Centralized and Decentralized Equilibria in Healthcare Provider Network Breadth[*]

Natalia Serna

June 11, 2025

Abstract

This paper shows that quality competition enhances welfare in markets with adverse selection when there are no prices. The intuition is that with price competition, a firm can enter offering a cheap, low-quality product, stealing consumers who are profitable. Without prices, offering low quality no longer attracts profitable consumer types. This is a common feature of health insurance markets where premiums tend to be heavily regulated, eliminating competition in that dimension. I test the theory using a structural model of the Colombian health insurance market where quality is defined as the insurers' provider network breadth. Counterfactual simulations show that the social planner would choose network breadth for hospital admissions equal to 60%, twice as broad as in the observed equilibrium but not complete. Collusion between insurers generates an equilibrium that is farther away from the social planner's. Certain network adequacy rules may help approximate the first-best.

Keywords: Health insurance, Adverse selection, Imperfect competition, Provider networks.

JEL codes: I11, I13, I18, L13.

1 Introduction

Since the seminal work of Rothschild and Stiglitz (1976), economists have been concerned with modelling endogeneous product characteristics in markets with adverse selection. Challenges that often arise in this type of work are the inexistence of equilibria and the difficulty of characterizing equilibria (when they exist) in an oligopoly. Several decades later, Azevedo and Gottlieb (2017) developed a framework in which an equilibrium always exists in perfectly competitive markets with adverse selection. Around the same time, Veiga and Weyl (2016) and Mahoney and Weyl (2017) characterized the consequences of competition in these markets, concluding that strong competition is harmful for welfare.

The conclusions from these papers rest on one crucial aspect of the market: that firms compete both on prices and quality.[1] Hence, the intuition for why competition lowers welfare is that a firm can always enter with a cheap, low-quality product, stealing consumers who are more sensitive to price than quality and who are potentially profitable. However, in practice, many markets such as health insurance, have stringent price regulations that essentially eliminate price competition. Some examples include Medicaid managed care and Medicare Advantage in the US, where premiums are zero.[2] In this paper, I show that removing prices from these models generates opposite predictions about how competition affects quality, because offering low-quality no longer attracts the most profitable consumers.

Specifically, I develop a theoretical framework in which insurers compete on quality, receive fixed per-capita transfers from the regulator, and pass-through their average costs to enrollees in the form of cost-sharing. Then, I test the predictions of this model in the health insurance market of Colombia. My paper is the first to empirically characterize equilibrium quality under the social planner and to evaluate the consequences of counterfactual imperfect competition and quality regulations on welfare in health insurance markets without prices. Contrasting prior literature, I find that imperfect competition

---

[1]I use the term "quality" to refer broadly to non-price product characteristics that consumers value such as coverage generosity or provider network breadth in health insurance markets.

[2]Specifically, in the Medicaid program states cannot charge premiums to individuals with incomes below 150% of the federal poverty line (Guth et al., 2021) and in Medicare Advantage 67% of plans do not charge any premiums other than the Part B premiums required from all individuals that are eligible for Medicare (Freed et al., 2025).

generally harms welfare because firms do not internalize any of the value that consumers accrue from high-quality products, leading to an equilibrium in which firms choose low quality. These findings have important implications for the design of policies that encourage competition in health insurance markets.

In the theoretical framework, I capture adverse selection in two ways: first, through the dependence of insurers' costs to the types of consumers that they enroll and second, through changes in total insurance demand and the composition of consumer types in demand when quality varies. I assume quality is a scalar, and start by characterizing the monopolist's optimal quality. Increasing quality raises the monopolist's average cost among existing consumers but also attracts marginal consumers with relatively strong preferences for quality versus cost-sharing. If these marginal consumers are unprofitable under the imperfect transfers from the regulator, then the monopolist chooses lower quality than in markets where adverse selection is weaker.

Then, I solve the social planner's problem, which involves choosing quality to maximize the sum of consumer welfare, insurer profits, and government spending. Marginal consumers are generally more "profitable" for the social planner than for the monopolist because the planner completely internalizes their valuation for quality. Therefore, equilibrium quality under the social planner is higher than under an insurance monopoly. The difference in optimal quality between the two scenarios suggests that competition could bring market outcomes closer to the planner's solution. Examining the equilibrium in an insurance duopoly indeed reveals that competition enhances welfare. This finding is similar to Shaked and Sutton (1982)'s solution in markets with perfect information where quality competition, holding prices fixed, generates an equilibrium in which at most two firms have positive market share and one of them offers a high-quality product.

To test the predictions of the model, I use data from the Colombian health care system. In this system, private insurers offer a single national health insurance plan. Insurers do not charge premiums or set their own cost-sharing rules, but they can design their provider networks for every health service offered in the plan. Hence, my measure of insurer quality is the fraction of providers that offer a particular service and are covered by the insurer, which I refer to as service network breadth. In this case, insurer quality

3

is a vector of network breadths for multiple services, expanding the scope of the scalar quality considered in the theoretical framework.[3]

My data encompasses individual-level enrollment and health claims from all consumers in Colombia's contributory health system between 2010 and 2011, which covers those who pay taxes and their dependents. I also have information on insurers' provider listings for every health service. Using these data, I provide descriptive evidence that supports the predictions of the theoretical framework, namely, that insurers' networks of covered providers are generally narrower (i.e., of lower quality) in more concentrated markets, where insurers likely exert market power. To simulate counterfactual scenarios that mirror the social planner and the insurance monopoly from the theoretical framework, I transition to the structural model.

I borrow the model and estimates from Serna (2024). On the demand side, new consumers select their insurer based on expected out-of-pocket expenses and the breadth of the provider networks. On the supply side, insurers compete by simultaneously choosing their provider network breadths across multiple services to maximize the net present value of their profits. Demand estimates show substantial adverse selection on service network breadth. Sicker, relatively unprofitable individuals have higher willingness-to-pay for network breadth in services they are more likely to claim compared to healthy, relatively profitable individuals. Cost estimates show that insurers enjoy economies of scale in the number of covered providers and in some cases also enjoy economies of scope from offering high network breadth across multiple services.

The observed equilibrium in service network breadth—where 12 insurers compete—is asymmetric, with some insurers choosing broad networks and others choosing narrow networks for the same service within the same market. Given the zero premiums and stringent regulations on cost-sharing, I test whether the level of competition among insurers for each health service factors into equilibrium network breadth choices. If all consumers

---

[3]One limitation of the theoretical model is that quality competition when quality is a scalar inevitably results in a pooling equilibrium in which healthier consumers are made worse off by the presence of sick consumers (Rothschild and Stiglitz, 1976). The structural model relaxes this assumption and allows me to test how insurer competition impacts equilibrium quality under the scenarios considered in the theoretical framework. In this case, separating equilibria could exist where healthy consumers value quality in one service but sick consumers value quality in another service.

tend to prefer broader networks independent of their health status, then insurers may have incentives to offer broader networks to attract a bigger pool of enrollees in markets where they face stronger competition.

I begin by simulating a counterfactual scenario in which the social planner chooses the vector of service network breadth for each insurer to maximize consumer surplus subject to insurers' participation constraints in which they make zero profits. In line with the theoretical framework, I find that the social planner would choose much broader networks than in the observed scenario. For example, the planner would choose network breadth for general medicine and hospital admissions that are twice as broad, moving from 0.3 to 0.6 in both services. In this case, long-run consumer surplus would increase by about 19%, with equal gains for individuals with and without diagnoses. Notably, provider networks under the social planner are not complete because of the fixed administrative costs that the health system would incur. Thus, the social planner faces a trade-off between balancing administrative costs and offering comprehensive provider network coverage.

To explore whether a decentralized equilibrium can move market outcomes closer to the social planner's, I simulate a counterfactual scenario in which insurers maximize joint profits. My findings indicate that, irrespective of which insurers collude, imperfect competition generates an equilibrium in which networks are narrower than in the observed scenario and therefore farther away from the social planner's solution. For instance, simulations of the model assuming the bottom five insurers collude predict that average network breadth would fall between 3% and 16% depending on my assumptions regarding the merged firm's cost structure. This suggests that network coverage improves with stronger insurer competition in markets without premiums.

Finally, given the challenges of designing regulations that promote insurer competition, in the last part of the paper I investigate the impact of policies that more directly address risk selection incentives. I examine a network adequacy rule that forces insurers to offer the same network breadth for general medicine and hospital admissions in an attempt to implement the planner's solution. My findings reveal that average network breadth would increase 17% for hospital admissions and decrease slightly for the rest of services, resulting in marginal gains in consumer surplus. Therefore, while the network adequacy

rule enhances hospital coverage, it falls short of fully implementing the social planner's solution.

Literature. This paper contributes to the literature by studying how competition impacts quality in markets with adverse selection when there is no price competition. These features are common in health insurance markets, where premiums often tend to be heavily regulated and where consumers choose their insurer based on unobserved health status. On the theory side, I build on Stiglitz and Yun (2013), Veiga and Weyl (2016) (VW hereafter), and Mahoney and Weyl (2017) to derive optimal quality under several assumptions regarding insurer competition. While I do not explore the effects of perfect competition between insurers, Azevedo and Gottlieb (2017) and Fang and Wu (2018) characterize this type of equilibrium in markets where firms facing adverse selection compete on prices. Moreover, my paper relates to MacLeod (2021) by examining the choices of a social planner that maximizes consumer welfare.

On the empirical side, I add to the body of work that has examined the trade-offs to competition in markets with adverse selection when firms can charge prices (e.g., Kong et al., 2024; Saltzman et al., 2021; Cuesta and Sepúlveda, 2021; Einav et al., 2012; Capps et al., 2003; Cutler and Reber, 1998) as well as the implications of these prices for welfare (e.g., Bundorf et al., 2012; Cabral et al., 2018).[4] In health insurance markets, in particular, Ho and Lee (2017); Dafny et al. (2015) examine the impact of insurer competition on premiums. My work complements this prior literature by investigating how competition influences insurer quality when there are no premiums. Importantly, my paper is the first to empirically derive the social planner's solution and to evaluate network adequacy rules that target insurers' leverage for risk selection. The analogy in markets with premiums are counterfactuals where insurers can price discriminate (Handel et al., 2015) or where risk adjustment is made more granular (Brown et al., 2014).

I am able to provide characterizations of equilibria under the social planner and different scenarios of insurer competition thanks to the tractability of my model. Here too I contribute to previous and ongoing work on upstream-downstream bilateral negotiations (e.g., Cuesta et al., 2024; Ho and Lee, 2017; Gowrisankaran et al., 2015; Grennan, 2013) and

---

[4]These findings have been summarized in Einav et al. (2021).

on determinants of healthcare provider network exclusions (e.g., Ho and Lee, 2019; Ghili, 2022).

The rest of this paper is organized as follows. Section 2 presents a theoretical model of insurer competition on quality. Section 3 summarizes the empirical setting and data. Section 4 presents the structural model and estimates. Section 5 shows results of a centralized equilibrium where the social planner chooses network breadth. Section 6 derives equilibrium network breadth when insurers maximize joint profits. Section 7 provides results of a network adequacy rule. And Section 8 concludes.

2   Theoretical Model

To examine how insurer competition affects provider networks, I develop a simple model of competition in network breadth. Suppose there is a unit mass of consumers. Each consumer is characterized by a unit-dimensional type $\theta \in \mathbb{R}$. The consumer's type follows a distribution $F(\theta)$ with continuously differentiable density function $f(\theta) > 0$. The consumer's type denotes their sickness level, so higher θ means the individual is in worse health. Consumers can choose from a set of insurers $\{1, ..., j, ..., J\}$ that offer network breadth $H_j \in [0, 1]$.[5] For simplicity, I assume network breadth is a scalar—although I will relax this assumption in my empirical application in section 4.

The expected medical cost of a type-θ consumer is $c(H_j, \theta)$, which is twice-continuously differentiable in both terms. Additionally, assume $c_H(H_j, \theta) > 0, c_\theta(H_j, \theta) > 0, c_{HH}(H_j, \theta) > 0, c_{\theta\theta}(H_j, \theta) > 0$, and $c_{H\theta}(H_j, \theta) > 0$.[6] The consumer pays a fraction r of her expected medical cost. This coinsurance rate is fixed exogenously and does not vary across insurers. The structure of the expected medical cost captures adverse selection because different consumer types have different costs conditional on network breadth. The cost structure also captures moral hazard because the medical cost depends on network breadth

---

[5]Network breadth defined as the fraction of providers that the insurer covers can also be interpreted as a measure of insurer quality, similar to VW.

[6]The expected medical cost increases more rapidly with health status under a broad network than a narrow network presumably because the broad-network insurer covers higher quality providers that charge higher prices for the same procedure relative to a narrow-network insurer.

conditional on the consumer type. Consumer θ's utility function for contract H_j is:

$$U(H_j, \theta) = u(H_j, \theta) - rc(H_j, \theta)$$

where $u(H_j, \theta)$ is also twice-continuously differentiable in both terms. Suppose individuals can choose to stay uninsured, the utility of which equals zero. Consumers buy insurance if $U(H, \theta) \geq 0$. This inequality defines a set of buyers $B(H_j)$ with cutoff type $\theta^*(H_j)$ such that $\theta \in B(H_j) = [\theta^*(H_j), 1]$. This cutoff type is implicitly defined as:

$$u(H_j, \theta^*(H_j)) = rc(H_j, \theta^*(H_j))$$

Insurers offer one level of network breadth to all enrollees and make per-enrollee profits equal to $\pi(H_j, \theta) = R(\theta) - (1 - r)c(H_j, \theta)$. $R(\theta)$ is an imperfect risk-adjusted transfer from the government such that $R_\theta(\theta) > 0$ and $R_\theta(\theta) < c_\theta(H_j, \theta)$. Moreover, assume it is always profitable to serve the healthiest consumer under a broad network $R(0) > (1-r)c(1, 0)$, but unprofitable to serve the sickest consumer under a narrow network $R(1) < (1 - r)c(0, 1)$.

