

Tiered Hospital Networks, Health Care Demand, and Prices*

JOB MARKET PAPER

Elena Prager[†]

December 15, 2015

Abstract

Health insurers are increasingly using plan designs that incentivize consumers to shop for health care based on price. I study the effects of one such plan design on demand and equilibrium prices. Tiered hospital networks group hospitals by price ranking and vary consumers' out-of-pocket prices to reflect the price variation faced by the insurer. Proponents argue that tiered networks reduce health care spending by steering consumers toward lower-priced hospitals, and by giving insurers an additional bargaining lever in price negotiations with hospitals. To evaluate these claims, I estimate a structural model of health care demand and insurer-hospital bargaining over prices in the Massachusetts private health insurance market. The model extends the standard Nash bargaining framework to explicitly account for the multiplicity of possible tier outcomes. I find that the effects of tiered networks on demand alone can lead to moderate reductions in hospital spending of 0.7% to 1.8%. The effects on negotiated hospital prices are substantially larger, with an average price decline of 11% across hospitals. I conclude that insurance plan designs with demand-side incentives can have large health care spending reduction effects.

JEL codes: I11, I13, L11, L13

The latest version is available [here](#).

*I am indebted to my advisor, Bob Town, and to my dissertation committee, Mark Pauly, Katja Seim, and Ashley Swanson for their invaluable guidance and support. I also received helpful suggestions from participants in the Wharton health economics and industrial organization workshops, as well as Nora Becker, Shulamite Chiu, Guy David, Sunita Desai, Brian Finkelstein, Cinthia Konishi, Adam Leive, Dan Polsky, Preethi Rao, and Amanda Starc. Financial support from the Agency for Healthcare Research and Quality (dissertation grant R36-HS024164), the Social Sciences and Humanities Research Council of Canada, the Leonard Davis Institute of Health Economics Pilot Grant program, and the Ackoff Doctoral Student Fellowship is gratefully acknowledged. This research does not represent the official views of any of these funding agencies.

[†]The Wharton School, University of Pennsylvania. Address: Colonial Penn Center 402, 3641 Locust Walk, Philadelphia, PA 19104. Email: elprager@wharton.upenn.edu.

1 Introduction

Unlike in other markets, prices in health care have historically been neither observed nor paid directly by consumers. Instead, traditional health insurance plans charge consumers out-of-pocket prices that are opaque or at most loosely correlated with the differences in total price across health care providers.¹ Consequently, incentives for price competition between providers have been blunted (Gaynor 2006; Enthoven 2014; White et al. 2014). Insurance design innovations that aim to sensitize demand to health care prices, such as value-based insurance, narrow provider networks, high-deductible health plans, reference pricing, and tiered provider networks, attempt to rein in health care spending using market principles (Yong et al. 2010). These plan designs aim to inject price competition into the health care market by incentivizing consumers to select providers at least partially based on price. If successful, such plan designs can be expected to affect not only consumer decisions but also, by extension, equilibrium prices for health care. In this paper, I study both of these effects in the context of tiered provider networks.

Among recent innovations in insurance design, tiered provider networks most directly encourage health care providers to compete over patients. Insurance plans that use tiered provider networks rank providers based on price and place them into mutually exclusive groups, or *tiers*, that determine consumers' out-of-pocket payment for a particular provider. In contrast to traditional insurance plans, tiered networks vary consumers' out-of-pocket prices to reflect the variation in prices paid by insurers to providers. Plans with tiered provider networks began to take hold in the mid-2000s, as insurers sought new mechanisms for bolstering their bargaining power against increasingly consolidated providers (Yegian 2003; Robinson 2003; Sinaiko 2012). Among very large employers, 33% of the highest-enrollment health plans now include a tiered provider network, with 54% of all employers expecting tiered networks to be a very effective or somewhat effective measure for health care cost reduction (KFF 2014, 2015).

This paper evaluates the effects of tiered networks on both the demand side and the supply side of the health care market. Advocates of tiered networks argue that they reduce health care spending through two mechanisms: the direct effect of steering consumers toward lower-priced providers (Sinaiko 2012), and an indirect effect on prices (Fronstin 2003; Robinson 2003). If consumers indeed respond to the incentives in tiered provider networks, then non-preferred tier placement becomes an additional bargaining lever that insurers can use in price negotiations with providers. In evaluating the spending reduction effects of tiered networks and similar demand-side incentives, it is therefore necessary to consider their impacts on negotiated prices between insurers and providers in addition to their direct effects on demand. In this paper, I focus on the tiering of hospitals, whose price negotiations have an outsize importance to their bottom line (Gaynor et al. 2015).²

I evaluate the overall effect of tiered hospital networks by building and estimating a model of

¹Plans that use coinsurance charge consumers a percentage of the total negotiated price, which is perfectly correlated with price but which does not allow consumers to observe the coinsurance amount ex ante due to a lack of price transparency (White et al. 2014).

²At 5.6% of GDP, factors affecting hospital spending are also of independent interest (CMS 2014a).

insurer-hospital competition under tiered networks. The model describes bilateral Nash bargaining between insurers and hospitals over prices. The equilibrium price maximizes the Nash product of the insurer’s and hospital’s surpluses, which are in turn functions of negotiated prices, hospital tiers, insurance plan premiums, plan enrollments, and hospital utilization. Tiered networks introduce an additional set of incentives relative to traditional insurance plans. In agreeing to a lower negotiated price, a hospital trades off lower per-patient revenue against higher volume due to more preferred tier placement. Plan premiums and enrollments also respond to prices and tiers, affecting both insurers’ and hospitals’ volumes. I derive the equilibrium negotiated price for this model, which extends the existing Nash bargaining framework from the literature to account explicitly for the presence of tiered networks.

A credible model of the negotiation process between insurers and hospitals requires estimates of the demand-side response to hospital tiers, prices, and other insurance plan characteristics. I first estimate a discrete choice model of individual demand for hospitals, using inertia in insurance plan choices to address the potential endogeneity between plan choice and out-of-pocket hospital price. Next, I estimate a model of demand for insurance plans at the household level. I use the estimates from the hospital demand model to generate a measure of consumers’ valuation of plans’ hospital networks, measured by willingness-to-pay, which enters the plan demand model alongside detailed data on plan financial characteristics. Finally, I combine the estimates from the hospital and plan demand analyses with my structural model of insurer-hospital bargaining for a subset of hospitals. I solve for the hospitals’ marginal costs of treatment and use the estimates to conduct counterfactual analyses that evaluate the effects of tiered networks relative to non-tiered plans, both on patient sorting across hospitals and on hospitals’ negotiated prices with insurers.

My empirical strategy and identification rely on very detailed data. I estimate the model using comprehensive data on the private health insurance market in Massachusetts. I combine data on health care utilization and health insurance enrollment from the 2009–2012 Massachusetts All-Payer Claims Database (APCD); data on insurance plan characteristics from the Massachusetts Group Insurance Commission (GIC); and novel, hand-collected longitudinal data on Massachusetts insurers’ hospital tiers. I use the longitudinal tiered network data to cleanly identify a price coefficient in hospital demand, which is typically impeded by a lack of data on provider networks and out-of-pocket prices (Gaynor et al. 2015). The GIC data provide information required for the plan demand model, including plan characteristics and plan choice set data that are not observed in medical claims databases such as the APCD. The final key piece of data reported in the APCD is actual transaction prices paid to hospitals, which are critical to measuring spending and to the credibility of the bargaining model but which are typically not available in medical claims data (Reinhardt 2006; Gowrisankaran et al. 2015). The unique combination of longitudinal network data, detailed hospital choice and plan choice data for the same consumers, and accurate price data forms the backbone for the demand- and supply-side empirical analyses in this paper.

I find that both demand and prices are responsive to tiered hospital networks. On the demand side, I find that consumers’ probability of choosing a hospital is decreasing in out-of-pocket price.

The estimated elasticity of demand is in the range of -0.1 to -0.2 , consistent with the literature (Manning et al. 1987; Chandra et al. 2010; Trivedi et al. 2010; Buntin et al. 2011). In counterfactual analyses, I find that the effect of moving a consumer from a non-tiered plan to a plan with a \$500 spread in out-of-pocket price between the most and least preferred tiers is a 0.7% reduction in total hospital spending. Increasing the spread across tiers to \$1,250 results in a 1.8% reduction in hospital spending relative to the baseline of no tiers. These results support the claim that demand-side incentives can lower health care spending. However, the magnitude of the demand steering effect of tiered networks is modest.

The second set of results concerns the effect of tiered hospital networks on negotiated hospital prices. Relative to traditional health insurance plans, plans with tiered networks affect price bargaining by making a hospital’s patient volume a function of its tier. For a subset of hospitals, I repeat the counterfactual exercise comparing prices when the insurer does not use a tiered network to prices when the same insurer has some tiered network plans. In addition to allowing consumers to respond to these changes, this exercise allows negotiated prices, tiers, premiums, and enrollment to adjust. The approximate equilibrium effect of moving from exclusively non-tiered plans to some tiered plans with a \$500 spread is an average decline of 11% in hospital prices. These results suggest that demand-side incentives in health insurance may have large downward effects on prices by aggregating demand responses across individual consumers. Moreover, these price effects constitute a more important spending reduction mechanism than the demand steering effects alone.

This paper is related to several strands of the literature. Conceptually and methodologically, it builds on the growing literature on price bargaining in markets lacking posted prices (Capps et al. 2003; Ho 2009; Crawford and Yurukoglu 2012; Collard-Wexler et al. 2014; Grennan 2013).³ I extend the bargaining framework from this literature to contexts in which the space of possible distinct agreement outcomes is larger than one. Existing Nash bargaining models cannot accommodate the multiple possible outcomes inherent in a tiered hospital network because they allow for only two distinct outcomes of a negotiation, agreement and disagreement.⁴ In a tiered network, the agreement outcome between the insurer and the hospital nests multiple possible tier placements. I incorporate the structure of tiered networks by modeling all possible permutations of tier assignments for the hospital and its close competitors and using insurers’ tier determination functions to assign a probability to each permutation. I also contribute to this literature empirically by allowing equilibrium hospital networks to adjust in the counterfactual exercises.

Substantively, this paper is related to a large literature on health insurance design and its relationship to health care demand. The paper contributes to the literature on the elasticity of

³There is general interest in markets with negotiated prices among economists. Horn and Wolinsky (1988) model price negotiations as a Nash bargaining game, and Collard-Wexler et al. (2014) provide a foundation for its use in bilateral oligopoly settings. Variations of the bilateral Nash bargaining model have been operationalized empirically in the context of health insurers’ hospital networks (Ho 2009; Gowrisankaran et al. 2015; Ho and Lee 2015) and other applications (Crawford and Yurukoglu 2012; Grennan 2013). In addition, many papers that study hospital markets rely on the underlying structure of bargaining models without estimating the models structurally (Town and Vistnes 2001; Sorensen 2003; Capps et al. 2003; Lewis and Pflum 2013; Shepard 2014; Trish and Herring 2015).

⁴In the case of bargaining between insurers and hospitals, the two outcomes correspond to the inclusion or exclusion of the hospital from the insurer’s provider network (Town and Vistnes 2001; Capps et al. 2003; Ho 2006, 2009).

demand for medical care⁵ by cleanly estimating the demand elasticity on the intensive margin of consumer substitution across providers in response to variation in spot prices for care.⁶ The majority of the existing evidence on the elasticity of health care demand, including the landmark estimates from the RAND Health Insurance Experiment, measures elasticities on the extensive margin of whether to purchase any health care. To my knowledge, this paper is the first to cleanly estimate the price elasticity of demand across health care providers in response to differences in the out-of-pocket prices borne directly by consumers.⁷

The paper also extends the literature on provider network design by modeling both the demand and supply sides of a more complex network structure than has previously been considered. Narrow hospital networks, which exclude high-priced hospitals from the network altogether rather than placing them in a non-preferred tier, have been studied extensively (Cutler et al. 2000; Town and Vistnes 2001; Capps et al. 2003; Ho 2009; Gowrisankaran et al. 2015).⁸ This paper is most closely related to Gowrisankaran, Nevo and Town (2015), who are the first to incorporate consumer response to hospital prices into insurer-hospital bargaining; and Ho and Lee (2015), who are the first to jointly estimate hospital demand, plan demand as a function of hospital networks, and insurer-hospital bargaining. The literature on tiered hospital networks is much smaller, and focuses on the demand response to hospitals' categorical tiers (Scanlon et al. 2008; Frank et al. 2015). I build on these papers by evaluating the effect of tiered hospital networks on negotiated prices, and by measuring demand response to changes in out-of-pocket prices, rather than hospital tiers alone. Finally, this paper contributes to the broader policy debate about mechanisms for containing rapidly rising health care costs. As health insurance designs that expose consumers to out-of-pocket price variations in health care become more widespread, understanding consumer response to price differences across health care providers is increasingly important.

The paper proceeds as follows. Section 2 discusses the history and design of tiered provider networks and outlines the pertinent policy background. Section 3 describes the data and empirical setting. Section 4 presents a model of hospital-insurer price bargaining under tiered networks, and Sections 5–6 detail the empirical approach. The results of the estimation and counterfactual analyses are presented in Sections 7–8. Implications of my findings are discussed in Section 9.

⁵The landmark estimates of the elasticity of health care demand provided by the RAND Health Insurance Experiment are in the range of -0.1 to -0.2 ; more recent estimates for various classes of medical care mostly fall in the same range (Manning et al. 1987; Chandra et al. 2010; Trivedi et al. 2010; Buntin et al. 2011).

