI by incorporating plan network design and moral hazard aspects into a bilateral oligopoly model of insurers and hospitals.

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.

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.
- AMEMIYA, T. (1983). Chapter 6 Non-linear Regression Models. In *Handbook of Econometrics, Handbook of Econometrics*, vol. 1, Elsevier, pp. 333–389.
- 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.
- , LEVINSOHN, J. and PAKES, A. (1995). Automobile Prices in Market Equilibrium. *Econometrica*, **63** (4), 841–890.
- 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-Económica 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.

- 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., GAYNOR, M. and REENEN, J. V. (2018). The Price Ain't Right? Hospital Prices and Health Spending on the Privately Insured. *Quarterly Journal of Economics*.
- 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.
- DIAS, M. (2022). Selection in a Health Insurance Market with a Public Option: Evidence from Chile, 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.
- FIGUEROA, C. (2023). Adverse Selection and Equity in Insurance Markets with Guaranteed Renewable Contracts: Evidence from Chile. *Working paper*.
- 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., HANDEL, B., HENDEL, I. and WHINSTON, M. D. (2024). Optimal Long-term Health Insurance Contracts: Characterization, Computation, and Welfare Effects. *Review of Economic Studies*, **91** (2), 1085–1121.
- 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.
- 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, cited by: 9.
- LEE, R. S. (2013). Vertical Integration and Exclusivity in Platform and Two-sided Markets. *American Economic Review*, **103** (7), 2960–3000.
- 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.
- 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.
- 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.
- SERNA, N. (2023). Non-Price Competition, Risk Selection, and Heterogeneous Costs in Provider Networks, 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. *Review of Economic Studies*.
- 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. (2024). Quality Disclosure and Regulation: Scoring Design in Medicare Advantage. *SSRN Electronic Journal*.
- 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.
- YAKOVLEV, S. (2017). Convex Extensions in Combinatorial Optimization and Their Applications. In S. Butenko, P. M. Pardalos and V. Shylo (eds.), *Optimization Methods and Applications : In Honor of Ivan V. Sergienko's 80th Birthday*, Cham: Springer International Publishing, pp. 567–584.