Monopoly. The monopolist's problem is to choose network breadth to maximize profits given by:

$$\Pi(H) = \int_B \pi(H, \theta)f(\theta)\, d\theta$$

The first order condition (FOC) of this problem is:

$$\frac{d\Pi}{dH} = -\underbrace{(1 - r)\mathbb{E}\left[c_H(H, \theta)|B\right]}_{\text{Average marginal cost}} - \underbrace{\left[R(\theta^*) - (1 - r)c(H, \theta^*)\right]}_{\text{Profitability of marginal consumer}} \underbrace{\left(\frac{rc_H(H, \theta^*) - u_H(H, \theta^*)}{u_\theta(H, \theta^*) - rc_\theta(H, \theta^*)}\right)}_{\text{Selection effect}}$$

The first term on the right-hand side of the equation above captures the change in the insurer's costs among the existing set of enrollees when it increases network breadth. The second term is similar to VW's covariance term: it captures the impact of attracting new marginal consumers, weighted by their profitability. If the marginal consumer is profitable but their valuation for additional network breadth is lower than the marginal increase in out-of-pocket costs, then the insurer will lower its network breadth to attract

this consumer. Denote by H^*_{mon} the monopolist's optimal network breadth.

Social planner. A social planner who maximizes the sum of consumer surplus and insurer profits, minus government spending has the following objective function:

$$W(H) = \int_B [w(H, \theta) - c(H, \theta)] f(\theta)d\theta$$

where $w(H, \theta)$ is consumer surplus. In this objective function, risk-adjusted payments cancel out because they are linear transfers from the government to the insurer. The FOC of the planner's problem is:

$$\frac{dW}{dH} = \underbrace{-\mathbb{E}[c_H(H, \theta)|B] + \mathbb{E}[w_H(H, \theta)|B]}_{\text{Average marginal surplus}} - \underbrace{[w(H, \theta^*) - c(H, \theta^*)]}_{\text{Surplus of marginal consumer}} \underbrace{\left(\frac{rc_H(H, \theta^*) - u_H(H, \theta^*)}{u_\theta(H, \theta^*) - rc_\theta(H, \theta^*)}\right)}_{\text{Selection effect}}$$

To interpret this expression, suppose there are no externalities across consumers, so that $w = u$. The social planner internalizes the value of providing network breadth to all consumers captured by the term $\mathbb{E}[u_H(H, \theta)|B]$, which contrasts with the monopolist's solution where this value is zero. This is an extreme case of Spence (1975)'s distortion because the monopolist cannot set premiums. In VW's setting, the monopolist internalizes some of the value of providing network breadth—the value that marginal consumers derive.

In my model, the contribution of marginal consumers to welfare differs from their contribution to private profits because the insurer receives risk-adjusted transfers from the government that do not necessarily reflect the consumer's willingness-to-pay for network breadth or their health status, so in general we can expect $w(H, \theta) > R(\theta)$.[7] This means that in a world with complete insurance generosity, $r = 0$, the planner will indistinctly choose broader networks than the monopolist. With incomplete generosity this is also true if the marginal consumer's valuation for network breadth is higher than the marginal

---

[7]In countries that use risk adjustment, the formula typically does not account for the insurer's quality other than through the correlation between the consumer's type and their valuation for quality.

cost of providing network breadth. In that case, we obtain the following condition:

$$H^*_{planner} > H^*_{mon} \text{ and } \theta^*(H^*_{planner}) < \theta^*(H^*_{mon})$$

where $H^*_{planner}$ is the optimal network breadth for social planner. The condition indicates that with a monopolist insurer network breadth is underprovided, fewer consumers buy insurance than is socially efficient, and these consumers are relatively sicker compared to the social planner's solution. The comparison between the two equilibria in my setting contrasts with VW in which the selection distortion has the same impact on equilibrium quality for the monopolist as it does for the social planner.

Are the differences between the monopolist and the social planner in my framework reasonable in the real world? There are several health systems where premiums and cost-sharing are heavily regulated, essentially eliminating the possibility that insurers compete on these dimensions of plan design. Examples include Medicaid managed care and Medicare Advantage in the U.S., and the Netherland's, Switzerland's, and Colombia's health systems. Risk adjustment is also perhaps the most common tool that regulators use to mitigate risk selection incentives. Therefore, insurer marginal revenues depend on the consumer's type through these risk-adjusted transfers, something that markets where firms compete on premiums do not necessarily consider because premium discrimination across health status is generally not allowed (and defeats the purpose of insurance).

Duopoly. Now consider the case of an insurance duopoly. Two insurers $j \in \{a, b\}$ compete by choosing network breadth $H_j \in [0, 1]$. The consumer type $\theta'(H_a, H_b)$ that is indifferent between enrolling with a and b satisfies:

$$u(H_a, \theta'(H_a, H_b)) - rc(H_a, \theta'(H_a, H_b)) = u(H_b, \theta'(H_a, H_b)) - rc(H_b, \theta'(H_a, H_b))$$

I will denote the consumer that is indifferent between the two firms as the "marginal" consumer. Let $\theta^*(H_a, H_b)$ denote the lowest type willing to enroll with insurer a, defined previously by:

$$u(H_a, \theta^*(H_a, H_b)) = rc(H_a, \theta^*(H_a, H_b))$$

I will denote the lowest type that buys insurance as the "minimum" consumer, which maps to the monopolist's "marginal" consumer. Define the set of buyers for insurer j as $B_j = \{\theta : \theta \in [\underline{\theta}, \overline{\theta}]\}$, which given the marginal type implies $B_a = \{\theta : \theta \in [\theta^*(H_a, H_b), \theta'(H_a, H_b)]\}$ and $B_b = \{\theta : \theta \in [\theta'(H_a, H_b), 1]\}$. Note that I have implicitly assumed that consumers have stronger preferences for insurer b and that consumers of higher type will always prefer to buy insurance. While these assumptions are not without loss of generality, they are reasonable. For instance, making consumers have stronger preferences for one insurer will have similar equilibrium implications as assuming insurers have different cost structures. Insurer j's profit function is:

$$\Pi_j(H_j, H_{-j}) = \int_{B_j} \left[R(\theta) - (1-r)c(H_j, \theta) \right] f(\theta) d\theta$$

And the corresponding FOC is:

$$\frac{d\pi_j}{dH_j} = -(1-r)\mathbb{E}[c_H(H_j, \theta)|B_j] + \left[R(\overline{\theta}) - (1-r)c(H_j, \overline{\theta})\right]\frac{d\overline{\theta}}{dH_j} - \left[R(\underline{\theta}) - (1-r)c(H_j, \underline{\theta})\right]\frac{d\underline{\theta}}{dH_j}$$

Let $\Delta f_\theta = f_\theta(H_a, \theta') - f_\theta(H_b, \theta')$ for any function f. Writing the FOC for each insurer separately yields:

$$\frac{d\pi_a}{dH_a} = -\underbrace{(1-r)\mathbb{E}[c_H(H_a, \theta)|B_a]}_{\text{Average marginal cost}} - \underbrace{\left[R(\theta^*) - (1-r)c(H_a, \theta^*)\right]}_{\text{Profitability of minimum consumer}}\underbrace{\left(\frac{rc_H(H_a, \theta^*) - u_H(H_a, \theta^*)}{u_\theta(H_a, \theta^*) - rc_\theta(H_a, \theta^*)}\right)}_{\text{Selection effect}}$$

$$+ \underbrace{\left[R(\theta') - (1-r)c(H_a, \theta')\right]}_{\text{Profitability of marginal consumer}}\underbrace{\left(\frac{rc_H(H_a, \theta') - u_H(H_a, \theta')}{\Delta u_\theta - r\Delta c_\theta}\right)}_{\text{Competitive effect}}$$

$$\frac{d\pi_b}{dH_b} = -\underbrace{(1-r)\mathbb{E}[c_H(H_b, \theta)|B_b]}_{\text{Average marginal cost}} - \underbrace{\left[R(\theta') - (1-r)c(H_a, \theta')\right]}_{\text{Profitability of marginal consumer}}\underbrace{\left(\frac{rc_H(H_b, \theta') - u_H(H_b, \theta')}{r\Delta c_\theta - \Delta u_\theta}\right)}_{\text{Competitive effect}}$$

To analyze the impact of competition on equilibrium network breadth consider insurer a's

FOC. If insurer b increases its network breadth, it will steal the relatively sicker consumers from a—who are healthier than b's current pool of enrollees—, thus $\theta'(H_a, H_b)$ decreases, reducing a's market share. At the same time, increasing b's network breadth may decrease the minimum consumer willing to participate in the market $\theta^*(H_a, H_b)$, expanding a's market share and making its pool of enrollees relatively healthier. If the *minimum* consumer type is profitable and has a weak valuation for network breadth, then insurer a will lower its network breadth to attract healthier types (selection effect). If the *marginal* consumer is profitable and has a weak valuation for network breadth, then insurer a will increase its network breadth to attract the relatively healthier types from b (competitive effect). Hence, the impact of competition on insurer a's equilibrium network breadth will depend on the relative magnitudes of the selection and competitive effects.

For insurer b, an increase in a's network breadth will shrink its market size, shifting the enrollee composition towards sicker types. If the marginal consumer is profitable and has a strong taste for network coverage, then b's best response is to increase network breadth. Denote by $\overline{H}^*_{comp}$ the average equilibrium network breadth across all insurers in a competitive market, then we can expect the following relation:

$$H^*_{planner} > \overline{H}^*_{comp} > H^*_{mon} \tag{1}$$

While the predictions of the model regarding whether network breadth decisions are strategic complements or substitutes are ambiguous (depending on the selection and competitive effects), *on average* competition generates an equilibrium in which networks are broader than the monopolist's solution. This contrasts with findings in VW and Mahoney and Weyl (2017) where competition is harmful in markets with adverse selection. In VW, competition implies that an insurer can always enter the market and offer a cheap, low-quality product, attracting all the healthy, relatively profitable types. In my setting, this pricing channel does not exist, thus an insurer that enters with a low-quality product will very likely have zero market share, unless out-of-pocket costs are very sensitive to network breadth.

Parametric specification. To settle final intuition on the effects of interest, I provide a

parametric specification and graphical representation of my model under the assumption that insurers can choose a binary network breadth $H_j \in \{0, 1\}$, where $H_j = 0$ denotes a narrow network and $H_j = 1$ denotes a broad network. This representation uses the following inputs:

$$u(H_j, \theta) = 1 + 0.5H_j + \log(2\theta)$$
$$c(H_j, \theta) = 0.2 + 0.2H_j + (2 + 0.2H_j)\theta^2$$
$$R(\theta) = 0.4 + 0.4\theta^2$$
$$F(\theta) = U[0, 1]$$
$$r = 0.5$$

Note that the functional form for $u(H_j, \theta)$ requires θ to be defined over $[\epsilon, 1]$ for a small ϵ. The top panels of Figure 1 depict the marginal consumer under each network in the dashed black line. This marginal consumer is higher (i.e., relatively sicker) when networks are narrow because only sick consumer derive sufficient value from enrolling. The middle panels present the monopolist's revenue and cost curves and extend the marginal consumer from the top panels. A visual comparison of the monopolist's profits across the two scenarios reveals that the profit maximizing choice of network breadth is $H = 0$. Given the binary nature of the game, this solution shows that as long as the monopolist makes a profit, there exists an equilibrium where one low-quality contract is offered and only the "high types" enroll—similar to Rothschild and Stiglitz (1976)'s result of nonexistence of a pooling equilibrium.

The bottom panels of Figure 1 depict the social planner's problem, extending the marginal consumer and reproducing in blue the welfare function (absent externalities) and in red the total cost curve. Here too a visual inspection indicates that social surplus is higher under a broad network than a narrow network because, unlike the monopolist, the social planner internalizes the full value of providing network breadth.