⁶I abstract from the nonlinearity of marginal price induced by variation in marginal tax rates and nonlinear contracts such as deductibles and out-of-pocket maximums (Gruber and Poterba 1994; Finkelstein 2002; Kowalski 2012; Einav et al. 2013; Abaluck et al. 2015).

⁷Gowrisankaran et al. (2015) and Ho and Pakes (2013) study provider choice under differential pricing, but in these papers, consumers are responding to price via coinsurance or because their choices are mediated by physician referrals. There are also estimates of consumer response to price transparency initiatives, but these are difficult to generalize because the price transparency initiatives usually involve a concerted patient information campaign that is not typical in other contexts (Christensen et al. 2013; Robinson and Brown 2013; Wu et al. 2014).

⁸Several papers have studied the mechanisms through which narrow networks reduce spending using a reduced-form approach (Cutler et al. 2000), including a study by Gruber and McKnight (2014) of the GIC health insurance plan market used in this paper. Papers using a structural approach to study narrow networks have estimated models of insurer-hospital bargaining; these papers typically do not allow for a consumer response to hospital prices.

2 Background

2.1 Tiered Provider Networks

Plans with tiered provider networks were introduced in the early 2000s, as insurers sought new mechanisms for bolstering their bargaining power with respect to increasingly consolidated providers (Yegian 2003; Robinson 2003; Sinaiko 2012). Tiered networks allowed insurers to maintain some of the bargaining leverage associated with health maintenance organizations (HMOs), which used the threat of contract termination to drive down negotiated prices but which experienced a backlash of public opinion in the 1990s (Cutler et al. 2000; Town and Vistnes 2001; Ho 2009). Detractors argued that HMOs’ savings came at the expense of patient choice, access to care, and continuity of care (McCanne 2013; Martin 2014).

Tiered provider networks combine the cost control mechanisms of narrow networks with patient choice and explicit price information for consumers. In a tiered network, almost all providers in the market remain in the consumer’s choice set, but a higher out-of-pocket price is associated with the use of higher-priced providers. Providers are placed into non-overlapping groups, or *tiers*, that determine consumers’ out-of-pocket prices for treatment. The out-of-pocket price faced by enrollees is then constant among providers within a tier, but varies across tiers. Throughout the paper, I distinguish between the out-of-pocket price faced by insured consumers and the full price negotiated between providers and insurers, which I call simply “price”.

The concept of tiering in health care is not new; insurers have been grouping prescription drugs into tiers on their drug formularies since at least the 1990s, and by 2000 the fraction of insurers using tiered formularies reached 80% (Motheral and Fairman 2001). The application of tiering to provider networks did not become widespread until the mid-2000s (Sinaiko 2012). Insurers can tier their hospital networks, their physician networks, or both (Sinaiko 2012). Motivated by the nearly one third share of total health care spending going towards hospital costs, insurers and employers have been particularly interested in tiering as a means for controlling hospital spending (Fronstin 2003; Gaynor et al. 2015). The typical tiered hospital network has three tiers, with most or all hospitals in the market included in one of the three tiers (Fronstin 2003). In my data, out-of-pocket price differentials for a single hospital admission between the most and least preferred tiers range from \$200 to as much as \$1,250.

Since their introduction in the early 2000s, the penetration of tiered-network plan designs has continued to rise. Health care system experts, insurers and employers increasingly see the use of tiered networks and other value-based plan designs as integral to cost control (Robinson 2003; KFF 2014; Stremikis et al. 2010). As of 2015, 33% of the highest-enrollment health plans offered by very large employers and 7% of plans offered on the health insurance exchanges include a tiered provider network, and multiple states expect growth in tiered-network plans (KFF 2014; Corlette et al. 2014; McKinsey 2015; KFF 2015). Moreover, some states have been directly involved in promoting the adoption of tiered provider networks.

The literature on tiered provider networks is small. Scanlon et al. (2008) and Frank et al. (2015)

each examine the demand side of a specific tiered hospital network program, and find evidence that hospital tiers steer patients to preferred tiers with lower out-of-pocket prices. Sinaiko and Rosenthal (2014) study tiered physician networks and find that physician tiers are only effective at steering new patients who do not have existing relationships with their physician. These papers estimate consumer response on the margin of provider tier rather than out-of-pocket price to the consumer. I build on these papers by estimating the demand response to changes in out-of-pocket prices, rather than categorical changes in hospital tiers alone. Moreover, to my knowledge, this paper is the first to study the effect of tiered hospital networks on the negotiated prices themselves.

2.2 The Massachusetts Health Care Market

The empirical application in this paper is the private health insurance market in Massachusetts, which provides an especially appropriate setting for studying tiered hospital networks. Its largest insurers have a substantial fraction of enrollees in plans using tiered networks, which is helpful for both a sufficient sample size and for identifying variation in tier prices over time, across insurers, and across plans within an insurer. Furthermore, since Massachusetts insurers were early adopters of tiered provider networks, the market has had an opportunity to adjust to the presence of these plans and reach an equilibrium to which a structural model can be applied. Combined with the state’s detailed health care data, these features of the market motivate the choice of Massachusetts as the empirical setting for this paper.

In 2006, Massachusetts passed a landmark health care overhaul which aimed to expand health insurance coverage and access to care. The Massachusetts reform subsequently served as the blueprint for the federal Patient Protection and Affordable Care Act (ACA) passed in 2010 (Kolstad and Kowalski 2012). Although the 2006 legislation succeeded in broadening insurance coverage in Massachusetts, policymakers remained concerned about the state’s high overall health care spending. Not only was the state’s per capita health care spending 15% higher than the national average, driven largely by high hospital spending, it had also grown faster than national health care spending since 2002 (DHCFP 2010). Based on recommendations by the Massachusetts Division of Health Care Finance and Policy, the state implemented additional reforms aimed at measuring and reducing health care spending in 2010 and again in 2012 (Massachusetts 2010, 2012a; Wrobel et al. 2014; CHIA 2015b). These reforms included, among other provisions,⁹ the creation of the All-Payer Claims Database used in this paper and requirements for insurers to offer value-based insurance designs (DHCFP 2010).

Since 2011, Massachusetts legislation has required all large insurers to offer at least one narrow- or tiered-network plan in at least one geographic area (Massachusetts 2010). The regulation does not require insurers to offer tiered-network plans; they may instead offer narrow-network plans. However, all three of the state’s largest insurers—Blue Cross Blue Shield of Massachusetts, Harvard

⁹Other notable pieces of the legislation consisted of health care price transparency requirements and the encouragement of vertical integration between providers in the form accountable care organizations (created under the moniker “Alternative Quality Contract” (Song et al. 2012)).

Pilgrim Health Care, and Tufts Health Plan—have offered both tiered- and narrow-network plans since before the regulation went into effect in 2011. These insurers now have 10–35% of their commercial enrollees in tiered-network plans. State regulation also outlines a method for insurers to calculate comparable prices across providers by adjusting for disease and patient mix; insurers are required to report these prices to the state’s Center for Health Information and Analysis (CHIA) and are expected to use them for determining providers’ network status.

Outside of state legislation, the push toward tiered networks in Massachusetts has been led by the Massachusetts Group Insurance Commission (GIC), which administers health insurance and other benefits for state and municipal employees, retirees, and their dependents.¹⁰ The GIC insures some 300,000–350,000 individuals per year throughout my sample period, corresponding to approximately 8% of the total commercially insured population in Massachusetts. The volume of covered lives on the GIC, along with the substantial fraction of the state budget devoted to it, makes the GIC an important and active player in the Massachusetts health insurance landscape (DHCFP 2010; Wrobel et al. 2014). The GIC was among the earliest adopters of tiered provider networks, introducing its first tiered hospital network plan in July 2003 and rolling out tiered physician networks in July 2006 GIC (2008, 2009). For the insurers of interest in this paper, Harvard Pilgrim Health Care and Tufts Health Plan, nearly 100% of tiered provider network plan enrollment comes from the GIC in the early part of the sample period, falling to roughly 90% by 2013 (Boros et al. 2014).

Massachusetts requires insurers operating tiered-network plans to “clearly and conspicuously indicate” consumers’ out-of-pocket prices for each tier (Massachusetts 2012b). Insurers provide this information to enrollees as part of the schedule of benefits documentation for each plan. At the insurer level, they also publish lists of hospitals and their network tiers each year, which can be easily accessed through their websites for the current year. These lists include each hospital’s tier, so consumers do not need to search for multiple providers’ network status in order to comparison-shop. This is in contrast to the difficulty of learning out-of-pocket prices for hospital care in advance in traditional plan types: even savvy consumers who ask for price quotes typically get poor response rates (Bebinger 2014).

3 Data

The data used in this paper are compiled from multiple sources. Data on health care utilization and health insurance enrollment come from the 2009–2012 Massachusetts All-Payer Claims Database (APCD); data on insurance plan characteristics and choice sets are drawn from the employee benefit guides of the Massachusetts Group Insurance Commission (GIC), a large employer group; and longitudinal data on hospitals’ placement in insurers’ tiered and narrow networks were hand-collected from the current and archived network lists of several Massachusetts insurers.

¹⁰This is the same employer group studied by Gruber and McKnight (2014) in evaluating the impact of narrow networks and by Sinaiko and Rosenthal (2014) in studying patient response to physician tiering.

3.1 Medical Claims and Hospital Price Data

Medical claims data are drawn from the Massachusetts Center for Health Information and Analysis’ (CHIA) All-Payer Claims Database (APCD) (CHIA 2014). The APCD consists of comprehensive data on interactions with the health care system of all privately insured residents of Massachusetts in the 2009–2012 period.

The APCD medical claims data are extremely detailed. They include information on physician visits, outpatient hospital visits, inpatient hospital admissions, and prescription drugs. The data include patient demographic information such as gender, date of birth, and five-digit zip codes of residence. I match patients to zip-level demographic characteristics from the U.S. Census Bureau and use the patient address information to calculate driving distance from patients to providers. The APCD allows me to track patients across years, and often across insurers, using longitudinal patient identifiers. In addition, it links patients insured as dependents to the primary enrollee in the insurance plan, allowing household units to be identified when modeling insurance enrollment decisions. To my knowledge, only one other study has estimated individual demand for providers and household demand for health insurance in the same population (Ho and Lee 2015). This link between the two stages of demand is key to an accurate model of the health insurance market, where plan enrollment decisions are often made at the level of the family rather than the individual.

Like other medical claims databases, the unit of observation in the APCD is the claim line, which is the smallest unit of service for which an insurer or patient is billed separately from other units of service. A single hospital visit, for example, can have many claim lines for drugs, operating room supplies, anesthesia, and physician fees. In the analysis, I aggregate information across claim lines to the level of the hospital admission. For each claim, the principal diagnosis is reported along with up to twelve secondary diagnoses. Similarly, for visits involving procedures, a principal procedure code is reported along with up to six secondary procedures. Summary statistics for the admissions included in the hospital demand model are reported in Table 7. Diagnoses and procedures are reported in the International Classification of Diseases, Clinical Modification (ICD-9) classification system, which consists of approximately 14,000 distinct diagnosis codes and 4,000 procedure codes. I assign diagnoses to diagnostic categories and severity levels using the Clinical Classifications Software (CCS) categorizations from the Agency of Healthcare Research and Quality (Table 8). Hospitals are identified in the data using fuzzy matching on hospital names and addresses, plus a final round of manual checks to correct errors and exclude mistakenly attributed onsite facilities or physician groups that are not involved in inpatient care. The APCD is supplemented with hospital characteristics data from the American Hospital Association Annual Survey Database; hospital quality data from the Centers for Medicare and Medicaid Services Hospital Compare database; and hospital financial and casemix data from state public use files published by the Massachusetts Center for Health Information and Analysis.

The APCD reports several key price variables. Most importantly, it reports allowed amounts, which are actual amounts paid by insurers and patients to health care providers. The majority of claims databases only report charge prices, which are not reflective of actual transaction prices

(Reinhardt 2006). The majority of existing empirical work on insurer-hospital strategic interactions has been limited by its inability to measure actual dollar flows from payers and patients to providers (Gowrisankaran et al. 2015). The typical approach for overcoming these data limitations has been to infer the break-down of price negotiations when a hospital is excluded from an insurer’s provider network. By contrast, the price information in the APCD allows price negotiations between insurers and providers to be examined directly, irrespective of variation in network status. In addition to amounts paid by insurers, the APCD separately reports patients’ out-of-pocket payments for care, a key identifying variable in estimating hospital demand in tiered-network plans.

The health care utilization data from the APCD are used to estimate hospital demand in conjunction with the hospital network data described below. The accurate price information from the APCD is used to measure negotiated prices between insurers and hospitals, and to estimate the structural model of insurer-hospital bargaining over prices.

3.2 Premiums and Choice Sets Data

Data on insurance plan availability and characteristics are drawn from the Massachusetts Group Insurance Commission (GIC) for the subset of consumers in the APCD who are insured through the GIC.¹¹ The GIC is the benefits administrator for the state of Massachusetts, some municipalities, and a number of other public entities. It insures some 300,000–350,000 people per year during my sample period, consisting of GIC-covered employees, retirees, and their dependents. This enrollment volume makes the GIC the clearinghouse for 8% of the state’s 4.4 million commercially insured lives (CHIA 2015b). The volume of covered lives on the GIC, along with the substantial fraction of the state budget devoted to it, makes the GIC an important and active player in the Massachusetts health insurance landscape (DHCFP 2010; Wrobel et al. 2014). My sample of GIC enrollees observed in the APCD includes approximately 90,000 state and municipal employees and 120,000 dependents. The remaining individuals insured through the GIC are retired government employees and their surviving spouses. The demographic characteristics for the GIC enrollees in my sample are shown in Table 1. Approximately 60% of primary enrollees insure their dependents as well. The majority of the primary enrollees live in the Boston area or elsewhere in eastern Massachusetts. Approximately half of the enrollees are first observed in the GIC prior to the start of the medical claims data in 2009.