Network breadth defined as an index over the unit interval summarizes potentially complex bilateral negotiations between insurers and providers over healthcare prices and

network inclusions. This representation of the bargaining game captures relevant features of a bargaining environment, such as allowing insurers' cost to increase with the inclusion of a provider if enrollment rises and to decrease with the exclusion of a provider if providers are substitutable. Under these conditions, the comparison of equilibrium network breadth under an insurance monopoly, the social planner, and the duopoly derived above likely holds in markets with these kinds of bilateral negotiations.
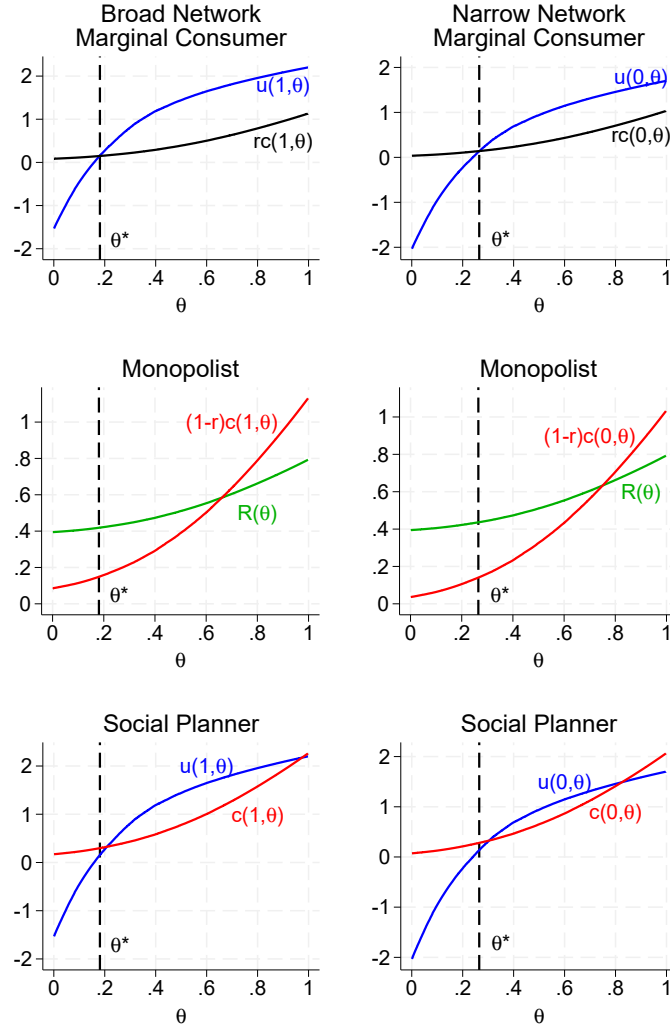
FIGURE 1: Monopolist's and Social Planner's Problem

Note: The first column of the figure shows the utility and cost curves under a broad network for consumers in the top panel, the monopolist insurer in the middle panel, and the social planner in the bottom panel. The second column shows the utility and cost curves under a narrow network in the same scenarios. θ^* denotes the marginal consumer who is willing to purchase insurance.

However, summarizing insurers' networks with a single index does not capture other

elements of the bargaining game, such as changes in insurers' and providers' disagreement payoffs. This raises the question of whether the comparison between the three scenarios holds after allowing for changes in these outside options. A simple example shows that if changes in the disagreement payoffs have the same directional effect on insurers' costs as the factual network then the predictions likely hold. Suppose insurer j excludes a provider from its network and as a result the prices that insurer $k \neq j$ pays to this provider increase. Then, in equilibrium insurer j should pay higher prices to the provider when including it in the network. Insurer j's costs are therefore higher when including the provider relative to excluding it, and this effect is even bigger when accounting for the externalities imposed by other insurers. A social planner would also internalize this externality in addition to the consumers' value for having broad networks. Thus, allowing insurers' costs to be endogenous to the disagreement payoffs likely results in the same ranking of equilibrium network breadth as shown in equation (1).

Another potential limitation of the results presented in this section is the assumption of unit-dimensional quality. With scalar quality a pooling equilibrium in which consumers of low θ are made better off relative to uninsurance does not exist. However, in reality, firms may differentiated along multiple dimensions of quality. For example, in health insurance, firms may differ in their provider networks, prior authorization requirements, claim denials, etc. If consumers of different type have different preferences over each dimension of quality, then we can have a separating equilibrium even in the insurance monopoly. In my empirical application I will relax this assumption and consider firms that differ in a vector of quality measures.

3   Data and Descriptive Evidence

3.1   Background

In the following sections, I apply the theoretical framework to the Colombian health insurance market. This market is divided into two main schemes: contributory and subsidized. The contributory scheme covers individuals who pay payroll taxes along with

their dependents, while the subsidized scheme is designed for low-income households. Colombia's insurance system operates under a managed care competition model, where private insurers offer a single national health insurance plan that has near-universal coverage. Premiums are set to zero, and both cost-sharing and benefits are heavily regulated. In the contributory scheme, individuals pay coinsurance rates and copays that are indexed to their monthly income, whereas healthcare is free for those in the subsidized scheme, aside from minimal copays for doctor visits.[8]

Private insurers are responsible for collecting payroll taxes and remitting contributions to the central government, which subsequently redistributes funds to the insurers using a risk adjustment formula. The formula compensates insurers in advance (ex-ante) based on the sex, age, and geographic location of their enrollees, but it does not account for specific diagnoses. Additionally, while the government provides ex-post compensations for certain chronic diseases, both forms of risk-adjusted payments are insufficient for effectively managing risk selection incentives.[9]

While insurers are prohibited from charging premiums or establishing their own cost-sharing rules, they compete for enrollees by determining which providers to cover and the number of providers available for each health service offered under the national insurance plan. For example, an insurer may choose to provide coverage for cardiac care at a particular provider while excluding renal care. Insurers also negotiate the prices of health services with in-network providers. Although the government has implemented some network adequacy rules for specific services such as primary care, oncology, and

---

[8]In 2011, for individuals with incomes below 2 times the monthly minimum wage (MMW) the coinsurance rate is 11.5% of the health service price, the copay is 1,900 COP, and the out-of-pocket maximum is 57.5% of the MMW. For individuals making between 2 and 5 times MMWs, the coinsurance rate is 17.3%, the copay is 7,600 COP, and the out-of-pocket maximum is 230% of the MMW. Finally, for individuals who earn more than five times the MMW, the coinsurance rate is 23%, the copay is 20,100 COP, and the out-of-pocket maximum is 460% of the MMW.

[9]Using health claims and enrollment data from year $t - 2$, the government calculates the ex-ante risk adjustment transfers for year t by computing the average annual health care cost per risk pool. Risk pools are defined by a combination of sex, age group (1-4, 5-14, 15-18, 19-44, 45-49, 50-54, 55-59, 60-64, 65-69, 70-74, 75+), and municipality of residence (insurers get 6% more for individuals who reside in the main capital cities of the country and 10% more for those who reside in peripheral areas like the amazon). The ex-post risk adjustment mechanism is known as the High-Cost Account. This is a zero-sum mechanism that compensates insurers with an above-average share of enrollees with chronic diseases with funds coming from insurers with a below-average share. The chronic diseases considered in this mechanism are: renal disease (since 2007), HIV-AIDS (since 2016), and certain cancers (since 2010).

urgent care, these rules do not encompass the entirety of services covered in the national plan. Overall, health service coverage in Colombia is extensive.

Consumers are free to choose any of the insurers that operate in their municipality of residence, and insurers typically participate in the majority of municipalities within a state.[10] Although there is no designated open enrollment period, consumers are allowed to switch insurers if they have been enrolled with their incumbent insurer for at least 12 (non-continuous) months. When making these enrollment decisions, consumers consider the set of providers that each insurer offers in their state of residence. Even though premiums are fixed and networks vary across insurers, switching is rare in this market: only about 6% of individuals switched their insurer between 2010 and 2011. Moreover, in 2011, only 4 out of the 23 insurers in the contributory scheme also operated in the subsidized scheme.

3.2   Data

I use individual-level enrollment and health claims data from all participants in Colombia's contributory scheme from 2010 to 2011, which includes approximately 24 million enrollees. The enrollment files provide detailed information on each enrollee's sex, age, municipality of residence, insurer, and length of enrollment within a year. These data enable me to calculate the ex-ante risk-adjusted transfers that each insurer received for its enrollees since the government's formula is public.

The health claims data includes the date the claim was filed, the insurer that processed the claim, the provider that delivered the service, the associated health service, diagnosis codes, and the negotiated price for each health service. Using this information, I can determine the consumers' health status by analyzing the diagnoses they received throughout the year, as well as compute insurers' total healthcare cost incurred in each individual. Anonymized individual identifiers are the same across datasets allowing me to merge enrollment with claims.

In addition to the enrollment and claims data, I have obtained provider listings from the

---

[10]Municipalities in Colombia are similar to counties in the U.S. and states are similar to Metropolitan Statistical Areas.

National Health Superintendency for insurers participating in the contributory scheme from 2010 to 2011. These listings detail the hospitals, clinics, and physician practices covered by each insurer, along with the specialties for which they are in network. I match the specialties in these provider listings with the relevant procedure codes in the health claims data based on the anatomical areas they pertain to. Examples of services include cardiac care, renal care, and hospital admissions. A complete list of these services is provided in Appendix Table 1.

The provider listings report the Colombian Tax Identification Number (TIN) for every provider. Each TIN may be associated with multiple facilities, each of which is assigned a unique provider identifier from the Colombian Ministry of Health and Social Protection. This provider identifier in turn matches the health claims data. I complement the information from the provider listings by incorporating providers from the claims data that do not appear in the listings but have submitted more than 10 claims for a specific insurer-service (results are robust to this threshold in the number of claims).

Using this final list of in-network provider-services, I calculate each insurer's service network breadth, defined as the fraction of providers in a market that offer a particular service and are covered by the insurer.[11] I define markets as Colombian states, recognizing that enrollees in more remote municipalities often travel to their state's capital city for care, thereby accessing their insurer's network in their state of residence. Insurers are required to cover at least one provider for each health service included in the national insurance plan. However, because consumers can access networks across different markets, some insurers may choose not to cover certain services in specific markets, potentially for profit-driven motives.

Throughout the analysis, I assume that service network breadth is the primary choice variable for insurers in this health system, as both premiums and cost-sharing are subject to strict regulation. This characterization of the networks assumes that providers are homogeneous conditional on the services they provide. For example, my model assumes that Stanford hospital and UCSF hospital, both of which can provide brain surgery, are

---

[11]In the construction of service network breadth, providers that do not deliver a particular service are excluded from the denominator of that service.
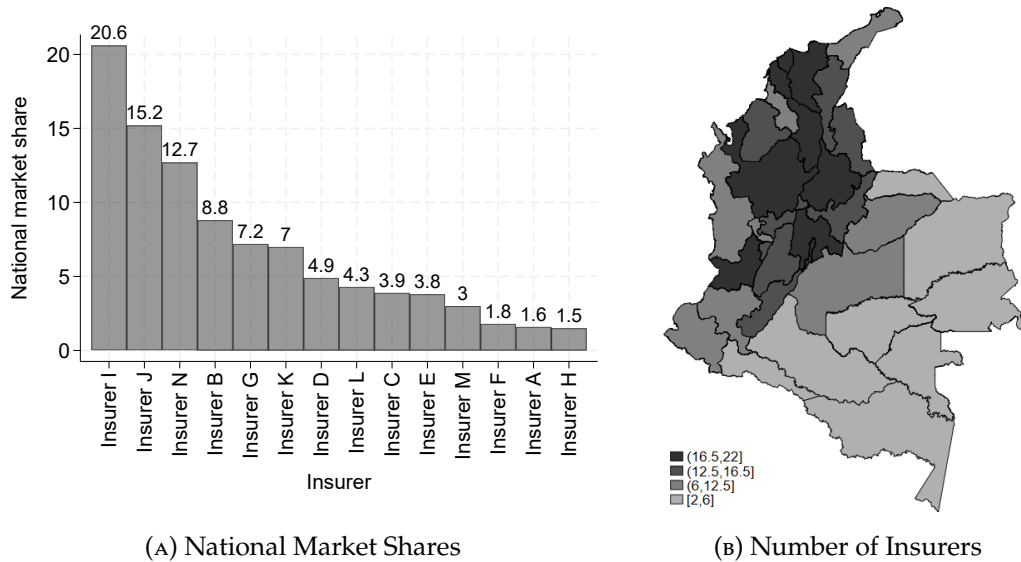
18

(A) National Market Shares   (B) Number of Insurers

FIGURE 2: Description of Markets

Note: Panel A presents the national market share on the full sample of enrollees of the top 14 insurance companies in the contributory system in 2011. Panel B presents the number of active insurers in every state in 2011. Darker colors represent higher numbers.

homogeneous in the provision of this service, but allows these two hospitals to differ from other providers that cannot render brain surgeries. Using service network breadth as a summary measure of insurer quality is appropriate in the Colombian setting, since there is relatively minimal variation in quality across providers for a specific service, as illustrated in Appendix Figure 1.[12]

During the sample period, 23 insurers participated in the contributory scheme, and 14 of these accounted for approximately 97% of enrollees. My analysis focuses on these 14 insurers. Figure 2 illustrates the structure of the contributory scheme in 2011. Panel A shows that the market is highly concentrated, with the three largest insurers covering 49% of enrollees. Notably, half of Colombian states had fewer than seven insurers, and around eight states exhibited an insurance duopoly, as shown in Panel B.

The considerable levels of market concentration in the contributory scheme raise important concerns about how competition affects insurers' network coverage decisions. The theoretical model presented earlier suggested that low competition can lead to provider networks that are narrower than socially desirable because insurers do not internalize any of the value to consumers of offering broad networks. In Colombia, insurers may wield

---

[12]If providers differed substantially in quality conditional on the service, as in the U.S., using service network breadth to characterize insurers' contracts would provide a lower bound on consumer surplus.

TABLE 1: Switch-in Rates

Insurer	All enrollees (1)	Continuously enrolled (2)
Insurer A	0.032	0.030
Insurer B	0.037	0.028
Insurer C	0.026	0.016
Insurer D	0.046	0.038
Insurer E	0.045	0.031
Insurer F	0.037	0.026
Insurer G	0.056	0.039
Insurer H	0.024	0.016
Insurer I	0.024	0.014
Insurer J	0.025	0.010
Insurer K	0.029	0.008
Insurer L	0.046	0.018
Insurer M	0.028	0.008
Insurer N	0.017	0.004

Note: Table shows the fraction of consumers that switch into each insurer in 2011 relative to 2010. Column (1) uses the full sample of enrollees without taking into account their enrollment spells. Column (2) conditions on enrollees with continuous enrollment spells in each year, that is, consumers who are enrolled 365 days.

market power by targeting the most profitable consumers and effectively "locking them in." Table 1 presents evidence of significant consumer inertia in line with these market power incentives. In the full sample of enrollees (column 1), only 1% to 6% of consumers changed their insurer between 2010 and 2011. For enrollees with continuous enrollment spells in both years (column 2), switching rates are even lower, ranging from 0.4% to 4%.

3.3 Service Network Breadth and Market Structure

Figure 3 describes my measure of service network breadth. Panel A indicates that insurers' network coverage decisions seem to be influenced by profit motives. This panel presents the average profit per enrollee—calculated as the risk-adjusted transfer minus total healthcare costs—across different bins of service network breadth conditional on individuals who file claims. Highly profitable individuals who do not file claims are excluded from this figure. The data shows that making any claim is associated with lower profits and that services for which insurers provide broad networks are associated with higher profits.

Panel B further illustrates that network breadth varies substantially both across ser-

vices and insurers, suggestive of selection incentives. This panel presents the distribution of residuals from a regression of service network breadth on service fixed effects (across insurer) and insurer fixed effects (across services) conditional on the largest market, the capital city of Bogotá. By focusing on a single market, this figure eliminates potential variation in network breadth that arises from differences in the set of insurers that participate in every market and, instead, depicts whether differences across insurers arise from the set of services that they decide to cover and how much they cover. The fact that the distribution of residuals has greater variance when controlling for the service suggests there is some non-random selection into insurers. This variation in residual service network breadth, given the stringent regulation of premiums and cost-sharing, stems from both differences in consumer preferences and insurers' cost structures, as discussed in Serna (2024).

Panel C uses information from all markets to determine whether market structure influences network breadth decisions. In this case, residual network breadth varies more substantially across markets than across services, suggesting the degree of competition also plays a role in insurers' choices. Interestingly, the distribution of residuals that is identified from variation in network breadth across insurers in Panel C is shifted to the left relative to Panel B, indicating that after accounting for the competitive landscape, insurers generally choose narrower networks.

(A) Network Breadth (B) Residual Networks, Bogotá (C) Residual Networks, All Markets
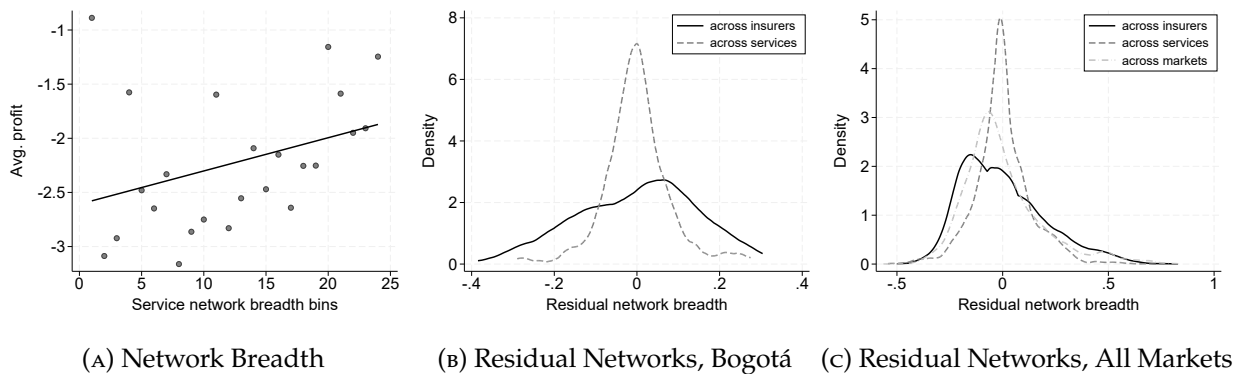
FIGURE 3: Distribution of Network Breadth

Note: Panel A shows the distribution of service network breadth in 2011 in black (left vertical axis) and the average profit conditional on consumers who make claims for each service in red (right vertical axis). Panel B shows the distribution of residuals from a regression of service network breadth in 2011 on insurer-by-service fixed effects in black, market-by-service fixed effects in dark gray, and market-by-insurer fixed effect in light gray.

21

To more directly examine the direction in which market structure influences network breadth, Figure 4 displays a scatterplot of average service network breadth in a market (averaged across insurers and services) against the Herfindahl-Hirschman Index (HHI). The HHI is calculated based on insurer market shares on the total number of claims. I focus on markets with enough variation in service network breadth which include the 13 metropolitan areas in the country. Panel A illustrates that markets with higher insurer concentration generally exhibit lower average service network breadth, consistent with the predictions of the theoretical model. Panels B and C break down the correlation into specific services such as general medicine and cardiac care, respectively. Insurers may face different risk selection incentives across these two services because they have different probabilities of being claimed. There is a strong negative correlation between HHI and average network breadth in general medicine, while the correlation is almost null conditional on cardiac care.

| (A) All Services | (B) General Medicine | (C) Cardiac Care |
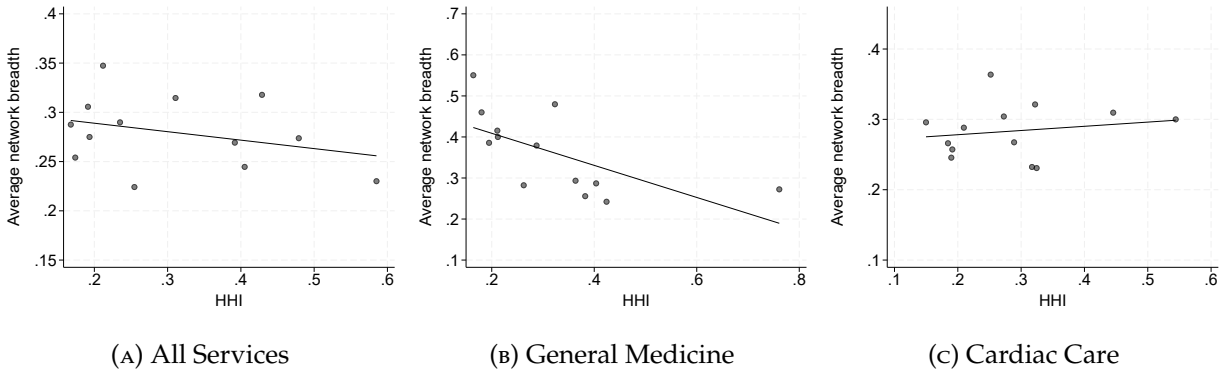
FIGURE 4: HHI and Average Service Network Breadth

Note: Scatter plot of average service network breadth in a market (across insurers and services) and the Herfindahl-Hirschman Index based on insurer market share in the total number of claims. Every dot is a market conditional on the 13 metropolitan areas in the country. The solid line represents a linear fit. Panel A uses network breadth across all services, Panel B uses network breadth in general medicine, and Panel C uses network breadth in cardiac care.

4 Empirical Model

Building on the descriptive evidence, this section introduces an equilibrium model of insurer competition on service network breadth. The model enables me to assess the impact of market power on provider networks, building on estimates from Serna (2024).

That paper offers a comprehensive overview of the model, identification strategy, and estimates. I summarize these modeling aspects in Appendix 2, and outline below the key empirical micro-foundations for the theoretical model presented in section 2.

Insurers compete for new enrollees who are make their first enrollment choice, because the Colombian insurance market is characterized by substantial consumer inertia. After consumers are "locked-in," insurers take into account the disease and age progression of their enrollees to choose the vector of service network breadth. Insurers make simultaneous choices of service network breadth to maximize the present discounted value of their profits. Take one market m, insurer j's profit function in this market is:

$$\Pi_j = \sum_\theta \pi_{j\theta}(\mathbf{H_j}, \mathbf{H_{-j}})N_\theta + \sum_{s=t+1}^T \zeta^s \sum_\theta (1 - \rho_{\theta'})\mathcal{P}(\theta'|\theta)\pi_{j\theta'}(\mathbf{H_j}, \mathbf{H_{-j}})N_{\theta'} - C_j(\mathbf{H_j}, \xi_{\mathbf{j}})$$

where θ is the consumer's sickness level, which is unobserved to insurers. Higher θs denote sicker individuals. $\mathbf{H_j} = \{H_{jk}\}_{k=1}^{|K_m|}$ is insurer j's network breadth across all services k, ζ is a discount factor, ρ is the probability that the consumer drops out of the contributory system, $\mathcal{P}$ is the transition probability from type θ in period t to type θ' in period $t + 1$, N_θ is the fixed market size of type-θ consumers, and C_j is the insurers' fixed cost, which mainly captures administrative expenses related to billing and auditing activities.

The profit per consumer type θ is:

$$\pi_{j\theta}(\mathbf{H_j}, \mathbf{H_{-j}}) = (R_\theta - (1 - r_\theta)AC_{j\theta}(\mathbf{H_j}))s_{j\theta}(\mathbf{H})$$

Here, R_θ is the risk-adjusted transfer from the government plus revenues from copayments, $AC_{j\theta}$ is insurer j's average cost for a type-θ consumer, r_θ is consumer θ's coinsurance rate, and $s_{j\theta}$ is insurer j's demand from type-θ consumers. Finally, $\mathbf{H} = \{\mathbf{H_j}\}_{j=1}^{|J|}$. Assume the demand and average cost functions are twice-continuously differentiable and that $\frac{\partial AC_{j\theta}}{\partial \theta} > 0$, $\frac{\partial AC_{j\theta}}{\partial H_{jk}} > 0$, $\frac{\partial C_j(\cdot)}{\partial H_{jk}} > 0$, $\frac{\partial^2 C_j(\cdot)}{\partial H_{jk}^2} > 0$, $\frac{\partial s_{j\theta}}{\partial H_{jk}} > 0$, and $\frac{\partial s_{j\theta}}{\partial H_{-jk}} < 0$.

Insurers compete in every market by choosing their service network breadth to maxi-

mize profits. The FOC of the insurer's problem is:

$$\frac{\partial \Pi_j}{\partial H_{jk}} = \underbrace{\sum_\theta \left((R_\theta - (1 - r_\theta)AC_{j\theta})\frac{\partial s_{j\theta}}{\partial H_{jk}} - (1 - r_\theta)s_{j\theta}\frac{\partial AC_{j\theta}}{\partial H_{jk}} \right) N_\theta}_{\text{Current profit derivative (CP)}} \tag{2}$$

$$+ \sum_{s=t+1}^{T} \zeta^s \underbrace{\sum_\theta (1 - \rho_{\theta'})\mathcal{P}(\theta'|\theta)\left((R_{\theta'} - (1 - r_{\theta'})AC_{j\theta'})\frac{\partial s_{j\theta'}}{\partial H_{jk}} - (1 - r_{\theta'})s_{j\theta'}\frac{\partial AC_{j\theta'}}{\partial H_{jk}} \right) N_{\theta'}}_{\text{Future profit derivative (FP)}} - \frac{\partial C_j}{\partial H_{jk}} = 0$$

Consider the first line of equation (2). Adverse selection manifests as the covariance between the consumer's valuation for network breadth and the insurer's average marginal cost, represented by $AC_{j\theta}\frac{\partial s_{j\theta}}{\partial H_{jk}} > 0$ and $s_{j\theta}\frac{\partial AC_{j\theta}}{\partial H_{jk}} > 0$, respectively. This covariance is positive and becomes more pronounced the broader is the network, because consumers of higher type have stronger preferences for network breadth. Therefore, adverse selection changes the composition of consumer types that enroll and incentivizes insurers to offer narrower networks.

The FOC also provides intuition on how market concentration—and perhaps market power—impacts service network breadth. Suppose for simplicity that insurers have the same average cost structure $AC_{j\theta} = AC_\theta$, and focus on the effects of a change in service network breadth weighted across insurers by their market share $s_{j\theta}$. We can rewrite equation (2) as:

$$\sum_\theta (R_\theta - (1 - r_\theta)AC_\theta)\left(\sum_j \frac{\partial s_{j\theta}}{\partial H_{jk}} s_{j\theta} \right) N_\theta - \sum_\theta (1 - r_\theta)\frac{\partial AC_\theta}{\partial H_{jk}} \overbrace{\left(\sum_j s_{j\theta}^2 \right)}^{HHI_\theta} N_\theta \tag{3}$$

$$+ \sum_{s=t+1}^{T} \zeta^s \sum_{j\theta} s_{j\theta} FP - \sum_{j\theta} s_{j\theta}\frac{\partial C_j}{\partial H_{jk}} = 0$$

Equation (3) illustrates that market concentration exacerbates the adverse selection effect when firms exhibit homogeneous cost structures. Specifically, the HHI has a multiplicative effect on the increase in insurers' average costs from providing broader networks. This

24

suggests that concentrated markets with adverse selection likely have narrower networks compared to less concentrated markets. However, in scenarios with heterogeneous costs and preferences, the effects of market concentration on network breadth become less clear. Furthermore, since market concentration may not accurately reflect true market power, empirical counterfactual analyses directly targeting market power—such as examining potential collusion among insurers—are needed to assess its impact on provider networks.

In the profit function, insurer demand follows a random utility representation. A new enrollee i of type θ has the following utility from enrolling with insurer j in market m:

$$u_{ijm} = \underbrace{\beta_{im} \sum_k q_{\theta k} H_{jkm} - \alpha_i c_{\theta jm}(\mathbf{H_{jm}}) + \phi_j}_{v_{\theta jm}} + \varepsilon_{ijm}$$

where $q_{\theta k}$ is the probability that a type-θ consumer claims service k, $c_{\theta jm}$ is the expected out-of-pocket cost at insurer j, $\mathbf{H_{jm}} = \{H_{jkm}\}_{k=1}^{K_m}$ with K_m denoting the set of services available in market m, ϕ_j is an insurer fixed effect, and ε_{ijm} is a type-I extreme value shock. Consumer types are defined by combinations of sex, age group (19-24, 25-29, 30-34, 35-39, 40-44, 45-49, 50-54, 55-59, 60-64, 65-69, 70-74, 75 or more), and diagnosis (cancer, cardiovascular disease, diabetes, renal disease, pulmonary disease, other disease, no diseases).[13] The coefficients in the utility function are given by $\beta_i = (x_i \; x_m)\beta$, $\alpha_i = x_i'\alpha$; the vector x_i includes dummies for sex, age group, and diagnoses, and x_m are market dummies.