I use data on the GIC’s health plan offerings, premiums, and plan characteristics such as deductibles for GIC fiscal years 2009–2011, which cover the calendar period July 2008–June 2012.¹² The plan offerings and their premiums for a sample enrollment year are described in Table 2. The employee portion of premium contributions is 25% of the total premium.¹³ Two levels of premiums

¹¹I am grateful to GIC Budget Director Catherine Moore for detailed information on the institutional setting and goals of the GIC.

¹²Data from July 2012 onward excluded because the GIC implemented a premium discount program that affected employees differently depending on characteristics I do not observe in the APCD (Gruber and McKnight 2014). The plan demand analysis therefore relies on GIC data through June 2012.

¹³Employees hired prior to July 2003 only pay 20% of the total premium cost. In the analyses, I therefore exclude GIC enrollees who were enrolled prior to 2007 (the earliest enrollment data in the APCD) in order to reduce noise in

Table 1: Characteristics of GIC health insurance enrollees

	Individuals	Families
% of households	39.5	60.5
% of total enrollment	17.8	82.2
Median family size	1	3
Mean family size	1	3.2
% female	59.5	50.3
Mean age	48.1	35.7
Median age	49	39
% entering before 2009	47.3	56.2
% Western Mass.	19.8	18.2
% Central Mass.	12.2	13.1
% Northeast Mass.	28.1	29.4
% Metro Boston	25.4	20
% Southeast Mass.	14.6	19.3

Summary statistics for Massachusetts Group Insurance Commission (GIC) health insurance enrollees. Column 1 is single enrollees; column 2 is enrollees with dependents. 60% of enrolled households include dependents, who are typically younger than primary enrollees. Approximately half of households are enrolled in the GIC prior to the start of the data in 2009.

Table 2: Plans available on the GIC, fiscal year 2011

Plan	Tiered	Network	Copays (\$)	Indiv. Prem. (\$)	Family Prem. (\$)
Fallon Direct		Narrow	200	1,265.16	3,007.44
Fallon Select		Inclusive	250	1,513.44	3,603.48
Harvard Pilgrim Independence	Yes	Inclusive	250/500/750	1,829.52	4,439.16
Harvard Pilgrim Primary Choice	Yes	Narrow	250/500/—	1,456.20	3,527.40
Health New England		Narrow	250	1,262.52	3,099.36
Neighborhood Health Plan		Narrow	250	1,261.08	3,307.92
Tufts Navigator	Yes	Inclusive	300/700/700	1,760.16	4,244.52
Tufts Spirit	Yes	Narrow	300/700/—	1,401.24	3,372.96
UniCare Basic		Inclusive	200	2,765.52	6,424.20
UniCare Community Choice	Yes	Inclusive	250/500/750	1,240.44	2,948.16
UniCare PLUS		Inclusive	250	1,703.52	4,036.92

GIC plans for fiscal year 2011 (July 2010–June 2011). Copays are for hospital inpatient services (across tiers 1/2/3). Premiums are the employee’s annual premium contribution, 25% of the total premium. Each plan has two levels of premiums, one for individual and one for family coverage, that do not vary with age or geography. Plan availability varies across counties and over time.

are set for each plan: one for individual coverage and another for family coverage (defined as two or more enrollees), with no variation in these two premium amounts across the entire state for each fiscal year. Plan characteristics, such as out-of-pocket prices and hospital networks, change over time. Plans on the GIC use copays, which are fixed dollar amounts paid out-of-pocket by consumers when they use health care. For example, inpatient copays in the Harvard Pilgrim Independence plan start at a flat \$300 per admission in fiscal year 2009, move to a tiered structure of \$250/\$500/\$750 across the three hospital tiers in 2010, and increase to \$275/\$500/\$1,500 in 2016.

Six insurers offer a total of eleven plans through the GIC (Table 2). The key insurers of interest in this paper are Harvard Pilgrim Health Care and Tufts Health Plan, although other insurers are also included in the analyses. These two insurers are the second- and third-largest in the state, with 20% and 14% of commercial enrollment, respectively (CHIA 2013).¹⁴ Harvard Pilgrim and Tufts each offer two plans through the GIC, one using a broad tiered hospital network and the other using a narrow version of their tiered network. These narrow-network plans were introduced in July 2010, and are studied extensively in Gruber and McKnight (2014), who also provide a more detailed description of the GIC market. The broad tiered-network plans by Harvard Pilgrim and Tufts have the two highest market shares among employees insured through the GIC, with a combined share ranging from 49% to 59% of employee enrollees throughout the sample period. Of the seven plans offered by other insurers, only one (UniCare) uses a tiered hospital network, and this plan has less than 10% market share on the GIC. UniCare does not contribute data to the APCD, so its enrollees are excluded from the analyses and UniCare plans are assigned to be the outside option.

Plans on the GIC market are fairly standardized: deductible levels, prescription drug copays, and some other plan characteristics vary little or not at all across plans within a fiscal year. This type of standardization is found in many health insurance markets, including Medigap, state health insurance exchanges, and large employers (Starc 2014; Ericson and Starc 2015; Handel 2013). Such markets can shed light on plan competition on the health insurance exchanges set up under the Affordable Care Act, where insurers have responded to the standardization of benefits by competing more aggressively on network design (Davis 2013). The primary differences between plans on the GIC come from the insurer brands, provider networks, and copay structures for physician and hospital care.

The GIC plan data are used to estimate plan demand. The use of the GIC plan data to supplement the APCD claims data provides crucial information on plan characteristics, premiums, and plan choice sets, which are unobserved in the APCD and in claims data more generally. Indeed, very few papers have jointly estimated demand for hospitals and demand for plans, due in part to this common data limitation (Ho 2006; Shepard 2014; Ho and Lee 2015). Without the accurate construction of the plan choice set, discrete choice models of plan demand would be highly unreliable

premium measurement.

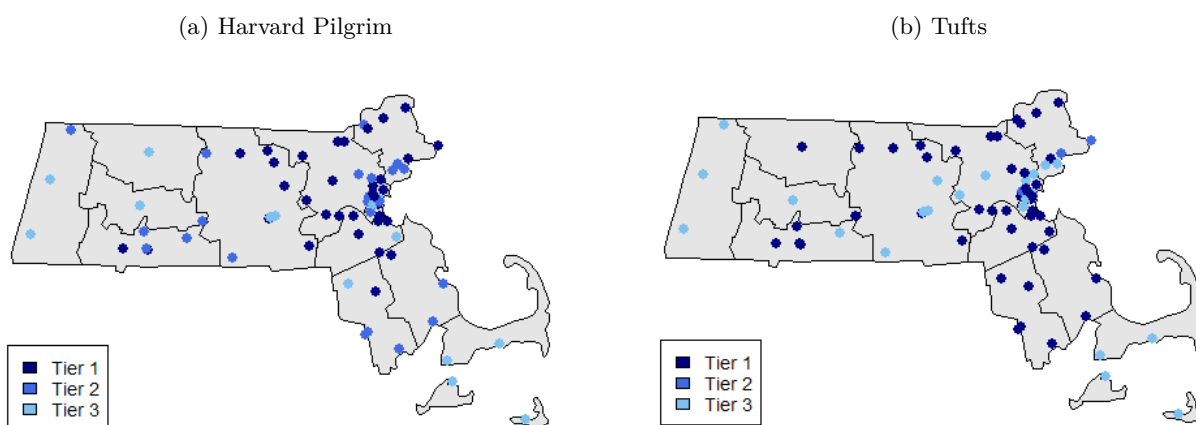
¹⁴The largest insurer in Massachusetts is Blue Cross Blue Shield (BCBS), with 45% of the commercial market (CHIA 2013). BCBS does not participate in the GIC market and is excluded from the analyses. Its tiered hospital network is studied by Frank et al. (2015).

(Train 2002). I observe not only the GIC’s full plan offerings, but also the variation in the subset of the plans available across Massachusetts’ fourteen counties and over time, including the introduction of the two new narrow-network plans in July 2010. While I have plan enrollment data through 2012, I censor the sample for plan demand estimation because in fiscal year 2012, the GIC introduced a premium discount program that affected employees differentially based on employee characteristics that are not observed in the APCD (Gruber and McKnight 2014). These data allow me to estimate a model of plan demand and allow for enrollment adjustments in response to changing hospital networks in the analysis.

3.3 Network Data

I have compiled a unique dataset tracking Massachusetts hospitals’ placements in several insurers’ tiered and narrow networks for the period 2009–2015. Network data were hand-collected from insurers’ current and archived plan documentation.¹⁵ My data cover Harvard Pilgrim’s and Tufts’ tiered networks, as well as all GIC insurers’ narrow networks. Data on insurers’ provider networks have to date been difficult for researchers to obtain, especially retrospective data that can be merged into claims databases, which has limited the scope of questions the literature has been able to address (Gaynor et al. 2015). To my knowledge, this paper is the first to use longitudinal tiered provider network data from multiple insurers, and indeed among the first to use longitudinal data on any type of provider network.

Figure 1: Massachusetts insurers’ hospital tiers (2012)



Maps of Harvard Pilgrim’s and Tufts’ tiered hospital networks in 2012. Each dot represents a general acute care hospital in Massachusetts. Contours represent Massachusetts counties. All hospitals are included in both insurers’ tiered networks, but hospitals’ tiers are not necessarily consistent across insurers.

A map of Harvard Pilgrim’s and Tufts’ network tiers for 2012, the most recent year for which claims data are available, is shown in Figure 1. Figure 7 in the appendix also shows the tiers for

¹⁵For three of the insurers—Health New England, Neighborhood Health Plan, and UniCare—data on narrow networks were supplemented with data collected by the GIC. I thank Cindy McGrath at the GIC for sharing these data for previous years.

Blue Cross Blue Shield, the state’s largest insurer, for comparison. All Massachusetts hospitals are in-network for these tiered-network plans. Table 3 reports the distribution of hospitals across tiers for 2012. The analysis will be restricted to the state’s 61 general acute care hospitals, which have a total of 72 distinct campuses.¹⁶ In the table, tier 1 denotes the most preferred tier with the lowest out-of-pocket price, and tier 3 the least preferred tier. The relative size of the tiers differs across insurers, and hospitals belonging to the same system are not necessarily in the same tier within an insurer. There is merger and acquisition activity within the time period covered by the tiering data, but changes in ownership or system status do not seem to affect tier assignments (almost all the acquired hospitals begin in the most preferred tier). Table 4 reports the distribution of hospital characteristics across tiers. Hospitals in the least preferred tier, tier 3, are disproportionately large. Academic medical centers (AMCs) are more commonly in tier 1 or tier 3 than in the middle tier. A non-negligible fraction of hospitals is found in each tier in both the Boston area and less urban parts of Massachusetts.

The longitudinal nature of the data provides identifying variation for estimating demand response to hospital out-of-pocket price. Hospitals’ tiers vary cross-sectionally across insurers and over time within an insurer. In addition, there is variation from differences in tier copays across plans within an insurer-year. Table 5 shows the contemporaneous variation in a hospital’s tier assignment across Harvard Pilgrim’s and Tufts’ tiered networks. Each cell (i, j) in the table denotes the percentage of hospitals, among those in Harvard Pilgrim’s row i tier, that are in Tufts’ column j tier in the same year. Although some hospitals consistently occupy high or low tiers across insurers, half (49%) of hospitals are in different tiers in Harvard Pilgrim’s and Tufts’ networks. Of those, a fifth (10% of the total) are in tier 1 for one insurer and tier 3 for the other.

Hospitals also change tiers within an insurer’s network over time. By law, tier assignments can change at most annually (Massachusetts 2010). Table 6 shows the transition matrix of hospitals’ tier placements over time within the same insurer’s network. Each cell (i, j) in the table denotes the percentage of hospitals starting in the tier in row i in 2010 that have moved to the tier in column j by 2014. Within an insurer over time, there is movement of hospitals across tiers in both directions; this movement is typically not consistent across insurers. Depending on the tier in the baseline year, 27-36% of hospitals in an insurer’s tiered network are moved to a different tier by the end of the sample period. Of the hospitals whose tier assignment changes, the majority move to an adjacent tier; that is, there is little movement between tiers 1 and 3. Hospitals occasionally move out of and then back into their initial tier during the sample period.

In addition to the variation in a hospital’s tier across insurers and over time, consumers’ out-of-pocket costs for care from a given hospital can vary across plans within an insurer. For example, Harvard Pilgrim Health Care’s family of tiered-network plans includes plans with copays for tiers 1, 2, and 3 of \$250, \$500, and \$750, respectively; and other plans with copays of \$300, \$300, and \$700. In both cases, the identity of hospitals in each tier is unchanged within an insurer-year, but

¹⁶Satellite campuses of hospitals are excluded from these summary statistics, but enter into the demand estimation as separate choice alternatives to account for the fact that their location and available services can differ from the hospital’s primary campus.

Table 3: Distribution of hospitals across tiers, 2012

# of Hospitals in	HPHC	Tufts
Tier 1	28	39
Tier 2	20	2
Tier 3	13	20
Total	61	61

Counts of hospitals in each tier for a sample year. HPHC is Harvard Pilgrim. Satellite campuses are excluded.