Consumers' out-of-pocket cost are a function of service network breadth because they pay a fraction of the health service prices that their insurer negotiates with in-network

---

[13]In cases where a single individual has multiple health conditions, I assign the diagnosis that accounts for the highest share of their healthcare cost.

providers: $c_{\theta jm} = r_\theta AC_{\theta jm}(\mathbf{H_{jm}})$. In turn, insurers' average cost function $AC_{\theta jm}$ is:

$$\log(AC_{\theta jm}(\mathbf{H_{jm}})) = \tau_0\left(\sum_k^{K_m} q_{\theta k}A_k\right) + \tau_1\left(\sum_k^{K_m} q_{\theta k}H_{jkm}\right)$$

$$+ \frac{1}{2K_m}\tau_2\sum_k^{K_m}\sum_{l\neq k}^{K_m} q_{\theta k}q_{\theta l}H_{jkm}H_{jlm} + \lambda_\theta + \eta_m + \delta_j + \epsilon_{\theta jm} \qquad (4)$$

In the first term on the right-hand side of equation (4), A_k represents the government's reference price for service k. This price is used to reimburse providers for events not covered by health insurance (such as car accidents, natural disasters, and terrorist attacks), and it serves as the baseline in insurers' bilateral negotiations with providers. The second term captures whether insurers with broad networks negotiate higher prices with in-network providers compared to those with narrower networks, thereby summarizing the bargaining environment. The third term introduces the potential for insurers to benefit from economies of scope across services, which helps explain why some insurers choose relatively broad networks across multiple services. Moreover, λ_θ is a consumer type fixed effect, η_m is a market fixed effect, and δ_j is an insurer fixed effect. Finally, I assume $\epsilon_{\theta jm}$ is a mean-zero shock independent of ε_{ijm}.

My specification of insurers' average cost per enrollee is guided by trends observed in the raw data illustrated in Appendix Figure 2. This figure shows a positive relationship between log average costs and service network breadth, along with a negative relationship with the interaction between network breadth across pairs of services.

Given the distribution of the preference shock, insurer j's demand in market m among type-θ enrollees is

$$s_{j\theta m} = \frac{\exp(\nu_{\theta jm})}{\sum_{g=1}^{|J_m|}\exp(\nu_{\theta gm})}$$

where $|J_m|$ is the set of insurers in market m.

Finally, I parameterize insurers' fixed cost as:

$$C_{jm}(\mathbf{H_{jm}}, \boldsymbol{\xi_{jm}}) \equiv \sum_k \left(\omega H_{jkm} + \xi_j + \xi_m + \xi_{jkm}\right)H_{jkm}$$

26

where ξ_j and ξ_m are insurer- and market-specific cost components, $\xi_{jm} = \{\xi_{jkm}\}_{k=1}^{K_m}$, and ξ_{jkm} is an unobserved (to the econometrician) cost component. The fixed cost allows me to rationalize insurers that choose broad networks despite selection incentives and imperfect competition.

4.1 Estimation Results

Serna (2024) estimates the demand model on data from new enrollees with complete enrollment spells in 2011. Cost functions and market sizes in the profit function use information from all the continuously enrolled. Demand and cost estimates are provided in Appendix 2.3. Appendix Table 2 shows that consumers prefer broad networks and dislike out-of-pocket expenses. The preference for network breadth is lower among individuals without diseases, and the disutility for out-of-pocket costs is lower among individuals with chronic health conditions. These parameter estimates imply substantial heterogeneity in willingness-to-pay for service network breadth across consumers. For instance, patients with renal disease are willing to pay almost 7 times more for an additional provider in the network for renal care relative to a healthy patient, consistent with adverse selection on provider networks.

Estimation results for the average cost function are provided in Appendix Table 3. Average costs per enrollee are increasing in service network breadth at a decreasing rate. Thus, broad-network insurers negotiate higher service prices with in-network providers and enjoy some economies of scope across services. These scope economies might come from price discounts at providers where insurers cover several services. Appendix Tables 4 and 5 present estimates of dropout and transition probabilities, which are computed non-parametrically, outside of the model. These probabilities factor into the estimation of insurers' fixed costs in Appendix Table 6 using the FOC. Findings show that fixed costs vary significantly across insurers. In particular, insurers' fixed cost structure explains half of the variation in total profits when an insurer unilaterally increases network breadth for a particular service, while heterogeneity in willingness-to-pay coming from the demand function explains the other half. This suggests that adverse selection—sicker, less prof-

itable individuals choosing insurers with greater coverage in certain services—and cost incentives weigh equally on insurers' network breadth choices.

5 Centralized Equilibrium

The first step to assess the impact of insurer competition on network breadth is to establish the optimal service network breadth that a social planner would choose for each insurer. Although deriving a social welfare function and interpreting its implications are inherently complex tasks, I approximate the social planner's problem using the empirical model of section 4. One limitation is that potential externalities across consumers in their use of provider networks are not captured in the model, but a social planner would likely observe and consider them. For example, broad-network insurers potentially have lower congestion, improving consumers' ability to schedule doctor appointments on time. While these types of externalities are not captured, the model allows me to tractably solve the social planner's problem, something that no study to date has implemented.

The social planner's objective is to maximize consumer surplus subject to insurers' participation constraints holding fixed total risk-adjusted transfers. My proxy for consumer surplus is the long-run expected utility obtained from the demand model:

$$CS_m(H_m) = \sum_{\theta} \left(EU_i(H_m) N_{\theta m} + \sum_{s=t+1}^{T} \zeta^s \sum_{\theta'} (1 - \rho_{\theta' m}) \mathcal{P}(\theta'|\theta) EU_i'(H_m) N_{\theta' m} \right)$$

where the short-run expected utility, following McFadden (1996), is

$$EU_i = \log \left(\sum_{j} \exp(\beta_{im} \sum_{k} q_{\theta k} H_{jkm} - \alpha_i c_{\theta jm}(H_{jm}) + \phi_j) \right)$$

The social planner solves the following optimization problem per market:

$$\max_{H_m} CS_m(H_m) \tag{5}$$

$$\text{s.t } \overline{\Pi}_{jm}(H_m) \geq 0 \ \forall j$$

where $\overline{\Pi}_{jm}(H_m)$ is the average profit per enrollee. The participation constraints imply that insurers are in perfect competition and that if a potential entrant enters, both it and the existent set of insurers would incur losses.

The welfare maximization problem in equation (5) resembles the theoretical model of section 2 in that the social planner maximizes the sum of consumer surplus and insurer profits, while holding risk-adjusted transfers fixed, which essentially makes these transfers irrelevant for the solution. As in section 2, risk-adjusted payments cancel out in the social surplus function because they are linear transfers from the government to the insurer. Moreover, note that in solving equation (5) I do not guarantee that the social planner's solution is an equilibrium of the game between insurers, rather, the planner takes insurers' technology as given. Finally, the maximization problem differs from the theoretical model in that the planner's solution is implemented through several insurers and must guarantee that these insurers are willing to participate.

To reduce the computational burden, I solve the social planner's problem only in the capital city of Bogotá. Additionally, because the optimization routine involves searching over 240 parameters (20 services for each of 12 insurers), I redefine the procedure over 24 parameters, which correspond to network breadth for general medicine and hospital admissions for each insurer, holding network breadth for the rest of services fixed at their values in the observed equilibrium. I focus on general medicine and hospital admissions because these are services commonly used by both healthy and sick individuals. Redefining the social planner's problem in this way means that the solution will reflect a partial equilibrium.

Results are presented in Table 2. Panel A shows the percentage change relative to the observed scenario in average network breadth (across insurers and services), average costs per enrollee, total average costs, and long-run consumer surplus for individuals with and without diagnoses. Panel B presents the percentage change in average network breadth for specific services. I find that the social planner would choose networks for general medicine and hospital admissions that are almost twice as broad as in the observed scenario. In the case of general medicine this is an average increase from 0.33 to 0.61 across insurers. Network breadth for hospital admissions similarly moves from 0.33 to

TABLE 2: Networks, Costs, and Welfare for Social Planner

Variable	Centralized equilibrium
Panel A. Overall	
Average network breadth	11.09
Average cost per enrollee	8.02
Total average cost	7.93
Consumer surplus (with diagnoses)	18.74
Consumer surplus (without diagnoses)	19.09
Panel B. Service network breadth	
General medicine	84.96
Hospital admission	92.46
Other services	0.00

Note: Table presents the percentage change between the social planner's solution and the observed scenario in total average cost, total network formation costs, long-run consumer welfare for the healthy and sick, network breadth for general medicine, and network breadth for hospital admissions. The counterfactual is calculated with data from Bogotá only.

0.63. Appendix Table 7 shows the value of the participation constraints, corroborating that for most insurers the average per-enrollee profit is near zero.

The increase in coverage for general medicine and hospital admissions raises insurers' total average cost by 8% and long-run consumer surplus by around 19%. This indicates that enhancing coverage for these widely utilized services more than offsets the welfare losses consumers experience from higher out-of-pocket expenses. Importantly, the resulting gain in consumer surplus is relatively uniform across individuals with different health statuses: general medicine is more commonly used by those without diagnoses, while hospital admissions are more frequent among those with existing diseases. Nevertheless, both services see a similar expansion in network breadth, contributing equally to the observed surplus gains.

The trade-off between total costs and network breadth highlights one reason why the social planner's solution may not be attainable in practice for health systems with managed care. A policy that imposes complete network coverage in some services is costly and may generate incentives for insurers to drop coverage of other services altogether. Although my counterfactual results in Table 2 can not speak to these latter incentives, adverse selection suggests that this is one way in which insurers may respond to network adequacy rules requiring complete networks in highly claimed services.

6 Collusive Equilibrium

I now turn to quantifying how changes in the level of insurer competition affect service network breadth relative to the social planner's benchmark ("first best"). If achieving the first-best solution is impractical due to administrative costs or other hassle costs, two key questions arise: First, can a decentralized equilibrium, in which insurers compete on service network breadth, achieve the first-best outcome? Second, if so, what level of competition is necessary to attain this first-best solution?

To address these questions, I use the empirical model to simulate a counterfactual scenario in which there is an insurance monopoly and in which insurers collude, approximating the solution outlined in the theoretical model. The impact of joint profit maximization on service network breadth is not immediately clear given the substantial heterogeneity in preferences and costs. On the one hand, we might predict that collusion would lead to narrower networks, as the colluding firms would internalize the negative externality they impose on their competitors' demand. On the other hand, economies of scope and scale could yield cost efficiencies that encourage colluding firms to expand their network breadth in one service but not another.

To derive the potential impact of imperfect competition from the econometric model, take one market with two firms j and g. When firms collude, they solve following optimization problem:

$$\max_{H_j, H_g} \ \Pi_j(\mathbf{H_j}, \mathbf{H_g}, \mathbf{H_{-jg}}) + \Pi_g(\mathbf{H_j}, \mathbf{H_g}, \mathbf{H_{-jg}})$$

where $\mathbf{H_j} = \{H_{jk}\}_{k=1}^{K}$ and $\mathbf{H_{-jg}}$ denotes the vector of network breadth for all other firms besides j and g. In the FOC for the merged firm, the derivative of per-enrollee profits with respect to H_{jk} is:

$$\frac{\partial \pi_\theta^*}{\partial H_{jk}} = (R_\theta - (1 - r_\theta) AC_{\theta j}^*) \frac{\partial s_{\theta j}^*}{\partial H_{jk}} - (1 - r_\theta) s_{\theta j}^* \frac{\partial AC_{\theta j}^*}{\partial H_{jk}} + (R_\theta - (1 - r_\theta) AC_{\theta g}^*) \frac{\partial s_{\theta g}^*}{\partial H_{jk}} \quad (6)$$

The upper-script (*) denotes objects that are evaluated in equilibrium. The first term to the

right-hand side of equation ([6]) maps to the theoretical model's selection effect; it captures changes in the composition of marginal consumers weighted by their profitability. The second term describes how collusion may affect the colluding firm's cost structure. If $\frac{\partial AC_{\theta j}}{\partial H_{jk}} > \frac{\partial AC^*_{\theta j}}{\partial H_{jk}}$, then the new equilibrium may be characterized by broader networks because the colluding firm enjoys greater economies of scope. The third term captures the externality that firm j imposes on firm g's per-enrollee profits and maps to the theoretical model's competitive effect. Because $\frac{\partial s_{\theta g}}{\partial H_{jk}} < 0$, the merged firm internalizes the reduction in g's demand when j increases its network breadth. Therefore, collusion can lead the merged firm to choose narrower networks relative to the scenario where firms compete separately.

The ambiguous predictions of how imperfect competition impacts insurers' equilibrium choices also rest on my assumptions on how the merged firm's cost structure relates to firm j's and g's costs. In my empirical analysis, I approximate the merged firm's costs in several ways to test the importance of average and fixed cost heterogeneity as well as the stability of my results. In the first scenario labelled "Average," I assume the merged firm is an average of individual firms within the collusive agreement, so its average cost per enrollee and its fixed cost have a firm fixed effect of $\delta_{\text{merger}} = N^{-1} \sum_{j \in \text{merger}} \delta_j$ and $\xi_{\text{merger}} = N^{-1} \sum_{j \in \text{merger}} \xi_j$, respectively, where N is the number of firms in the agreement. In the second scenario labelled "P25 FE," I assume the merged firm is as efficient as the firm in the 25th percentile of the distribution of firm fixed effects among those in the collusive agreement. Finally, I assume the merged firm does not accrue any cost efficiencies, so $\delta_{\text{merger}} = \max_{j \in \text{merger}}\{\delta_j\}$ and $\xi_{\text{merger}} = \max_{j \in \text{merger}}\{\xi_j\}$. I denote this last scenario as "Max FE."