Table 4: Hospital characteristics by tier, 2010-2014

	% of All Hospitals	Beds (tier means)	% of System Hospitals	% of AMCs	% of Boston HRR Hospitals	% of Non-Boston HRR Hospitals
Tier 1	51.6	240.9	41.1	32.5	54.7	44.1
Tier 2	23.9	286.7	22.2	30.8	22.5	26.8
Tier 3	24.5	318.2	36.7	36.8	22.8	29.1
Count	61.0	53.0	31.0	14.0	41.0	20.0

Hospital characteristics weighted by tier frequency across insurers and years. Final row reports hospital counts. Hospitals in the least preferred tier (tier 3) are larger and have a higher proportion of academic medical centers (AMCs). Hospitals both in and outside of Boston are present in all three tiers.

Table 5: Hospital contemporaneous tier differences across insurers (%), 2011-2014

HPHC \ Tufts	Tier 1	Tier 2	Tier 3	Total
Tier 1	81.0	5.0	14.0	100.0
Tier 2	67.0	9.6	23.4	100.0
Tier 3	23.1	7.7	69.2	100.0

Percent of hospitals in row insurer's tier that are in column insurer's tier in the same year. Satellite campuses are excluded. Half of hospitals are in different tiers across insurers.

Table 6: Hospital tier changes within insurers (%), 2010-2014

From \ To	Tier 1	Tier 2	Tier 3	Total
Tier 1	68.2	25.8	6.1	100.0
Tier 2	31.8	63.6	4.5	100.0
Tier 3	3.0	24.2	72.7	100.0

Percent of hospitals transitioning from row tier in 2010 to column tier in 2014. Satellite campuses are excluded. One quarter to one third of hospitals in each tier change tiers over time. Most movement is to adjacent tiers.

the associated copay structure can differ across plans. Larger differences in out-of-pocket costs are also observed. Among high-enrollment products, the largest differences are in Tufts Health Plan plans with copays of \$250, \$750, and \$1,500 across hospitals in tiers 1, 2, and 3. The combination of cross-sectional variation in hospital tiers across insurers, variation over time within an insurer, and variation in copays across plans within an insurer-year provides helpful identifying variation for estimating hospital demand.

The tiered network data are used to estimate hospital demand as a function of out-of-pocket price. Clean identification of a price coefficient in hospital demand is typically impeded by a lack of data on insurers' provider network arrangements, especially retrospective data that can be merged into data on medical care usage (Gaynor et al. 2015). I overcome this identification challenge using my longitudinal tiered network data, which allow me to infer consumers' out-of-pocket prices for hospitals from which they are not observed to seek care.

4 Model Overview

I model consumer choice of hospitals, household choice of insurance plans, and price negotiations between insurers and hospitals. The model and empirical approach extend the literature on price negotiations under narrow networks¹⁷ to explicitly account for the multiplicity of possible tier outcomes in a tiered network and the concomitant variation in consumers' out-of-pocket prices for care. In the model, insurers and hospitals bargain bilaterally over prices, which determine hospitals' tiers in insurer networks. The equilibrium prices are a function of demand response to hospital tiers and out-of-pocket prices in hospital choice and plan choice. The model assumes there are no information asymmetries in the market.

The model takes as given the product menu offered by insurers and the mapping from a hospital's price relative to other hospitals in the insurer's network to its tier. Fixing the product characteristics is equivalent to assuming that plan characteristics other than the hospital network and premium are determined separately from hospital prices. In my setting for estimating plan choice, the GIC is intimately involved in setting plan characteristics for all insurers offering GIC plans, leaving little room for insurers to reoptimize their plan characteristics in response to changes in hospital networks. Many plan characteristics, such as deductibles, out-of-pocket payments for prescription drugs, and even the ratio of individual to family premiums, are the same for all six insurers participating on the GIC. More generally, health insurance plans are complex products with high fixed costs of redesigning product characteristics, making the assumption that product characteristics can be held fixed as hospital prices change a reasonable first-order approximation.

The market is modeled according to the following three-stage game.

1. In stage 1, insurers and hospitals engage in simultaneous, bilateral negotiations over prices. The price determines the hospital's tier according to a stochastic mapping from prices to tiers, and insurers set premiums for their plans accordingly.

¹⁷See Town and Vistnes (2001); Capps et al. (2003); Ho (2009); Ho and Lee (2013); Gowrisankaran et al. (2015).

2. In stage 2, households choose a health insurance plan given the expected utility from each plan’s network.
3. In stage 3, individual consumers get sick with some probability and, if sick, they choose a hospital given their plan’s network.

Stage 1 corresponds to the bargaining model, while stages 2 and 3 correspond to demand estimation. The bargaining component takes the demand component as an input, since the effect of tiered networks on bargaining will depend on the leverage insurers gain from demand response to a hospital’s tier placement. I therefore build up the model and estimation strategy starting from the last stage of the game.

5 Demand Estimation

5.1 Demand for Hospitals

The market consists of consumers (patients) i , who belong to households $\iota \in I$ with one or more members; insurers $\mathcal{M} \in M$; insurance plans $m \in \mathcal{M}$; and hospitals $h \in H$. In the last stage of the model, consumers’ health risk is realized: consumer i enrolled in plan m becomes sick with diagnosis $d \in D$ with probability f_{id} , which is allowed to vary according to consumer characteristics. (A list of the symbols used throughout the paper is provided in the appendix on page 60.) The consumer must then choose a hospital for treatment to maximize her utility, which depends on the consumer’s characteristics, the hospital’s characteristics, and the out-of-pocket price for the hospital in the consumer’s health plan. Conditional on being sufficiently ill to require inpatient hospital treatment, patient i ’s utility from seeking treatment at hospital h is given by

$$u_{mhid} = -\alpha c_{mh} + \beta x_{hid} + \varepsilon_{mhid} \quad (1)$$

where c_{mh} is the copay for treatment at hospital h under plan m ; α is out-of-pocket price sensitivity; x_{hid} is a vector of patient, illness, and hospital characteristics and their interactions, including hospital fixed effects; β is the associated coefficient vector; and ε_{mhid} is an idiosyncratic error term that is i.i.d. type 1 extreme value. The key parameter of interest is demand sensitivity to out-of-pocket price α .

Hospital demand is estimated on approximately 30,000 inpatient hospital admissions of nonelderly, privately insured patients in Massachusetts between 2009 and 2012. These include all observed admissions of GIC enrollees in four tiered and five non-tiered GIC plans and an additional 8,000 admissions of patients in Harvard Pilgrim’s and Tufts’ tiered plans offered outside the GIC. The non-GIC enrollees are included to provide additional variation in hospital tier copays across plans. I exclude claims for admissions originating from the emergency department (ED) or via transfers from other hospitals, for two reasons. First, the notion of patients’ hospital choice is of questionable validity for such hospitalizations. Second, Massachusetts legislation requires care originating in the

ED to be covered at the lowest patient out-of-pocket price regardless of provider tier (Massachusetts 2010).

Patient and hospital characteristics in x_{hid} include patient demographics, diagnosis category, hospital characteristics, and distance. Distance to a hospital has been found to be an important determinant of hospital choice (Kessler and McClellan 2000; Town and Vistnes 2001; Capps et al. 2003). The demand model uses driving distance from the centroid of the patient’s zip code to the hospital’s street address and the square of the distance, calculated using Bing Maps driving directions. Patient demographics such as age and gender are also included. Hospital characteristics include teaching status, number of beds, an indicator for whether the hospital is a secondary satellite campus, and hospital quality. Hospital quality is measured as perceived by patients using measures from the Hospital Consumer Assessment of Healthcare Providers and Systems (HCAHPS).¹⁸ In contrast to previous work on hospital choice, the inclusion of quality measures allows less of the heterogeneity in hospital preferences to be loaded onto hospital fixed effects. Summary statistics for the admissions included in the hospital demand model are shown in Table 7.

Diagnoses are grouped into diagnostic categories using the Clinical Classifications Software (CCS) developed by the Agency for Healthcare Research and Quality (AHRQ 2015). The CCS classification system assigns diagnosis codes to approximately 300 mutually exclusive diagnosis groups, which are further aggregated into eighteen broad diagnostic categories. The CCS diagnostic categories are described and their prevalence in the Massachusetts nonelderly population given in Table 8. The model includes interaction terms for CCS categories and indicators for the presence of related services at each hospital, drawn from the American Hospital Association Annual Survey of Hospitals. In particular, the demand estimation includes: cardiac CCS interacted with hospital catheterization lab; obstetric CCS interacted with neonatal intensive care unit; nervous, circulatory, and musculoskeletal CCS interacted with MRI; and nervous system CCS interacted with neurological services. These interactions allow hospital choice to vary according to whether specialized services relating to the patient’s diagnosis are available at the hospital.

This parameterization of hospital choice has several implications. The multinomial logit structure implies the independence of irrelevant alternatives (IIA) property of demand, which I mitigate by including detailed data at the consumer-hospital level, such as driving distance and interactions between diagnosis and hospital facilities. The model also treats choice of hospital as a composite measure of the patient’s preferences and other factors. In general, a patient’s choice of hospital can be mediated by unobserved factors, notably referrals by the patient’s physician (Kolstad and Chernew 2009; Ho and Pakes 2013). In this paper, the goal is to estimate the effect of tiered networks on actual market outcomes that may include physician referrals, so I treat the observed choice of hospital as the quantity of interest irrespective of the physician’s influence on the decision. If hospital choices are subject to unobserved influences not related to price, this will bias my

¹⁸The HCAHPS is a third-party national survey of patients that asks about their hospital experience, including responsiveness of medical staff, cleanliness, pain control, and overall rating (CMS 2014b). The HCAHPS scores capture patients’ perceptions of hospital quality and are highly correlated with other hospital reputation measures such as U.S. News rankings.

Table 7: Inpatient admissions for hospital demand model

Mean age	41.6	—	—
% female	64.1	—	—
% chronic	34.6	—	—
% in tiered plans	65.6	—	—
	Non-tiered	Tier 1	Tier 2, 3
% of admits	34.4	31.2	68.8
Mean distance	15.1	11.5	15.9
Mean copay (\$)	240.2	268	614.8

Summary statistics for admissions used to estimate the hospital demand model. Two-thirds of admissions are from enrollees in tiered plans. First column of second panel reports non-tiered plans' share of admissions and characteristics. Columns 2 and 3 report tiered plan admissions. Patients travel farther to hospitals in higher-copay tiers.

Table 8: Descriptions and prevalence of CCS diagnostic categories

Code	Description	Share
1	Infectious and parasitic diseases	1.9
2	Neoplasms	4.9
3	Endocrine; nutritional; and metabolic diseases and immunity disorders	3.9
4	Diseases of the blood and blood-forming organs	0.9
5	Mental illness	9.8
6	Diseases of the nervous system and sense organs	2.7
7	Diseases of the circulatory system	10.2
8	Diseases of the respiratory system	7.5
9	Diseases of the digestive system	10.0
10	Diseases of the genitourinary system	3.9
11	Complications of pregnancy; childbirth; and the puerperium	13.5
12	Diseases of the skin and subcutaneous tissue	2.1
13	Diseases of the musculoskeletal system and connective tissue	5.4
14	Congenital anomalies	0.5
15	Certain conditions originating in the perinatal period	13.1
16	Injury and poisoning	7.1
17	Symptoms; signs; and ill-defined conditions	2.1
18	Residual codes; unclassified; all E codes	0.3

Clinical Classifications Software (CCS) diagnostic categories. First column is Level 1 code (the broadest level), second column is description, third column is % share of nonelderly hospital discharges in Massachusetts.

estimate of out-of-pocket price sensitivity toward the null.

Conditional on a diagnosis and insurance plan structure, consumers choose a hospital to maximize utility as a function of all the choice variables just described. Because the error ε_{mhid} is assumed i.i.d. type 1 extreme value, the consumer’s probability σ_{mhid} of choosing hospital h under plan m and diagnosis d then becomes

$$\sigma_{mhid} = \frac{\exp(-\alpha c_{mh} + \beta x_{hid})}{\sum_{h' \in H} \exp(-\alpha c_{mh'} + \beta x_{h'id})}. \quad (2)$$

In this set-up, patients will value more highly those network arrangements that set low out-of-pocket prices c_{mh} for nearby, high-quality, or otherwise desirable hospitals.

Consumer valuation of a hospital network is measured by willingness-to-pay (WTP). An individual consumer’s ex ante dollarized valuation of plan m ’s tiered hospital network is the expected utility of seeking care at various hospitals at the out-of-pocket prices dictated by the tiers in m :

$$W_{mi} = \frac{1}{\alpha} \sum_{d \in D} f_{id} \ln \left(\sum_{h \in H} \exp(-\alpha c_{mh} + \beta x_{hid}) \right). \quad (3)$$

This expression is the familiar log-sum equation for expected consumer surplus for a logit model, modified in that an additional expectation is taken over the probability of consuming any care, expressed in f_{id} . This modification gives rise to the willingness-to-pay for a hospital network as defined in Capps et al. (2003), here with the additional complication that networks can vary in out-of-pocket prices across hospitals. The availability of a direct estimate of the price responsiveness parameter α allows the WTP to be expressed in dollars, rather than in utils as is the case in settings that lack out-of-pocket price variation. The identification of W_{mi} within versus across consumers is discussed in the appendix (page 62).

The calculation of WTP requires each consumer’s ex ante distribution of diagnosis probabilities f_{id} for the upcoming year. I calculate these probabilities separately for each sex–10-year age band cell and each CCS diagnostic category using data on all non-transfer hospital admissions of Massachusetts residents from the 2010 HCUP State Inpatient Database.¹⁹ Since patient covariates such as distance to hospitals also vary across zip codes, WTP for a given hospital network takes on a separate value for each gender-age group–zip code triplet. Allowing for this granular variation in consumers’ preferences and admission probabilities at the diagnostic category level allows the WTP measure to capture rich variation across consumers.