For the sake of tractability, I conduct these counterfactual analyses in the city of Bogotá. Results are summarized in Table [3]. Panel A shows the percentage change relative to the observed scenario in average network breadth, average costs per enrollee, total average costs, and long-run consumer surplus for individuals with and without diagnoses. Panel B shows the percentage change relative to the observed scenario in network breadth for specific services. Columns (1) to (3) assume all 12 insurers maximize joint profits under the different assumptions regarding the merged firm's cost structure. Columns (4) to (6)

show results assuming only the bottom 5 insurers collude (Insurers C, D, E, L, M).

In line with the intuition derived from equation (6) and from the theoretical model in section 2, joint profit maximization leads to lower average network breadth in equilibrium because the merged firm internalizes the negative externality it imposes on its competitors. However, as the cost efficiencies achieved by the merged firm increase, the reduction in network breadth from imperfect competition decreases. Under an insurance monopoly, average network breadth across insurers and services falls by 69% relative to the observed scenario when the monopolist is as efficient as the firm in the 25th percentile of the distribution of firm fixed effects. As this efficiency decreases by imposing the average fixed effect and the maximum fixed effect, reductions in average network breadth are as large as 84%. In these cases, the market essentially unravels.

Although network breadth falls across the board, column (1) shows that reductions are larger among entry-level services with high baseline network breadth such as general medicine and laboratory testing relative to complex care such as renal and cardiac care. Given that individuals with diagnoses have higher claim probabilities across all services, their long-run surplus falls by a slightly greater magnitude than for individuals without diagnoses.

When the bottom 5 insurers engage in joint profit maximization, I find qualitatively similar results. In column (4) assuming the merged firm is as efficient as the firm in the 25th percentile of the distribution, average network breadth decreases 3.2%, with reductions being larger among services that mostly individuals with diagnoses tend to claim. The decrease in coverage generates lower average costs per enrollee and lower total average costs because the direct effect of network breadth on average costs is greater than the impact of scope economies. In this case, long-run consumer surplus for individuals with and without diagnoses increases by a moderate amount because of the lower healthcare costs that are passed-through to consumers. As the merged firm becomes more inefficient in columns (5) and (6), network breadth falls by a greater magnitude and consumers experience declines in surplus. Importantly, while network breadth decreases when the bottom 5 insurers collude, in line with the theory, these reductions are economically small, suggesting the observed scenario is not far from a situation in which the number of firms

33

TABLE 3: Networks, Costs, and Welfare under Decentralized Equilibria

Variable	Monopoly			Collusion		
	P25 FE	Average FE	Max FE	P25 FE	Average FE	Max FE
	(1)	(2)	(3)	(4)	(5)	(6)
Panel A. Overall						
Average network breadth	-68.74	-84.17	-84.35	-3.24	-4.07	-16.02
Average cost per enrollee	-24.81	-20.97	-0.07	-4.85	-0.84	7.38
Total average cost	-22.92	-19.57	2.24	-3.56	0.52	3.40
Consumer surplus (with diagnoses)	-35.75	-41.67	-42.26	1.58	1.28	-2.77
Consumer surplus (without diagnoses)	-34.40	-39.95	-40.36	1.48	1.24	-2.69
Panel B. Service network breadth						
Otorhinolaryngologic care	-83.94	-95.49	-95.59	-4.37	-5.17	-17.28
Cardiac care	-61.70	-79.42	-79.64	-4.35	-5.27	-17.62
Gastroenterologic care	-73.07	-88.78	-88.97	-4.32	-5.24	-17.12
Renal care	-62.85	-81.86	-82.09	-4.60	-5.57	-18.55
Gynecologic care	-71.92	-89.01	-89.17	-3.45	-4.39	-16.64
Orthopedic care	-64.35	-81.99	-82.24	-4.40	-5.32	-17.60
Imaging	-73.98	-83.90	-84.03	-1.77	-2.17	-10.89
General medicine	-77.98	-91.10	-91.19	-2.79	-2.95	-10.53
Laboratory	-74.28	-84.61	-84.74	-2.29	-2.60	-11.03
Hospital admission	-67.54	-82.79	-82.99	-2.77	-3.55	-13.89

Note: Panel A presents the percentage change relative to the observed scenario in average network breadth, average cost per enrollee, total average cost, and long-run consumer surplus for sick and healthy individuals, in the scenario without risk adjustment in column (1), the scenario with improved risk adjustment in column (2), the scenario with homogeneous average costs in column (3), and the scenario with homogeneous network formation costs in column (4). Panel B presents the percentage change relative to the observed scenario in average network breadth for a few service categories. Simulations use data from Bogotá.

halves and these firms have market power. Appendix Figure 3 presents the distribution of service network breadth in each scenario.

Table 4 explores what happens with each firm in the collusive agreements. Panel A presents the percentage change relative to the observed scenario in average network breadth (across services) and total variable profits for the counterfactual in which there is an insurance monopoly with average efficiency.[14] Panel B presents these statistics when the bottom 5 insurers maximize joint profits and I impose the average firm fixed effect in their cost structure. Each firm that makes up the insurance monopoly substantially reduces their network breadth relative to the observed scenario, and in some cases shuts down. These reductions range from 75% for Insurer C to 100% for insurer N. Consistent

---

[14]The change in total variable profits is calculated as the change in total revenues minus the change in total variable costs.

TABLE 4: Networks and Profits for Colluding Firms

Insurer	Network breadth (1)	Variable profits (2)
Panel A. Monopoly, Avg FE		
Insurer A	-85.02	39.93
Insurer B	-95.22	8.98
Insurer C	-74.90	17.14
Insurer D	-80.18	7.64
Insurer E	-83.07	50.39
Insurer G	-87.09	16.59
Insurer I	-97.04	11.69
Insurer J	-100.00	24.23
Insurer K	-66.47	7.70
Insurer L	-100.00	35.98
Insurer M	-61.27	7.65
Insurer N	-100.00	17.22
Panel B. Collusion, Avg FE		
Insurer C	2.55	-3.16
Insurer D	-8.52	-9.74
Insurer E	24.16	21.17
Insurer L	-96.64	5.26
Insurer M	16.09	-25.76

Note: Table presents the percentage change in average network breadth, total profits, and short-run average cost per enrollee for the insurers that collude.

with joint profit maximization generating higher profits for each individual firm, Panel A, column (2) shows that variable profits increase between 7% and 50% across insurers. For insurers J, L, and N, for which network breadth completely collapses, the change in total variable profits essentially represents their scrap value.

In Panel B, I find that collusion among the bottom 5 insurers, results in one of these insurers shutting down (Insurer L), and the rest absorbing its demand and increasing network breadth. The reduction in coverage documented in Table 3 is therefore explained by the best response of insurers that are not in the collusive agreement. Column (2) shows that only 2 out of the 5 insurers that maximize joint profits see increases in total variable profits. However, the average change in variable profits weighted by demand is positive and equal to 1.32%.

The findings in this section show that collusion exacerbates risk selection incentives. Insurers in this market engage in risk selection by offering narrower networks for less

profitable services. Thus, the significant decline in network breadth across services when firms maximize joint profits suggests lower levels of competition enable risk selection. Conversely, findings also suggest that a market equilibrium with strong competition between private health insurers, even if premiums and cost-sharing are regulated, can more closely approximate the social planner's solution. Table 2 showed that the social planner would choose around 60% coverage for general medicine and hospital admissions (holding other services fixed), while table 3 indicates that network breadth for these two services when the bottom 5 insurers collude would be 28 percentage points farther away from the first-best.

7 Network Adequacy Rules

Encouraging competition among insurers to achieve broader provider networks can be challenging from a policy perspective. As an alternative, the social planner can develop regulations that directly address risk selection within the existing market structure, allowing insurers to endogenously respond to the regulations. In Colombia, where insurers can differentiate networks based on health services, restricting their leverage across services may help mitigate risk selection incentives.

In this section, I examine the impacts of a network adequacy rule that forces insurers to offer the same network breadth for hospital admissions and general medicine, which resembles the social planner's solution derived in section 5.[15] Given the model estimates, network adequacy rules that mandate coverage of specific providers will likely result in greater coverage after accounting for endogenous supply responses because insurers enjoy substantial economies of scale in the number of covered providers.

Formally, I impose that each insurer selects identical networks for general medicine and hospital care, i.e., $H_{\text{gen. med.}} = H_{\text{hosp}}$ without placing restrictions on the ultimate network breadth levels. Instead, the model determines insurers' optimal responses given this network adequacy constraint. However, the solution depends on assumptions regarding the

---

[15]Health plans in the US Health Insurance Exchanges are considered to have broad hospital networks when they cover more than 70% of hospitals in a market (Bauman et al., 2014).

fixed costs insurers face for these two services. If insurers choose to cover both hospital admissions and general medicine through the same provider, it is reasonable to assume they incur one average fixed cost across these services rather than separate ones. Conversely, if insurers contract with additional providers for only one of the services, they would face distinct fixed costs for general medicine and hospital care. In my implementation, I adopt the former assumption.

Table 5 presents the results of this exercise. Panel A shows that network breadth would increase across insurers on average. This is because network breadth for hospital admissions is prioritized relative to network breadth for general medicine, resulting in a 17% increase for the former and a 1% decrease for the latter as seen in Panel B. However, insurers also respond to the policy by decreasing network breadth for other services, although the reduction is economically small. The increase in coverage for hospital admissions coupled with the slight decrease among other services results in slightly higher total average costs for insurers. Despite reductions in coverage for other services, long-run consumer surplus rises for both types of consumers, with increases being larger among those with chronic diseases. This suggests that guaranteeing broad networks across *all* services is not needed to improve welfare but perhaps only across a few key services.

Eliminating insurers' leverage for risk selection generates an equilibrium that is closer to the social planner's solution from Table 2 but does not fully implement it. Consumer surplus increases under the network adequacy rule but not by a similar magnitude as in the social planner's solution. One concern with fully implementing the centralized outcome is that the health system may incur substantial administrative costs. For example, the network adequacy rule considered in this section raises insurers' fixed costs by 10% on average. The lesson from this exercise is that designing network adequacy regulations should consider its impacts on health system fiscal sustainability.

TABLE 5: Networks, Costs, and Welfare Under Network Adequacy

Variable	Network adequacy
Panel A. Overall	
Average network breadth (all)	0.72
Average network breadth (rest)	-0.22
Average cost per enrollee	-0.04
Total average cost	0.15
Consumer surplus (with diagnoses)	0.26
Consumer surplus (without diagnoses)	0.19
Panel B. Service network breadth	
Otorhinolaryngologic care	-0.22
Cardiac care	-0.21
Gastroenterologic care	-0.21
Renal care	-0.22
Gynecologic care	-0.22
Orthopedic care	-0.21
Imaging	-0.16
General medicine	-1.11
Laboratory	-0.15
Hospital admission	17.07

Note: Panel A presents the percentage change in average network breadth and long-run consumer welfare for sick and healthy individuals, in the scenario with a network adequacy prohibiting discrimination of networks across services. Panel B presents the percentage change in mean network breadth by service category.

7.1 Discussion

Does the theoretical framework developed in this paper and the empirical application have implications outside Colombia? Many health insurance markets are characterized by strict price regulations that essentially eliminate insurer competition on that dimension. In Colombia, the health system without premiums was designed to allegedly promote competition on quality. I summarize quality with a measure of provider network breadth, which captures in a reduced-form way the more complex bilateral negotiations between insurers and providers. Competition in provider networks under fixed prices is not unique to Colombia. For example, in Medicaid managed care in the US, premiums and cost-sharing are indexed to the enrollee's income and in most cases are equal to zero. Managed care companies compete by choosing which providers to include in their networks and receive risk-adjusted transfers from state governments that fail to adequately compensate for quality (if at all). Medicare Advantage has a similar setting although insurers tend to be much more efficient. This suggests that lessons from this paper could be considered

for these markets as well. Even though the specific estimates and magnitudes would not apply, the general direction of the impact of competition and network adequacy rules on equilibrium quality contributes to the bigger discussion of whether to use and promote competition in insurance markets.

8   Conclusion

This paper revisits the implications of imperfect competition on equilibrium quality in markets with adverse selection and without prices. I build on the insights developed by Veiga and Weyl (2016) and Mahoney and Weyl (2017) to show that without prices, imperfect competition can harm social welfare because firms do not internalize consumers' valuation for quality. Focusing on health insurance markets and defining quality as the breadth of health insurers' network of covered providers, I find that a social planner would choose broader networks than an insurance monopoly, while the duopoly brings network breadth closer to the planner's solution.

I test the predictions of the theory with a structural model of the health insurance market in Colombia, borrowing estimates from Serna (2024). I empirically characterize optimal network breadth as chosen by the social planner and contrast it with different decentralized equilibria in which insurers compete. As in the theoretical framework, I find that under an insurance monopoly, network breadth collapses and this effect holds under different assumptions about the monopolist's cost structure. Instead, a social planner would choose networks that are twice as broad as in the observed scenario. For example, the planner would increase network breadth for general medicine from 0.3 to 0.6. Finally, I evaluate the impact of a network adequacy rule forcing insurers to provide the same network breadth for general medicine and hospital admissions. Findings show that network breadth would increase 17% for hospital admissions and decrease 1% for general medicine, resulting in marginal improvements in consumer surplus but not fully implementing the planner's solution.

My results indicate that network breadth is an increasing function of the degree of health insurer competition in markets where premiums are heavily regulated. Nonethe-

less, policies that approximate the social planner's solution might involve substantial administrative costs for these types of health systems. Therefore, the normative implications of greater coverage and higher administrative costs should be considered as well.

References

AZEVEDO, E. M. AND D. GOTTLIEB (2017): "Perfect Competition in Markets with Adverse Selection," *Econometrica*, 85, 67–105.

BAUMAN, N., E. COE, J. OGDEN, AND A. PARIKH (2014): "Hospital Networks: Updated National View of Configurations on the Exchanges," *McKinsey Report*, accessed through https://www.mckinsey.com/industries/healthcare/our-insights/hospital-net works-updated-national-view-of-configurations-on-the-exchanges/.

BROWN, J., M. DUGGAN, I. KUZIEMKO, AND W. WOOLSTON (2014): "How Does Risk Selection Respond to Risk Adjustment? New Evidence from the Medicare Advantage Program," *American Economic Review*, 104, 3335–3364.

BUNDORF, M. K., J. LEVIN, AND N. MAHONEY (2012): "Pricing and Welfare in Health Plan Choice," *American Economic Review*, 102, 3214–3248.

CABRAL, M., M. GERUSO, AND N. MAHONEY (2018): "Do Larger Health Insurance Subsidies Benefit Patients or Producers? Evidence from Medicare Advantage," *American Economic Review*, 108, 2048–2087.

CAPPS, C., D. DRANOVE, AND M. SATTERTHWAITE (2003): "Competition and Market Power in Option Demand Markets," *RAND Journal of Economics*, 737–763.

CUESTA, J. I., C. E. NOTON, AND B. VATTER (2024): "Vertical Integration and Plan Design in Healthcare Markets," National Bureau of Economic Research.

CUESTA, J. I. AND A. SEPÚLVEDA (2021): "Price Regulation in Credit Markets: A Trade-off Between Consumer Protection and Credit Access," *Available at SSRN 3282910*.

CUTLER, D. M. AND S. J. REBER (1998): "Paying for Health Insurance: The Trade-Off between Competition and Adverse Selection," *The Quarterly Journal of Economics*, 113, 433–466.

DAFNY, L., J. GRUBER, AND C. ODY (2015): "More Insurers Lower Premiums: Evidence from Initial Pricing in the Health Insurance Marketplaces," *American Journal of Health Economics*, 1, 53–81.

EINAV, L., A. FINKELSTEIN, AND N. MAHONEY (2021): "Chapter 14 - The IO of selection markets," *Handbook of Industrial Organization*, 5, 389–426.

EINAV, L., M. JENKINS, AND J. LEVIN (2012): "Contract Pricing in Consumer Credit Markets," *Econometrica*, 80, 1387–1432.

FANG, H. AND Z. WU (2018): "Multidimensional Private Information, Market Structure, and Insurance Markets," *The RAND Journal of Economics*, 49, 751–787.

FREED, M., J. BINIEK, A. DAMICO, AND T. NEUMAN (2025): "Medicare Advantage 2025 Spotlight: A First Look at Plan Premiums and Benefits," *Kaiser Family Foundation Report*.

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

GOWRISANKARAN, G., A. NEVO, AND R. TOWN (2015): "Mergers When Prices Are Negotiated: Evidence from the Hospital Industry," *American Economic Review*, 105, 172–203.

GRENNAN, M. (2013): "Price Discrimination and Bargaining: Empirical Evidence from Medical Devices," *American Economic Review*, 103, 145–177.

GUTH, M., M. AMMULA, AND E. HINTON (2021): "Understanding the Impact of Medicaid Premiums & Cost-Sharing: Updated Evidence from the Literature and Section 1115 Waivers," *Kaiser Family Foundation Report*.

HANDEL, B., I. HENDEL, AND M. D. WHINSTON (2015): "Equilibria in Health Exchanges: Adverse Selection versus Reclassification Risk," *Econometrica*, 83, 1261–1313.

HO, K. AND R. LEE (2017): "Insurer Competition in Health Care Markets," *Econometrica*, 85, 379–417.

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

Kong, E., T. Layton, and M. Shepard (2024): "Adverse Selection and (Un) Natural Monopoly in Insurance Markets," Working Paper.

MacLeod, W. B. (2021): "A Welfare Analysis of Competitive Insurance Markets with Vertical Differentiation and Adverse Selection," National Bureau of Economic Research.

Mahoney, N. and G. Weyl (2017): "Imperfect Competition in Selection Markets," *The Review of Economics and Statistics*, 99, 637–651.

McFadden, D. (1996): "Computing Willingness-to-Pay in Random Utility Models," *University of California at Berkeley, Econometrics Laboratory Software Archive, Working Papers*.

Rothschild, M. and J. Stiglitz (1976): "Equilibrium in Competitive Insurance Markets: An Essay on the Economics of Imperfect Information*," *The Quarterly Journal of Economics*, 90, 629–649.

Saltzman, E., A. Swanson, and D. Polsky (2021): "Inertia, Market Power, and Adverse Selection in Health Insurance: Evidence from the ACA Exchanges," National Bureau of Economic Research.

Serna, N. (2024): "Determinants of Healthcare Provider Networks: Risk Selection vs. Administrative Costs," .

Shaked, A. and J. Sutton (1982): "Relaxing Price Competition Through Product Differentiation," *The Review of Economic Studies*, 49, 3–13.

Spence, M. (1975): "Monopoly, Quality, and Regulation," *The Bell Journal of Economics*, 6, 417–429.

Stiglitz, J. and J. Yun (2013): "Optimality and Equilibrium in a Competitive Insurance Market under Adverse Selection and Moral Hazard," National Bureau of Economic Research.

VEIGA, A. AND E. G. WEYL (2016): "Product Design in Selection Markets," *The Quarterly Journal of Economics*, 131, 1007–1056.

Supplemental Appendix

Appendix 1   Additional Descriptives

APPENDIX TABLE 1: List of services

Service code	Description
01	**Neurosurgery:** Procedures in skull, brain, and spine
02	**Other neurologic care:** Procedures in nerves and glands
03	**Otorhinolaryngologic care:** Procedures in face and trachea
04	**Pneumologic care:** Procedures in lungs and thorax
05	**Cardiac care:** Procedures in cardiac system
06	**Angiologic care:** Procedures in lymphatic system and bone marrow
07	**Gastroenterologic care:** Procedures in digestive system
08	**Hepatologic care:** Procedures in liver, pancreas, and abdominal wall
09	**Renal care:** Procedures in urinary system
10	**Gynecologic care:** Procedures in reproductive system
11	**Orthopedic care:** Procedures in bones and joints
12	**Other orthopedic care:** Procedures in tendons, muscles, and breast
13	**Diagnostic aid:** Diagnostic procedures in skin and subcutaneous cellular tissue
14	**Imaging:** Radiology and non-radiology imaging
15	**Internal and general medicine:** Consultations
16	**Laboratory:** Laboratory and blood bank
17	**Nuclear medicine:** Nuclear medicine and radiotherapy
18	**Rehab and mental health:** Rehabilitation, mental health care, therapy
19	**Therapy (chemo and dialysis):** Prophylactic and therapeutic procedures
20	**Hospital admissions:** Inpatient services
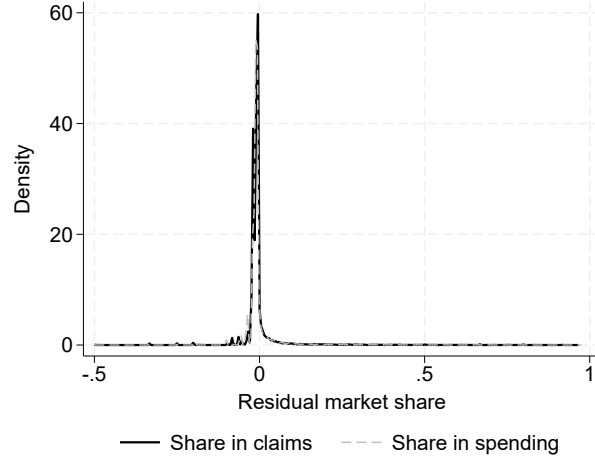
APPENDIX FIGURE 1: Residual Provider Market Share

Note: Figure presents the distribution of residuals from a regression of provider market shares in the number of claims (in black) and of provider market share in total health care spending (in gray) on market-by-service fixed effects.

(A) Network breadth

(B) Network breadth interactions
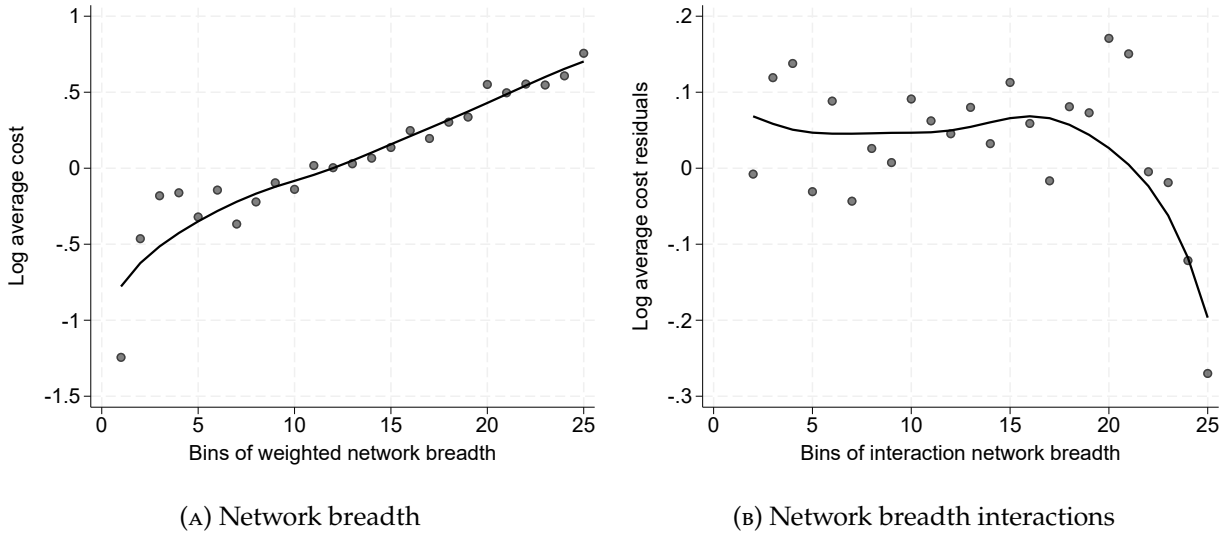
APPENDIX FIGURE 2: Empirical Relation of Log Average Cost per Enrollee

Note: Panel A presents a scatter plot of log average cost per enrollee averaged within 20 bins of weighted service network breadth, $\sum_k q_{\theta k} H_{jkm}$. Panel B presents a scatter plot of log average cost per enrollee averaged within 20 bins of weighted interactions of network breadth across pairs of services, $\sum_k \sum_l q_{\theta k} q_{\theta l} H_{jkm} H_{jlm}$. The solid line represents a linear fit.

45

Appendix 2   Model Summary

In this Appendix I describe additional details of the empirical model of insurer competition in service-level network breadth presented in Serna (2024).

2.1   Service Claim Probabilities

I estimate the claim probability, $q_{\theta k}$, outside of the model. I use data from all enrollees in 2010 and 2011 to estimate the following logistic regression:

$$\text{logit(any claims)}_{ik} = \psi_k + \psi_\theta + \epsilon_{ik}$$

The dependent variable is an indicator for whether patient i makes a claim for service k, and ψ_k and ψ_θ are service and consumer type fixed effects, respectively.

2.2   Identification

The main source of variation that identifies the preference for service network breadth in the demand model is the variation in claim probabilities $q_{\theta k}$ across consumer types and markets. These claim probabilities are plausibly exogenous to the extent that diseases considered in the model require explicit treatment guidelines and therefore do not vary with service network breadth. Insurer fixed effects also absorb some of the endogenous variation in service network breadth that stems from insurer competition in every market, allowing identification to come from the exogenous claim probabilities.

One concern related to identifying the coefficient on out-of-pocket costs in the demand model is variation in provider quality. For example, if an insurer covers a high-quality provider, then we would likely see high demand for that insurer (because consumers value having access to high-quality provider) as well as high out-of-pocket costs (because the provider has a relatively high bargaining power), which would bias α_i towards zero. Variation in provider quality introduces endogenous variation in service network breadth across insurers. Thus, the inclusion of insure fixed effects help isolate this source of endogenous variation. The coefficient α_i is then identified from exogenous variation in

income across consumers within a market, which generates variation in the coinsurance rates.

For the average cost function, coefficients are identified from variation in average costs within insurer and across consumer types. The rich set of fixed effects included in this function account for potential unobserved cost variation within consumer types. Intuitively, identification of the average cost parameters requires observing two insurers that are identical (in terms of the characteristics of their enrollees) except for their network breadth.

Identification of the fixed cost relies on systematic variation in marginal variable profits across services within an insurer. For example, if there are two insurers that have identical health risk, but one insurer offers low service network breadth and another offers high service network breadth across all services, then the model would rationalize these choices with high fixed costs for the former and low fixed costs for the latter. To identify ω, I rely on instrumental variables because insurers choose service network breadth with knowledge of their cost shocks ξ_{jkm}. My instrument is the average claim probability among healthy consumers. Note that claim probabilities affect marginal variable profits only through their interaction with service network breadth.