¹⁹This is equivalent to the assumption that that a consumer’s expectation of her health status for the upcoming year is a consistent predictor of her health status, given only her sex, her 10-year age group, and the fact of residing in Massachusetts. This assumption is more likely to hold for relatively healthy consumers who do not have highly informative personal experience to inform their ex ante expectations of diagnosis (Shepard 2014). Since my data consist of non-elderly, commercially insured, mostly employed individuals, they are healthier than the general population and good candidates for the assumption that their expected health status is approximately equal to the average health status for their age group. To the extent that there are deviations from the average health status, they will load onto the error term in the plan choice model.

5.2 Identification of Hospital Demand

Identification of preference parameters in the hospital choice model relies on cross-sectional and longitudinal variation in hospital networks in addition to differences in hospital and patient characteristics. The model includes hospital fixed effects, so identification comes from within-hospital variation across plans, patients, and years. For example, distance to the hospital and interactions of diagnosis with hospital characteristics vary across admissions by patient address and clinical characteristics. Hospital choice sets vary across plans, with some providing their enrollees access to all hospitals in the state and others using narrow networks.

Identifying variation for the coefficient of interest on out-of-pocket price comes from three sources that leverage the tiered hospital networks in the data. Hospitals move across tiers within insurers' networks over the course of the sample period (Table 6), generating a change of \$200 to \$1,500 in out-of-pocket price depending on the plan. Hospitals with higher negotiated prices are generally in less preferred tiers with higher out-of-pocket price. Table 9 reports mean negotiated prices for the hospitals in each tier, as a multiple of the mean price for hospitals in the most preferred tier (tier 1),²⁰ and the copays for those tiers in the respective insurers' largest tiered plans. In addition, within a year, there is substantial variation in out-of-pocket price arrangements across plans in the sample: copays for hospitals in tiers 1/2/3 range from \$200/\$400/\$400 to \$250/\$750/\$1,500. Finally, a hospital's tier assignment can vary contemporaneously across insurers (Table 5), which provides an important source of within-year identifying variation.

Table 9: Mean hospital prices by tier

Insurer	Price (x)			Copoly (\$)		
	Tier 1	Tier 2	Tier 3	Tier 1	Tier 2	Tier 3
HPHC	1	1.29	1.89	250	500	750
Tufts	1	1.12	1.23	300	700	700

Mean within-tier hospital price as a multiple of insurer's mean tier 1 price and out-of-pocket copays in the insurer's largest GIC plan (2011). Higher-priced hospitals are in less preferred tiers. HPHC is Harvard Pilgrim.

Hospital copays may be endogenous to hospital choice if consumers select into plans based on the network status of their preferred hospitals. If consumers are taking their preferences over hospitals into account when choosing a health insurance plan, then the copay arrangements of the plans into which they sort will not be exogenous. Indeed, empirical evidence suggests that plan choice and subsequent choice of provider can be correlated, at least among consumers with established relationships with their health care providers (Shepard 2014). In my setting, for example, a consumer who places high value on receiving treatment at Massachusetts General Hospital (MGH) for unobservable reasons such as a strong preference for academic medical centers may also choose plans that include MGH in the network at a low out-of-pocket price. The copays faced by consumers in

²⁰Privacy considerations in the data use agreement preclude the reporting of dollar amounts for negotiated prices.

the hospital demand stage, c_{mh} , may therefore be correlated with the error term ε_{mhid} , leading to a biased estimate of the price sensitivity coefficient. The primary concern is that consumers sort into plans based on which networks include their preferred hospitals at the lowest out-of-pocket price. If this is the case, then the estimate of consumers’ price sensitivity will be biased away from the null, in the direction of a more negative coefficient than the true price sensitivity.

To address the potential endogeneity from correlated plan and hospital choices, I leverage consumers’ high level of inertia in plan choices. Intuitively, the identification strategy uses enrollees’ past plan choices to deal with endogeneity in current plan characteristics. The identifying assumption is that conditional on current plan copays c_{mh} and preferences over hospitals captured by x_{hid} , consumers do not anticipate a plan’s future changes to the network or copay structure in period $t + 1$ when choosing a plan for the current enrollment period t . When consumers first enroll in insurance through the GIC, they are in an active-choice setting and may consider their valuation of each plan’s hospital network and other plan characteristics in choosing a plan. In subsequent enrollment periods, although premiums and plan characteristics change, many consumers remain in the same plan without conducting a full reevaluation of their choice sets. Over time, therefore, an inertial consumer’s plan characteristics increasingly approximate random assignment. I leverage this inertia by using the previous year’s copay in the consumer’s plan to deal with the endogeneity in that plan’s current copay.²¹ The identifying assumption would be violated if, for example, consumers are aware that the insurer intends to raise copays in the next enrollment year at the time that they purchase this year’s coverage, a year in advance of the publication of next year’s plan benefits descriptions. An analogous approach is employed by Abaluck et al. (2015) in the context of pharmaceutical coverage choice among seniors.

The use of previous plan choices to generate plausibly exogenous variation in current plan copays is only justified if there is, indeed, a high degree of inertia in plan choice. The fraction of consumers enrolled in each GIC plan in enrollment year 2010 who continued to be enrolled in the same plan in 2011 is reported in Table 10.²² Despite the introduction of two new plans in 2011, 92% of 2010 enrollees remain in the same plan in 2011. In another stark example, in enrollment year 2010, the Harvard Pilgrim Independence Plan switched from a standard hospital network with flat \$300 copays to a tiered hospital network for the first time. The new tiered network uses three hospital tiers with copays of \$250, \$500, and \$750 (Table 2). In spite of this substantial network change, at least 90% of those enrolled in the Harvard Pilgrim plan in enrollment year 2009, immediately prior to the introduction of hospital tiering, were still enrolled in the plan in 2010. These patterns are consistent with findings from the plan choice literature showing that consumers fail to re-optimize their plan choices over time, even as plan characteristics change (Handel 2013; Ericson 2014; Shepard 2014). Combined with these findings, the very high degree of inertia in this dataset motivates the identification strategy.

²¹I use the previous year’s copay rather than the copay in the consumer’s first year of enrollment in order to avoid losing a large portion of the sample when consumers have been enrolled since before the start of the data.

²²The GIC’s enrollment periods coincide with its fiscal years, which begin on July 1 and end on June 30 of the following calendar year.

Table 10: Plan enrollment inertia on GIC, fiscal years 2010–2011

Plan	2010 Enrolt.	2011 Enrolt.	% Inertial
Fallon Direct	3,034	3,913	88.40
Fallon Select	8,109	10,019	91.92
Harvard Pilgrim Independence	70,131	73,486	92.61
Health New England	20,779	21,482	87.43
Neighborhood Health Plan	2,759	3,616	93.33
Tufts Navigator	82,747	85,292	93.39
Mean across plans (weighted)			92.29

% of GIC enrollees remaining in their plans. Two new plans were introduced in 2011 (not shown). Plan enrollments are highly inertial even following a shock to the choice set. This inertia helps to identify the hospital demand model.

Since hospital choice is not linear in the endogenous variable (copay), the standard IV approach of substituting predicted values of the endogenous regressor into the second-stage equation would produce biased estimates (Terza et al. 2008). Instead, I employ a control function approach to deal with the endogeneity. The control function corrects for the correlation between copays c_{mh} and the error term ε_{mhid} by approximating the component of the error that is correlated with copays and including it as a separate regressor (Petrin and Train 2010). In practice, the endogenous variable is regressed on the exogenous variables and the “instrument”, and the residuals from this first-stage regression enter into the nonlinear second-stage model. This approach requires an exclusion restriction analogous to standard IV methods, namely, that the “instrument” affects hospital choice only through its effect on copay. If the assumptions are satisfied, then there exists some function of the first-stage residuals that produces a consistent estimate of the coefficient on the endogenous variable (Wooldridge 2010). Because the true functional form required for consistency is unknown, I allow the first-stage residuals to enter flexibly into the hospital choice model using up to a fifth-degree polynomial expansion.²³ The high degree of plan choice inertia in the data, along with the use of a control function, allow me to obtain a consistent estimate of price sensitivity in a nonlinear and potentially endogenous hospital choice setting.

5.3 Demand for Plans

In stage 2, households choose a health insurance plan given the household members’ expected hospital choices in stage 3. Plan choice is estimated at the level of the household, where each household ι includes the primary plan policy holder and may additionally include other household members. Premiums and willingness to pay for the hospital network are calculated therefore for all spouses and dependents as well as the household’s primary enrollee. The choice between purchasing

²³Some papers have used two-stage residual inclusion (2SRI), where the residuals are entered into the second stage linearly (see, for example, Terza et al. (2008)). However, the consistency result for control functions does not generally hold without a flexible specification for the residuals in the second stage (Wooldridge 2010).

individual insurance or family insurance is taken as given. To my knowledge, only Ho and Lee (2015) have jointly estimated provider choice and plan choice with the household, rather than the individual, as the unit of observation in the plan demand model.

Plan choice depends on the household’s total WTP for the hospital network $W_{m\iota}$, insurance premium $r_{m\iota}$, and other plan attributes X_m . Each household must choose a health insurance plan before household members’ health risk for the entire enrollment period is realized. Therefore, at the time of enrollment, the household projects the expected utility from each plan’s network to choose a utility-maximizing plan. Household ι ’s expected utility $U_{m\iota}$ from enrolling in plan m is given by

$$U_{m\iota} = -\delta_1 r_{m\iota} + \delta_2 W_{m\iota} + \gamma X_m + \zeta_{m\iota} \quad (4)$$

where $W_{m\iota} = \sum_{i \in \iota} W_{mi}$ is the household’s dollarized expected utility from using the hospitals in plan m ; $r_{m\iota}$ is the premium for plan m (which is a function of the family size of household ι); δ_1, δ_2 are premium and WTP sensitivities, respectively; X_m is a vector of the plan’s non-hospital care attributes and plan fixed effects, and γ the associated coefficient; and $\zeta_{m\iota}$ is an i.i.d. type 1 extreme value error term. For households with more than one member, W_{mi} is added up across all household members i and the plan premium $r_{m\iota}$ reflects family coverage premium levels.²⁴

The introduction of the plan’s non-hospital-related characteristics X_m is a recent addition to the literature on insurer-hospital bargaining and plan choice.²⁵ Detailed plan characteristics such as deductible levels and out-of-pocket prices are typically not observable in claims datasets, and are not included in the Massachusetts APCD. However, in the GIC plan benefits documentation, I observe longitudinal information on plan characteristics. The plan demand model includes an indicator whether the plan uses a tiered or narrow physician network; and copay levels for primary care visits, specialist physician visits, and mental health care. Deductibles are observed but excluded from the plan choice model because they do not vary across plans within an enrollment period. To my knowledge, this paper is the first to estimate plan demand allowing a dollarized measure network valuation to enter into plan utility, and using detailed plan financial characteristics. Accounting for both types of arguments in plan demand allows a comparison of consumers’ relative valuation of various plan characteristics in choosing plans.

Given parameterizations of $u_{mh\iota}$ and $U_{m\iota}$, households choose a plan $m \in M$ to maximize $U_{m\iota}$ before household members’ health risk is realized. The type 1 extreme value distribution of $\zeta_{m\iota}$

²⁴A list of the symbols used throughout the paper is provided in the appendix (page 60).

²⁵Ericson and Starc (2014) include plans’ actuarial values in their plan choice models, but earlier papers typically do not have access to claims data and so cannot explicitly account for a plan’s generosity. For many of these papers, this data limitation was not central to the analysis, since they do not rely on estimates of plan choice models for their results. An exception is Ho (2009), who estimates aggregate plan shares and includes plan quality measures, which capture a dimension similar to plan generosity. For the purposes of this analysis, it is important to account for the non-hospital attributes of a plan, since plan choice is being estimated directly. Using only the value of the hospital network $W_{m\iota}$ less the premium would load all other differences across plans into the error term $\zeta_{m\iota}$, which would likely lead to endogeneity.

yields the familiar probability of choosing plan m

$$s_{m\iota} = \frac{\exp(-\delta_1 r_{m\iota} + \delta_2 W_{m\iota} + \gamma X_m)}{\sum_{m'=1}^M \exp(-\delta_1 r_{m'\iota} + \delta_2 W_{m'\iota} + \gamma X_{m'})}.$$

I use the subset of consumers who purchase their coverage through the Group Insurance Commission (GIC) to estimate the plan demand model in order to construct well-specified choice sets. Reconstructing plan choice sets directly from the claims data is unreliable at best, as they do not include employer or group identifiers, and lack information on premiums and plan characteristics. One of the six insurers participating on the GIC is missing from the claims data: UniCare, which has a 32% market share on the GIC across three plans. I therefore estimate plan demand on those GIC enrollees who are enrolled in a plan offered by one of the other five insurers in the data. For these enrollees, the outside option is assumed to be one of UniCare’s three GIC plans. Although these plan records are missing from the claims data, I can include their counterfactual utility, because I observe their premiums and total enrollment in GIC publications. Enrollment in the GIC plans is summarized in Table 19 in the appendix. The assumption required in order to generalize plan demand estimates from the GIC market to the rest of the Massachusetts commercial market is that conditional on the choice set, employment in a state or municipal agency is orthogonal to plan preferences. In the counterfactual exercise, I remain agnostic about this assumption by holding non-GIC plan enrollments fixed as negotiated prices change.

Demand for plans is estimated on a subset of households, using only those for which plan choice sets and characteristics are observed because they purchase health insurance through the GIC. To deal with the high inertia in plan choices, I estimate plan choice for the subset of GIC enrollees who make an active choice, defined as those who are observed in the GIC market for the first time. The longitudinal nature of the enrollment data is key to this approach. In the counterfactual analyses, this restriction of the sample to first-time enrollees is equivalent to conducting medium-run counterfactuals where enrollment is allowed to fully adjust to the sizable shocks explored in the counterfactual scenarios. Table 19 in the appendix shows the number of first-time GIC policy holders (primary enrollees) and their dependents enrolling in each plan over the relevant sample period. There are approximately 36,000 new primary enrollees making an active choice of health plan prior to July 2011. In July 2011, the GIC implemented an incentive program to encourage employees to enroll in narrow-network plans. Employees were eligible for a premium discount if they enrolled in a narrow-network plan, with the size of the discount determined by employee type and tenure (Gruber and McKnight 2014). Since I do not observe employee types in the claims data, I further restrict the sample to plan choices made before the premium discount in order to eliminate measurement error in premiums. Due to the dimensionality of the data required to compute network WTP for each household,²⁶ the model is estimated on a 3% sample of households enrolling in a GIC plan for the first time in fiscal years 2009–2011, for a total of 1,217 households.

²⁶Each individual-diagnosis-hospital-plan combination requires a separate observation.

5.4 Identification of Plan Demand

Identifying variation for the plan demand model comes from several sources. The menu of plans available to consumers on the GIC varies across the state’s fourteen counties as well as over time, including the entry of two new plans in July 2010. Plan characteristics, such as copay levels and hospital networks, also change over time. Premiums are set for the entire state for each fiscal year. These institutional features of the GIC market provide plausibly exogenous variation in plan premiums.

Premiums vary nonlinearly with respect to family size: there is a discontinuous jump in moving from individual coverage to coverage with a dependent, but premiums do not increase further as family size gets larger than two. Figure 2 shows an example of the nonlinear variation in total family premium and per-person premium as family size increases. This type of variation is analogous to the discontinuous increase in premiums as a function of age leveraged by Ericson and Starc (2015). Moreover, the ratio of the family premium to the individual premium is plausibly exogenous, since a uniform ratio is set by the GIC’s actuaries for all GIC plans. Family premiums are approximately 2.4 times the same plan’s individual premium throughout my sample period. Similarly, premium differences between an insurer’s full-network plan and the same insurer’s narrow-network plan, when both are offered on the GIC, are 1.25 across insurers and years. These ratios are chosen by the GIC to maintain a consistent difference between family and individual coverage across plans in the interest of fairness to employees, and to encourage employees to enroll in narrow-network plans, respectively. In other words, they are not designed to accurately reflect cost differences across plan populations and household sizes.


The household-level WTP measure varies by family size and the age, sex, and zip code of residence of the household members (appendix Table 22). Unlike premiums, WTP increases linearly in family size conditional on household member characteristics. This difference in the margins over which premiums and WTP vary helps to identify the model. Older household members contribute larger WTPs due to their higher probabilities of hospital admission.²⁷ For two individuals of the same sex and age but residing in different parts of Massachusetts, variation in WTP comes from the fact that some consumers are geographically closer to a larger number of hospitals or more desirable hospitals (Figure 6a).

6 Bargaining Model

In stage 1 of the game outlined in Section 4, insurers and hospitals engage in simultaneous bilateral bargaining over prices. The bargaining model is solved to infer hospital marginal costs per patient, which are unobserved in the data, and used to simulate market responses to counterfactual policy scenarios. Both insurers and hospitals are assumed to be profit-maximizers and have full information about the rest of the market. Each insurer-hospital pair engages in simultaneous, bilateral

²⁷An exception to the overall monotonically increasing relationship between age and WTP is women of prime childbearing age, who are admitted at higher rates than women younger than 20 and older than 40.

Figure 2: Premiums for Harvard Pilgrim Independence plan in 2011, by family size



Total and per-person employee premium contributions for a sample GIC plan. Premiums per person are nonlinear in family size as a result of only two premium levels for each plan: individual coverage and family coverage. The margin of variation in premiums across households is therefore different from the variation in WTP across households (WTP is linear in household size, conditional on demographics). Nearly two thirds of households purchase family coverage; the figure shows the distribution of enrolled family sizes among those households.

negotiations over price. The equilibrium prices are those that maximize the Nash product of the insurer's and hospital's surplus.

A hospital's network status impacts the surpluses through several channels. First, utilization of the hospital by the insurer's enrollees will change depending on the hospital's network status, with the highest volume when it is in the most preferred to tier and a volume of zero when it is out of network. Second, the distribution of enrollment across plans will depend on consumers' valuation of the insurer's hospital network, which is higher when hospitals are included in more preferred tiers at lower out-of-pocket prices. When a hospital is dropped from the network or moved to a less preferred tier, some consumers' valuation of the network may fall enough to cause them to enroll in a different plan for whose network they have higher WTP. Finally, enrollment is also a function of premiums, which will change in response to the plan's costs as a function of hospital utilization and the share of the total negotiated price borne by patients out-of-pocket. To account for the latter two effects on hospital recapture of an insurer's patients through the patients' enrollment in other plans, I rely on my estimates of plan demand.²⁸

²⁸Few papers in the literature account for plan re-sorting in insurer-provider bargaining. The notable exceptions

The model explicitly accounts for the multiplicity of possible tier outcomes in a tiered network. While standard Nash bargaining models allow for only two distinct outcomes of a negotiation—agreement and disagreement—the agreement outcome between a hospital and an insurer using a tiered network nests multiple possible tier placements. To accommodate this feature of the bargaining game, I leverage the institutional features of the market to construct an approach to collapsing the multiplicity of outcomes to a single summary measure, as described in Section 6.1.

The disagreement outcome of the negotiations is taken to be termination of business between the two parties; that is, the hospital is excluded from the insurer’s network. This is the standard assumption in the existing literature studying narrow-network plans (Town and Vistnes 2001; Capps et al. 2003; Ho 2009; Ho and Lee 2013). Discussions with contracting managers of Massachusetts insurers and hospitals indicate that both sides are acutely aware of the volume and reputational repercussions of failing to reach an agreement. One contracting manager for a large health system colorfully describes failure to renew a contract as “the nuclear option” that no one wants to test. Contract termination does not appear to be an equilibrium outcome for the Massachusetts market.²⁹

The remainder of this section describes the structure of the bargaining model and the estimation approach. Due to the large dimensionality of the bargaining model with tiered networks,³⁰ I estimate the bargaining model for just one insurer, Harvard Pilgrim, and for a small subset of hospitals constituting a relatively self-contained submarket. Harvard Pilgrim is the largest insurer in my data and the second-largest overall in Massachusetts. Other insurers’ negotiated prices and product characteristics are held fixed at their observed equilibrium values, so competition across insurers is captured primarily through competition in the plan demand stage.

6.1 Mapping Prices to Tiers

I model each insurer-hospital pair’s negotiated price as a single base price that applies to all of the insurer’s enrollees³¹ and is scaled by the production resource intensity for each patient’s diagnosis. Each insurer-hospital pair negotiates a price schedule, which is a vector of prices for various treatments, and is collapsed to a base price according to a formula set by the state of Massachusetts. This base price is a casemix-deflated average price paid to the hospital for treating the insurer’s patients. The casemix adjustment converts all hospitals’ prices to a comparable scale by accounting for cross-hospital variation in the complexity of diagnoses and treatments for each hospital’s patient population.

are Ho (2009) and Ho and Lee (2015).

²⁹In one near counter-example in November 2011, BCBS and Tufts Medical Center, along with its affiliated physicians, announced that their contract would lapse due to unsuccessful negotiations (Weisman and Kowalczyk 2011). With both parties facing media scrutiny and intense pressure from patients, a new agreement was reached by December of the same year (Weisman 2011). Under certain market conditions, theory predicts that profit-maximizing insurers will include every hospital in their networks (Capps et al. 2003).

³⁰The computational challenges associated with the model’s dimensionality are detailed in the section on dimensionality reduction.

³¹Insurers typically have a single price schedule for all commercially insured patients and separate schedules for Medicare and Medicaid patients. Since this paper focuses solely on the commercial market, only the commercial enrollee price schedule enters into the model.

The price adjustment formula is fixed by the state and uses 3M’s All Patient Refined Diagnosis Related Groups (APR-DRGs) (Massachusetts 2010; CHIA 2015a). Insurers then rank hospitals by their base prices (simply called “prices” from here on), and determine hospitals’ tiers based on those prices. Some insurers make further adjustments to the assigned tiers based on hospitals’ geographic isolation or negotiated prices with the hospital system’s affiliated physician groups. However, such adjustments are generally minimal, affecting 0–13% of hospitals in an insurer’s network.³² Discussions with the provider contracting divisions of several anonymous, large Massachusetts insurers indicate that this is an accurate representation of their negotiations with providers and their network design.

In the estimation, I simplify the bargaining problem by using the base price as the immediate object of negotiations. For a given patient’s diagnosis, the price paid to the hospital is then the product of the base price and a production resource intensity multiplier for that diagnosis using the APR-DRG system. This parameterization of total price as a base price multiplied by a disease weight is motivated by institutional features of pricing in health care. Due to the large number of health care services for which prices must be agreed upon, prices are not typically negotiated service-by-service. Instead, payers such as Medicare and private insurers often negotiate a base price which is then scaled by a measure of resource intensity, such as a DRG weight or a Relative Value Unit (RVU). This parameterization of prices is also used by Gowrisankaran, Nevo and Town (2015) and Ho and Lee (2015).³³

The bargaining model must accommodate the key features of the data. First, unlike in a standard Nash bargaining model where only the discrete outcomes of agreement and disagreement are possible, the model must nest multiple discrete possibilities for the agreement outcome (corresponding to different hospital tiers). Second, the model must rationalize the fact that the distribution of hospital prices in the market is smooth. I accommodate both of these features by introducing a smooth stochastic mapping from negotiated price to hospital tier. The most straightforward extension of standard bargaining models to the context of tiered networks is price bargaining where the negotiated price maps deterministically into the hospital’s tier. However, this version of the model yields substantial price bunching just below the would-be cutoffs between tiers, which is inconsistent with the data.³⁴ Therefore, in the bargaining estimation, I rely on the stochastic tier version of the model described in this section.


Figures 3 and 4a show the distribution of Harvard Pilgrim’s negotiated prices with Massachusetts hospitals (excluding satellite campuses). Prices are reported as multiples of the insurer’s

³²In principle, insurers can use quality metrics in addition to price in setting hospital tiers (Massachusetts 2010). However, I find that in practice, including hospital quality measures does not change hospitals’ tier assignments relative to a baseline of using price alone. The overwhelming majority of Massachusetts hospitals score very high on most quality measures during the study period, typically above 90 points out of a possible 100, so that most of the variation across providers is in prices rather than measured quality.

³³While other payment arrangements also exist in the market, such as disease-specific negotiated prices and global payments, discussions with Massachusetts hospitals and insurers suggest that base prices scaled by diagnosis severities are often the starting point for negotiations even when other payment arrangements are used.

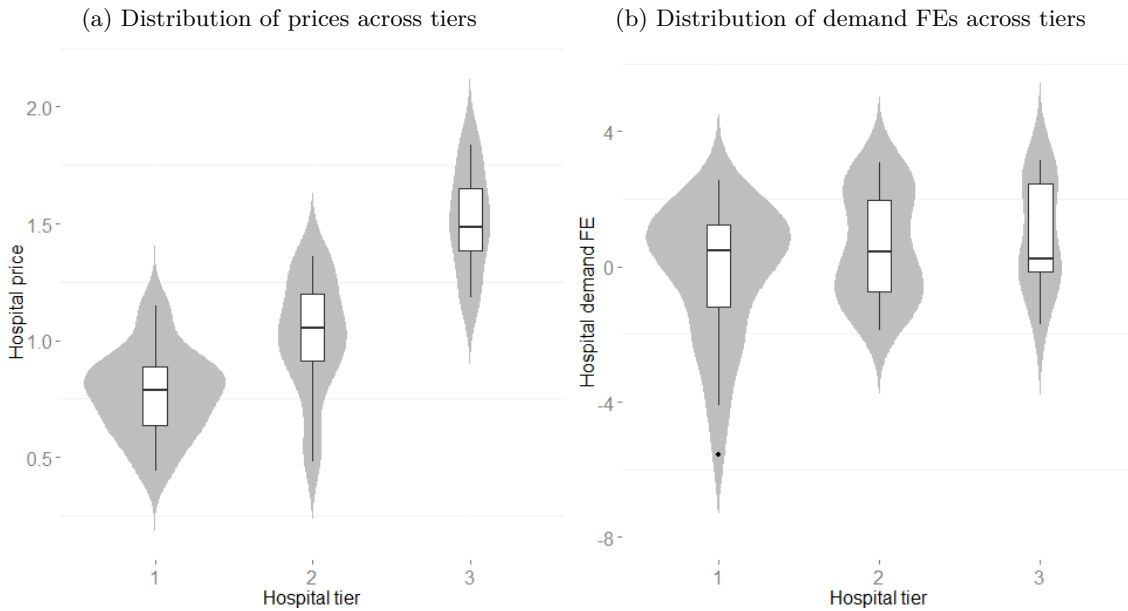
³⁴There exist parameter combinations that generate an interior solution for price that falls strictly between the tier cutoffs, but their ranges are narrow and would not hold for more than a small fraction of hospitals in a given market.

Figure 3: Harvard Pilgrim's distribution of negotiated hospital prices (2011)



The distribution of hospital prices in Harvard Pilgrim's network is smooth; notably, there is no bunching of prices in regions that are on the threshold between tiers.