2.3 Model Estimates

The following tables summarize the model estimates from Serna (2024). Appendix Table 2 shows estimation results for the insurer demand model. I find that consumers have preferences for broad networks and dislike out-of-pocket costs. The estimates imply that willingness-to-pay for service network breadth—defined as $\frac{1}{-\alpha_i} \frac{\partial s_{ijm}}{\partial H_{jkm}} \frac{H_{jkm}}{s_{ijm}}$—is higher for individuals with chronic conditions than for individual without diagnoses, consistent with adverse selection.

Appendix Table 3 shows results of insurers' average cost per consumer type. Average costs are increasing in service network breadth in line with broad-network insurers negotiating higher prices with providers, but at a decreasing rate, in line with insurers enjoying economies of scope across services.

Appendix Tables 4 and 5 present summary statistics of dropout and transition probabilities which enter the computation of future profits. These probabilities are estimated outside of the model as follows. For computing both probabilities, I use the enrollment data for all enrollees between 2010 and 2011. The probability that a consumer type θ drops out is the fraction of consumers type θ enrolled in 2010 but not enrolled in 2011. The transition probability is the fraction of consumers type θ in 2010 that turn into θ' in 2011.

APPENDIX TABLE 2: Insurer Demand

Variable	Network breadth		OOP spending (million)	
	coef	se	coef	se
Mean	3.429	(0.021)	-1.602	(0.117)
Interactions				
Male	0.543	(0.010)	0.121	(0.066)
Cancer	-0.601	(0.013)	0.003	(0.092)
Cardiovascular	-0.901	(0.011)	-0.205	(0.075)
Diabetes	-0.464	(0.023)	0.008	(0.105)
Other disease	-0.783	(0.015)	0.465	(0.068)
Pulmonary	-0.610	(0.031)	0.841	(0.095)
Renal	0.039	(0.037)	0.873	(0.069)
Age 19-24	0.055	(0.020)	0.566	(0.158)
Age 25-29	-0.575	(0.019)	0.326	(0.120)
Age 30-34	-0.560	(0.019)	0.338	(0.130)
Age 35-39	-0.456	(0.020)	-0.312	(0.215)
Age 40-44	-0.356	(0.020)	0.442	(0.161)
Age 45-49	-0.384	(0.019)	0.237	(0.141)
Age 50-54	-0.324	(0.020)	0.314	(0.133)
Age 55-59	-0.246	(0.021)	0.477	(0.122)
Age 60-64	-0.147	(0.023)	0.160	(0.121)
N		5200890		
Pseudo-R^2		0.112		

Note: Table presents a conditional logit model of insurer choice estimated by maximum likelihood on a random sample of 500,000 new enrollees. An observation is a combination of individual and insurer. Specification includes insurer fixed effects. Robust standard errors in parenthesis.

APPENDIX TABLE 3: Insurer Average Costs Per Enrollee

Variable	Log average cost per enrollee	
	coef	se
Service network breadth	0.274	(0.047)
Scope economies	-1.100	(0.580)
Reference price	3.372	(0.544)
Insurer FE		
Insurer A	0.156	(0.038)
Insurer B	-0.084	(0.022)
Insurer C	0.020	(0.025)
Insurer D	-0.137	(0.027)
Insurer E	0.279	(0.047)
Insurer F	-0.012	(0.062)
Insurer G	0.062	(0.033)
Insurer H	0.043	(0.040)
Insurer I	-0.002	(0.017)
Insurer J	0.099	(0.020)
Insurer K	-0.135	(0.033)
Insurer L	0.116	(0.035)
Insurer M	-0.059	(0.032)
Constant	-1.368	(0.097)
Consumer type FE	Yes	
Market FE	Yes	
N	18369	
R^2	0.611	

Note: Table presents OLS regressions of the log of average costs per consumer type on service network breadth, the measure economies of scope, and the service reference price. An observation is a combination of insurer, consumer type, market and year. Estimation uses data from all continuously enrolled individuals in 2010 and 2011. Specification includes consumer type fixed effects, market fixed effects, and insurer fixed effects. Robust standard errors in parenthesis.

APPENDIX TABLE 4: Dropout Probabilities

	mean	sd
Female	0.084	(0.131)
Male	0.106	(0.165)
Age 19-24	0.120	(0.177)
Age 25-29	0.087	(0.134)
Age 30-34	0.081	(0.135)
Age 35-39	0.085	(0.141)
Age 40-44	0.085	(0.146)
Age 45-49	0.085	(0.149)
Age 50-54	0.089	(0.153)
Age 55-59	0.091	(0.158)
Age 60-64	0.092	(0.158)
Age 65-69	0.096	(0.159)
Age 70-74	0.104	(0.160)
Age 75+	0.124	(0.164)
Cancer	0.048	(0.024)
Diabetes	0.027	(0.008)
Cardiovascular	0.028	(0.009)
Pulmonary	0.040	(0.015)
Renal	0.044	(0.018)
Other disease	0.026	(0.011)
Healthy	0.450	(0.073)

Note: Mean and standard deviation in parenthesis of dropout probabilities conditional on diagnosis in the first panel, age group in the second panel, and sex in the third panel.

APPENDIX TABLE 5: Transition Probabilities

Diagnosis	Stat	Cancer	Cardio	Diabetes	Renal	Pulmonary	Other	Healthy
Cancer	mean	0.316	0.017	0.139	0.014	0.007	0.047	0.460
	sd	0.067	0.014	0.090	0.013	0.006	0.019	0.176
Diabetes	mean	0.030	0.557	0.170	0.009	0.013	0.021	0.200
	sd	0.026	0.078	0.100	0.010	0.011	0.010	0.140
Cardio	mean	0.043	0.028	0.554	0.014	0.011	0.034	0.316
	sd	0.036	0.018	0.205	0.012	0.010	0.009	0.224
Pulmonary	mean	0.055	0.019	0.191	0.234	0.007	0.078	0.416
	sd	0.046	0.014	0.089	0.152	0.006	0.034	0.231
Renal	mean	0.044	0.036	0.214	0.012	0.371	0.058	0.265
	sd	0.035	0.030	0.132	0.013	0.062	0.031	0.154
Other	mean	0.056	0.016	0.156	0.023	0.008	0.343	0.398
	sd	0.040	0.013	0.106	0.020	0.004	0.058	0.095
Healthy	mean	0.055	0.012	0.108	0.014	0.004	0.045	0.762
	sd	0.042	0.008	0.068	0.014	0.003	0.021	0.109

Note: Table presents mean and standard deviation in parenthesis of transition probabilities across diagnoses. Summary statistics are calculated across sex-age combinations in each cell.

Using the demand, average costs, and dropout and transition probability estimates, I forward simulate the insurers' profit function and marginal variable profits for 100 periods. I estimate the remaining parameters associated with the fixed cost from insurers' FOC below:

$$\text{MVP}_{jkm} = \omega H_{jkm} + \xi_j + \xi_m + \xi_{jkm}$$

where

$$\text{MVP}_{jkm} \equiv \sum_i \left(\frac{\partial \pi_{ijm}}{\partial H_{jkm}} N_{\theta m} + \sum_{s=t+1}^{T} \zeta^s \sum_{\theta'} (1 - \rho_{\theta'm}) \mathcal{P}(\theta'|\theta) \frac{\partial \pi'_{ijm}}{\partial H_{jkm}} N_{\theta'm} \right)$$
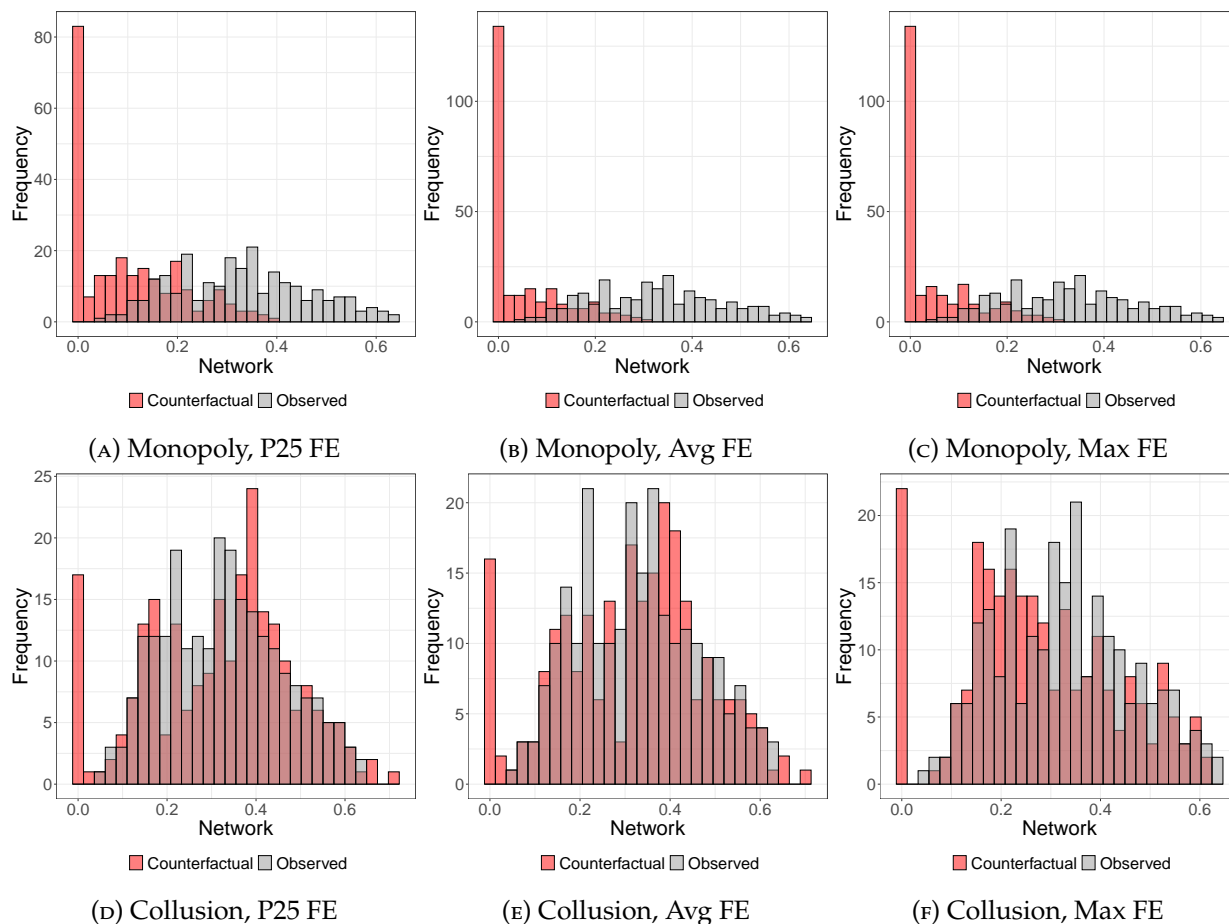
To accommodate the fact that marginal variable profits vary substantially across insurers and some values are relatively large (in the order of billions of COP), I estimate the parameters on the log of MVP_{jkm}. Appendix Table 6 presents the results.

APPENDIX TABLE 6: Model of Insurer Fixed Costs

Variable	Log marginal variable profit	
	coef	se
Service network breadth	29.19	(1.903)
Insurer FE		
Insurer A	7.762	(0.744)
Insurer B	3.067	(0.464)
Insurer C	6.419	(0.643)
Insurer D	5.391	(0.551)
Insurer E	6.955	(0.932)
Insurer F	8.307	(0.989)
Insurer G	9.376	(0.772)
Insurer H	1.441	(1.105)
Insurer I	2.107	(0.423)
Insurer J	4.741	(0.523)
Insurer K	5.479	(0.616)
Insurer L	-1.694	(0.604)
Insurer M	8.122	(0.833)
Constant	-8.799	(1.130)
Market FE	Yes	
First-stage F-stat	249.47	
N	2280	
Unadjusted R^2	0.670	

Note: Table presents 2SLS regression of the log of marginal variable profit on service network breadth. An observation is a combination of insurer, service, and market. The instrument for service network breadth is the average claim probability for each service among healthy consumers. Table reports the F-statistic for the first stage regression. Specification includes market fixed effects and insurer fixed effects. Robust standard errors in parenthesis.

Appendix 3 Additional Counterfactual Results

(A) Monopoly, P25 FE   (B) Monopoly, Avg FE   (C) Monopoly, Max FE

(D) Collusion, P25 FE   (E) Collusion, Avg FE   (F) Collusion, Max FE

APPENDIX FIGURE 3: Distribution of Service Network Breadth in Collusive Agreements

Note: Figure presents the distribution of service network breadth in the insurance monopoly in the first row and the scenario where the bottom 5 insurers collude in the second row. The first column imposes that the merged firm's fixed effects equals that of the firm in the 25h percentile of the distribution, the second column imposes the average firm fixed effect among those in the collusive agreement, and the third column imposes the maximum firm fixed effects. Histograms in gray represent the observed scenario and in red the counterfactual scenario.

APPENDIX TABLE 7: Insurer Participation Constraints in Centralized Equilibrium

Insurer	Profit per enrollee
Insurer A	-0.009
Insurer B	0.105
Insurer C	0.085
Insurer D	-0.007
Insurer E	-0.009
Insurer G	-0.008
Insurer I	0.261
Insurer J	-0.006
Insurer K	0.230
Insurer L	0.002
Insurer M	-0.001
Insurer N	-0.002

Note: Table shows insurers' total profit divided by their total demand in the social planner's scenario.