Figure 4: Harvard Pilgrim's negotiated hospital prices (2011) and hospital demand fixed effects



Figures 4a and 4b show the price and hospital demand fixed effects distributions by tie, where tier 1 is the most preferred tier with the lowest out-of-pocket price. The black-bordered boxes represent the 25th percentile, median, and 75th percentile of the vertical axis quantity (hospital prices or demand fixed effects) in each tier; the gray shaded areas represent densities. The bulk of the price distribution for each tier does not overlap with the bulk of the prices in other tiers. The distribution of hospital demand fixed effects is comparable across tiers.

mean hospital price (data confidentiality considerations preclude reporting dollar amounts). There is no bunching of prices in regions around the threshold between tiers (Figure 3). While the bulk of the price distribution mass within each tier does not overlap with the bulk of prices in other tiers, there are no hard price cutoffs separating the tiers (Figure 4a). Furthermore, tiers are determined primarily by prices, rather than by consumer valuation of hospitals. Figures 4a and 4b show the distribution across tiers of hospital prices and consumer valuation of hospitals, as measured by hospital fixed effects from the demand model. Whereas the distribution of negotiated prices varies substantially across tiers, the distribution of hospital demand fixed effects is stable and provides little explanatory power for tier assignment. Motivated by these observations, I fit smooth functions mapping a hospital’s price as a multiple of the overall mean price across hospitals to the hospital’s probability of placement in each of three tiers. I use generalized logistic functions for the most and least preferred tiers, denoted $G^t(p) = \Pr(\text{tier} = t|p)$, where $t \in \{1, 3\}$ is the tier and p is price. For the middle tier, the probability is modeled as $G^2(p) = 1 - [G^1(p) + G^3(p)]$. Each tier mapping is continuous and differentiable in price, with the derivative with respect to price denoted $g^t(p) = \frac{\partial}{\partial p} G^t(p)$. The functions $G^1(p), G^3(p)$ are fit to Harvard Pilgrim’s observed tiers and prices for 2010–2013 using nonlinear least squares.

Figure 5: Mapping negotiated prices to tiers (Harvard Pilgrim)



Fitted generalized logistic functions mapping Harvard Pilgrim’s negotiated hospital prices to the hospitals’ probability of being in each tier, $G^t(p), t \in \{1, 2, 3\}$. The stochastic, smooth nature of the price-to-tier mapping implies that the model is smooth in negotiated prices.

The fitted price-to-tier mappings for Harvard Pilgrim are shown in Figure 5. Low prices imply a higher probability of the hospital ending up in the most preferred tier with the lowest out-of-pocket price (tier 1). As price increases, the probability of placement in the least preferred, highest out-of-pocket price tier (tier 3) approaches unity. The mapping from relative price to tier approximates the

insurer's tiering strategy in a computationally tractable fashion. The approximation relies on the assumption that the mapping from relative price to hospital tier remains constant in counterfactual scenarios. This assumption does not require the mapping from a hospital's raw price to its tier to remain unchanged in new equilibria. Rather, it only requires the less restrictive condition that the mapping from a hospital's price relative to other hospitals in the insurer's network remain constant. This allows for the mapping from raw prices to tiers to change as the distribution of prices shifts in new equilibria.³⁵ The mapping also relies on the assumption that hospitals' tier probabilities are independent from other hospitals in the network, conditional on the mean price.

The mapping is assumed to be observed by all insurers and hospitals at the time of price negotiations. Each player's expected surplus from agreement in negotiations is therefore an expectation over the possible tier placements resulting from the negotiated price. Uncertainty over tier placement is resolved after all bilateral insurer-hospital negotiations have concluded. The remainder of this section presents the model for insurer and hospital surplus, and the Nash bargaining solution under tiered hospital networks.

6.2 Insurer Objectives

The insurer's profits are determined by premiums, enrollment, negotiated hospital prices, hospital utilization, and consumer out-of-pocket payments. Each insurer $\mathcal{M} \in M$ may have multiple health insurance plans $m \in \mathcal{M}$. A plan's total enrollment is given by integrating a household ι 's probability of enrollment across households that have the plan in their choice sets, based on the distributions of sickness probabilities f_{id} and consumer characteristics x_{id} for household members $i \in \iota$. Note that plan shares $s_{m\iota}$, hospital shares σ_{mhid} , and premiums $r_{m\iota}$ depend on the characteristics of other plans in the market; these are suppressed from the notation for parsimony but enter into the empirical application. Consider first the insurer's profit when hospital tiers are already known, after the bargaining and tier determination stage of the game outlined in Section 4. Insurer \mathcal{M} 's total revenue from plan m is

$$Rev_m = \sum_{\iota \in I} (r_{m\iota} s_{m\iota})$$

where $r_{m\iota}$ is the plan's premium for household ι (individual or family premium, depending on the household's size), and $s_{m\iota}$ is household ι 's probability of enrollment in plan m . Note that the share of consumers enrolling in the plan depends on the other plan offerings in the market.

The plan's total cost is equal to the insurer's expected outlays for plan m 's enrollees to visit their chosen hospitals h

$$Cost_m = \sum_{\iota \in I} \left(s_{m\iota} \sum_{i \in \iota} \sum_{d \in D} \left(f_{id} \sum_{h \in H} (\sigma_{mhid} [l_d p_{\mathcal{M}h} - c_{mh}]) \right) \right)$$

³⁵This setup still assumes that the insurer's mapping is somewhat exogenous, rather than treating it fully as an equilibrium object. Future work could endogenize insurers' tier-setting formulas. Since little is understood about the effect of demand-side incentives on health care prices, the current setup represents a non-negligible step forward in the understanding of tiered networks and other innovations in health insurance design.

where σ_{mhid} is consumer i 's probability of going to hospital h when sick with diagnosis d , which the consumer contracts with probability f_{id} ; $p_{\mathcal{M}h}$ is the baseline negotiated price; l_d is the disease-specific multiplier adjusting the base price to the diagnosis; and c_{mh} is the consumer's out-of-pocket price portion for hospital h in plan m . The out-of-pocket price is set according to hospital h 's tier when m is a tiered-network plan, and is assumed not to vary across hospitals when m is a non-tiered plan.³⁶ Since the plans with tiered networks in my data use copays rather than coinsurance, the consumer's out-of-pocket price c_{mh} is an absolute amount that does not vary with price, conditional on the hospital's tier. For purposes of brevity, denote by $V_{mhi} = \sum_{d \in D} f_{id} \sigma_{mhid} l_d$ hospital h 's total expected casemix-adjusted volume for consumer i in plan m , and by $C_{mhi} = \sum_{d \in D} f_{id} \sigma_{mhid} c_{mh}$ consumer i 's total expected out-of-pocket payments to hospital h under plan m . Then insurer \mathcal{M} 's expected profit from plans $m \in \mathcal{M}$, given hospitals' tiers in its network, is

$$\Pi_{\mathcal{M}}^{known tiers} = \sum_{i \in I} \sum_{m \in \mathcal{M}} \left(s_{mi} \left[r_{mi} - \sum_{i \in i} \sum_{h \in H} (p_{\mathcal{M}h} V_{mhi} - C_{mhi}) \right] \right).$$

Now consider stage 1 of the game outlined in Section 4, before uncertainty over hospital tiers is resolved. The insurer's expected profit is now an expectation over the tier assignments for the hospitals in its network. Denote by $\tau \in T$ the possible permutations of hospitals' tiers in the insurer's network, with each hospital h 's tier denoted $t_h \in \tau$. Then the insurer's expected profit can be expressed as

$$\Pi_{\mathcal{M}} = \sum_{\tau \in T} \left\{ \prod_{t_h \in \tau} G_{\mathcal{M}}^{t_h}(p_{\mathcal{M}h}) \left\{ \sum_{i \in I} \sum_{m \in \mathcal{M}} s_{mi}^{\tau} \left[r_{mi}^{\tau} - \sum_{i \in i} \sum_{h \in H} (p_{\mathcal{M}h} V_{mhi} - C_{mhi}) \right] \right\} \right\} \quad (5)$$

where $\prod_{t_h \in \tau} G_{\mathcal{M}}^{t_h}(p_{\mathcal{M}h})$ is the probability of network tier permutation τ as a function of negotiated hospital prices. Due to the assumption of independence of $G_{\mathcal{M}}^t(p)$ across hospitals, this probability is simply the product of the individual hospitals' probabilities of the tier they occupy in a given tier permutation, $\Pr(tier_h = t_h \in \tau)$. The sum of the network tier probabilities across all tier permutations $\sum_{\tau \in T} \left\{ \prod_{t_h \in \tau} G_{\mathcal{M}}^{t_h}(p_{\mathcal{M}h}) \right\}$ is equal to one. Note that plan shares s_{mi}^{τ} , hospital shares σ_{mhid}^{τ} , copays c_{mh}^{τ} , and premiums r_{mi}^{τ} change with network tiers τ . In negotiations with a hospital, the insurer's objective is composed of its expected profit in the case of agreement less the expected profit in the case of disagreement, with profits defined by Equation 5.³⁷

³⁶In practice, some non-tiered plans in Massachusetts use coinsurance for hospital care, so that consumers' out-of-pocket cost of care is a fixed percentage of the negotiated price with a given hospital. This can vary across hospitals. However, I make the simplifying assumption that consumers' *expected* out-of-pocket cost for hospital care in a non-tiered plan does not vary across hospitals. This assumption is supported by the finding in the literature that consumers are generally uninformed about hospital prices before they are billed; and even savvy consumers who ask for price quotes typically get poor response rates (Bebinger 2014).

³⁷A list of the symbols used throughout the paper is provided in the appendix (page 60).

6.3 Hospital Objectives

Now consider a hospital's expected profit given a set of prices and networks in the market. The hospital's profits are determined by insurance plan enrollments, the hospital's share of each plan's enrollees, and negotiated prices. The hospital may treat patients from multiple insurers $\mathcal{N} \in M$, and changes in its network tier may affect insurance plan enrollments via consumers' changing valuation of the hospital network and via premiums. When tier assignments are known, the hospital's expected profit is

$$\begin{aligned}\Pi_h^{known\ tiers} &= \sum_{\iota \in I} \sum_{n \in \mathcal{N} \in M} \left[s_{n\iota} \sum_{i \in \iota} \sum_{d \in D} (f_{id} \sigma_{nhid} l_d [p_{\mathcal{N}h} - k_h]) \right] \\ &= \sum_{\iota \in I} \sum_{n \in \mathcal{N} \in M} \left[s_{n\iota} \sum_{i \in \iota} (p_{\mathcal{N}h} - k_h) V_{nhi} \right]\end{aligned}$$

where $V_{nhi} = \sum_{d \in D} f_{id} \sigma_{nhid} l_d$ is hospital h 's total expected casemix-adjusted volume for consumer i in plan n ; $p_{\mathcal{N}h}$ is the hospital's baseline negotiated price with insurer $\mathcal{N} \in M$, which applies to all of the insurer's plans $n \in \mathcal{N}$; k_h is the hospital's marginal cost of treating a patient with a diagnosis severity weight of one; and l_d is the disease-specific resource intensity use multiplier. Each hospital is assumed to have a baseline marginal treatment cost per patient k_h that is constant for all commercially insured patients with a diagnosis severity weight of one. For a given patient's diagnosis, the cost and price are scaled to $l_d k_h$ and $l_d p_{\mathcal{N}h}$, respectively, based on the resource intensity associated with diagnosis d . As is typical in the industrial organization literature, marginal costs are not observed. They are inferred from the solution to the bargaining model and used in the counterfactual exercises.

Now consider the stage of the game before uncertainty over hospital tiers is resolved. The hospital's expected profit is now an expectation over its own tier assignment and the tier assignments of other hospitals in Harvard Pilgrim's network, with other insurers' networks taken as fixed. As before, denote by $\tau \in T$ the possible permutations of hospitals' tiers in the network, with each hospital h 's tier denoted $t_h \in \tau$. Denote Harvard Pilgrim as insurer \mathcal{M} with plans m , and the other insurers in the market $\mathcal{N} \in M \setminus \mathcal{M}$. Then the hospital's expected profit can be expressed as

$$\begin{aligned}\Pi_h &= \sum_{\tau \in T} \left\{ \prod_{t_h \in \tau} G_{\mathcal{M}}^{t_h}(p_{\mathcal{M}h}) \left\{ \sum_{\iota \in I} \left[\sum_{m \in \mathcal{M}} s_{m\iota}^\tau \sum_{i \in \iota} (p_{\mathcal{M}h} - k_h) V_{mhi}^\tau \right. \right. \right. \\ &\quad \left. \left. + \sum_{n \in \mathcal{N} \in M \setminus \mathcal{M}} s_{n\iota}^\tau \sum_{i \in \iota} (p_{\mathcal{N}h} - k_h) V_{nhi}^\tau \right] \right\} \right\} \quad (6)\end{aligned}$$

where $\prod_{t_h \in \tau} G_{\mathcal{M}}^{t_h}(p_{\mathcal{M}h})$ is the probability of Harvard Pilgrim's network tier permutation τ as a function of negotiated hospital prices. In negotiations with Harvard Pilgrim, the hospital's objective is composed of its expected profit in the case of agreement less the expected profit in the case of disagreement, with profits defined by Equation 6.

6.4 Premium Setting

In negotiating new prices with hospitals in its network, an insurer can adjust plan premiums to reflect the new cost structure implied by hospital prices and tiers. I let Harvard Pilgrim adjust its premiums for GIC plans, the section of the market on which plan demand is estimated (see Section 3.2). The insurer offers two plans on the GIC market, one that includes all Massachusetts hospitals in its network and another using a narrow network (Table 2). There are four premiums associated with these plans: one for individual coverage and another for family coverage for each plan. Motivated by the institutional features of the GIC described in Section 5.4, I take the ratios of each premium in relationship to the full-network plan’s individual premium as exogenous. The ratio of a plan’s family premium to its individual premium is 2.4 and the ratio of the insurer’s narrow-network plan to the insurer’s full-network plan is 0.8. In the bargaining estimation, I fix these ratios between plan premiums but allow all four premiums to shift together in response to negotiated prices and networks.

Premium shifting as a function of prices and tiers is modeled assuming the insurer has a constant profit margin bound by regulation. The state of Massachusetts has minimum medical loss ratio (MLR) regulations, which set an upper bound on insurer profit margins by dictating a minimum fraction of premium revenue that insurers must spend on their enrollees’ medical care (Massachusetts 2010). The MLR regulations require insurers to spend at least 85% of premium revenue in large-group plans on medical expenses³⁸, and insurers whose medical spending falls short of the target are required to issue premium rebates to their enrollees. Moreover, the GIC’s actuaries are directly involved in premium-setting for the plans offered to its employees. On the GIC market, 90% of premiums are generally disbursed in the form of payment to health care providers (GIC 2011). I therefore assume that hospital price changes are passed through to GIC premiums at a rate of $1/0.9$ times the change in expected spending as a result of the price change. Simulations using my hospital and plan demand estimates indicate that the MLR does indeed bind for Harvard Pilgrim’s GIC plans at current market conditions, and the MLR assumption provides substantial modeling and computational advantages for the bargaining model.

In response to a change in the negotiated price $p_{\mathcal{M}h}$ with hospital h , Harvard Pilgrim’s baseline premium r changes as a function of the change in expected spending

$$\begin{aligned}\frac{\partial r}{\partial p_{\mathcal{M}h}} &= \frac{\partial}{\partial p_{\mathcal{M}h}} \left[\frac{1}{\lambda} \sum_{i \in \iota \in m} \left((p_{\mathcal{M}h} V_{mhi}^\tau - C_{mhi}^\tau) + \sum_{j \in H \setminus h} (p_{\mathcal{M}j} V_{mji}^\tau - C_{mji}^\tau) \right) \right] \\ &= \frac{1}{\lambda} \sum_{i \in \iota \in m} V_{mhi}^\tau\end{aligned}$$

where $\lambda = 0.9$ is the GIC-specific MLR and the tier structure is held fixed. Conditional on hospitals’ tiers, hospital shares within the plan remain unchanged as prices change because consumers only respond to the out-of-pocket copays dictated by tiers, so that $\partial V_{mji}^\tau / \partial p_{\mathcal{M}h} = 0 \forall j \in H \setminus h$. Next,

³⁸On the individual and small-group market, the MLR requires insurers to spend 90% of premium revenue on medical expenses.

consider the effect of the premium change on Harvard Pilgrim's enrollment. Since the price change is passed through to premiums for both of Harvard Pilgrim's plans, each of its $m \in \mathcal{M}$ plans' share responds not only to its own premium change but also to the insurer's other plan's $m' \in \mathcal{M} \setminus m$ premium change. Accounting for this multi-product feature of the insurer, the change in a Harvard Pilgrim plan's probability of enrolling household ι in response to negotiated price is

$$\begin{aligned} \frac{\partial s_{m\iota}^\tau}{\partial p_{\mathcal{M}h}} &= \frac{\partial s_{m\iota}^\tau}{\partial r_{m\iota}^\tau} \cdot \frac{\partial r_{m\iota}^\tau}{\partial p_{\mathcal{M}h}} \\ &= -\frac{\delta_1}{\lambda} (s_{m\iota}^\tau) \left[(1 - s_{m\iota}^\tau) \left(\sum_{i \in \iota' \in m} V_{mhi}^\tau \right) - \sum_{m'} s_{m'\iota}^\tau \left(\sum_{i \in \iota' \in m} V_{mhi}^\tau \right) \right] \end{aligned}$$

where $\delta_1 > 0$ is the coefficient on premiums in plan demand. This formula is similar to the standard derivative of market share with respect to price in multinomial logit demand models (Train 2002), with the addition of terms to account for the fact that Harvard Pilgrim is a multi-product firm whose prices (premiums) change in tandem for multiple products.

Other insurers on the GIC market may also experience a change in enrollment in response to Harvard Pilgrim's negotiated prices with hospitals, since the new Harvard Pilgrim premiums will affect consumers' probability of enrollment in all plans in the market. For a plan $n \notin \mathcal{M}$ offered by a different insurer, the change in enrollment by household ι is

$$\begin{aligned} \frac{\partial s_{n\iota}^\tau}{\partial p_{\mathcal{M}h}} &= \sum_{m \in \mathcal{M}} \left(\frac{\partial s_{n\iota}^\tau}{\partial r_{m\iota}^\tau} \cdot \frac{\partial r_{m\iota}^\tau}{\partial p_{\mathcal{M}h}} \right) \\ &= \sum_{m \in \mathcal{M}} \left(\delta_1 (s_{n\iota}^\tau) (s_{m\iota}^\tau) \cdot \frac{1}{\lambda} \left(\sum_{i \in \iota' \in m} V_{mhi}^\tau \right) \right) \\ &= \frac{\delta_1}{\lambda} (s_{n\iota}^\tau) \sum_{m \in \mathcal{M}} \left(s_{m\iota}^\tau \sum_{i \in \iota' \in m} V_{mhi}^\tau \right) \end{aligned}$$

The MLR assumption implies that all plan enrollments are differentiable in negotiated hospital prices. This structure provides substantial modeling and computational advantages for the bargaining model by abstracting away from Bertrand-Nash competition in premiums, for which there is no closed-form solution. Since my simulations suggest that the MLR binds for Harvard Pilgrim's GIC plans, I take this to be a reasonable approximation of the market. This approach provides a middle road between papers on insurer-hospital bargaining that abstract from plan demand altogether (e.g., Gowrisankaran et al. (2015)) and those modeling a full Bertrand-Nash game in premiums (e.g., Ho and Lee (2015)).

6.5 Nash Bargaining

The object of the negotiations is the price $p_{\mathcal{M}h}$ paid to hospital h for treating each of Harvard Pilgrim's patients, which enters into both parties' expected profits. Prices determine hospital tiers in Harvard Pilgrim's network, which in turn affect hospital volume conditional on plan enrollment,

insurer costs, and household WTP for plans. Harvard Pilgrim also adjusts plan premiums in response to negotiated prices and hospital tiers, and premiums affect both per-enrollee revenues and plan enrollments.

Denote by $H_{\mathcal{M}} \subseteq H$ the subset of Massachusetts hospitals that are in Harvard Pilgrim's network. The insurer's surplus from coming to an agreement with a given hospital, $S_{\mathcal{M}(h)}$, can then be denoted by $S_{\mathcal{M}(h)} = \Pi_{\mathcal{M}}(h \in H_{\mathcal{M}}) - \Pi_{\mathcal{M}}(h \notin H_{\mathcal{M}})$; similarly, hospital h 's surplus is $S_{h(\mathcal{M})} = \Pi_h(h \in H_{\mathcal{M}}) - \Pi_h(h \notin H_{\mathcal{M}})$. Since the equilibrium in this market is such that all hospitals reach agreements with all insurers, the second term is calculated assuming that all other hospitals $j \in H \setminus h$ remain in the insurer's network in the event of a disagreement.³⁹

The equilibrium negotiated price $p_{\mathcal{M}h}^*$ maximizes the Nash bargaining product $(S_{\mathcal{M}(h)})^{b_{\mathcal{M}(h)}} \cdot (S_{h(\mathcal{M})})^{b_{h(\mathcal{M})}}$, where $b_{\mathcal{M}(h)}$ and $b_{h(\mathcal{M})}$ are the insurer's and hospital's respective bargaining weights, normalized so that $b_{\mathcal{M}(h)} = 1 - b_{h(\mathcal{M})}$. The bargaining model is solved for unobserved hospital marginal costs per patient k_h using first-order conditions with respect to negotiated prices. The assumptions of a stochastic price to tier mapping and MLR premium setting yield a Nash bargaining solution that is continuous and differentiable in prices. Taking the logarithm of the Nash bargaining product, the first order condition for the price $p_{\mathcal{M}h}^*$ is then given by

$$\begin{aligned} b_{\mathcal{M}(h)} \frac{\partial}{\partial p_{\mathcal{M}h}^*} \log [\Pi_{\mathcal{M}}(h \in H_{\mathcal{M}}) - \Pi_{\mathcal{M}}(h \notin H_{\mathcal{M}})] &= -b_{h(\mathcal{M})} \frac{\partial}{\partial p_{\mathcal{M}h}^*} \log [\Pi_h(h \in H_{\mathcal{M}}) - \Pi_h(h \notin H_{\mathcal{M}})] \\ b_{\mathcal{M}(h)} \frac{\frac{\partial}{\partial p_{\mathcal{M}h}^*} \Pi_{\mathcal{M}}(h \in H_{\mathcal{M}})}{\Pi_{\mathcal{M}}(h \in H_{\mathcal{M}}) - \Pi_{\mathcal{M}}(h \notin H_{\mathcal{M}})} &= -b_{h(\mathcal{M})} \frac{\frac{\partial}{\partial p_{\mathcal{M}h}^*} \Pi_h(h \in H_{\mathcal{M}})}{\Pi_h(h \in H_{\mathcal{M}}) - \Pi_h(h \notin H_{\mathcal{M}})} \end{aligned} \quad (7)$$

The left-hand side of Equation 7 is the insurer's component of the first-order conditions, while the right-hand side is the hospital's component into which hospital costs k_h enter. Profits are defined as in Equations 5 and 6, and the first-order conditions are solved using the derivatives of premiums and plan shares with respect to price derived in Section 6.4.

Due to notational complexity, I present to the left-hand and right-hand side of the first-order conditions separately. To further simplify notation, denote by $\mathcal{G}_{\mathcal{M}h}^{\tau}$ the product of probabilities of network tier permutation τ for all hospitals except h , that is $\mathcal{G}_{\mathcal{M}h}^{\tau} = \prod_{t_j \in \tau \setminus t_h} G_{\mathcal{M}}^{t_j}(p_{\mathcal{M}j})$. Then the total probability of network arrangement τ is $G_{\mathcal{M}}^{t_h \in \tau}(p_{\mathcal{M}h}) \mathcal{G}_{\mathcal{M}h}^{\tau}$, and its derivative with respect to hospital h 's price is $g_{\mathcal{M}}^{t_h \in \tau}(p_{\mathcal{M}h}) \mathcal{G}_{\mathcal{M}h}^{\tau}$, where $g_{\mathcal{M}}^{t_h \in \tau}(p_{\mathcal{M}h}) = \frac{\partial}{\partial p_{\mathcal{M}h}} G_{\mathcal{M}}^{t_h \in \tau}(p_{\mathcal{M}h})$.⁴⁰

Insurer's Component of FOCs: The numerator of the insurer's component of the first-order conditions (the left-hand side of Equation 7) is

$$\begin{aligned} \text{insurer numerator} &= \\ \text{tier likelihood effect} &\left\langle \sum_{\tau} g_{\mathcal{M}}^{t_h \in \tau}(p_{\mathcal{M}h}) \mathcal{G}_{\mathcal{M}h}^{\tau} \left\{ \sum_{i \in I} \sum_{m \in \mathcal{M}} s_{mi}^{\tau} \left[r_{mi}^{\tau} - \sum_{i \in I} \sum_{h \in H} (p_{\mathcal{M}h} V_{mhi}^{\tau} - C_{mhi}^{\tau}) \right] \right\} \right\rangle \end{aligned} \quad (8)$$

³⁹Hospital systems in Massachusetts are prohibited by state regulation from negotiating their network status as a group.

⁴⁰A list of the symbols used throughout the paper is provided in the appendix (page 60).

$$\begin{aligned}
\text{plan enrollment effect} & \left\langle + \sum_{\tau} G_{\mathcal{M}}^{t_h \in \tau} (p_{\mathcal{M}h}) \mathcal{G}_{\mathcal{M}h}^{\tau} \left\{ \sum_{i \in I} \sum_{m \in \mathcal{M}} \left[r_{m\iota}^{\tau} - \sum_{i \in \iota} \sum_{h \in H} (p_{\mathcal{M}h} V_{mhi}^{\tau} - C_{mhi}^{\tau}) \right] \right. \right. \\
& \quad \cdot \left. \left[-\frac{\delta_1}{\lambda} (s_{m\iota}^{\tau}) \left[(1 - s_{m\iota}^{\tau}) \left(\sum_{i \in \iota} V_{mhi}^{\tau} \right) - \sum_{m' \in \mathcal{M} \setminus m} s_{m'\iota} (\sum_{i \in \iota} V_{m'hi}^{\tau}) \right] \right] \right\} \\
\text{premium and price effect} & \left\langle + \sum_{\tau} G_{\mathcal{M}}^{t_h \in \tau} (p_{\mathcal{M}h}) \mathcal{G}_{\mathcal{M}h}^{\tau} \left\{ \sum_{i \in I} \sum_{m \in \mathcal{M}} s_{m\iota}^{\tau} \left(\frac{1}{\lambda} - 1 \right) \left(\sum_{i \in \iota} V_{mhi}^{\tau} \right) \right\} \right.
\end{aligned}$$

which captures the multiple channels through which changes in the negotiated price affect insurer surplus. The first term captures the hospital tier effect: as price moves, so too does the probability that the hospital will be in a given tier in the insurer's network. An increase in the negotiated price reduces the probability of the hospital being assigned to a preferred tier. Less preferred tier assignment affects the insurer's surplus by reducing the hospital's volume of the insurer's patients and raising the portion of price borne out-of-pocket by consumers. The second term captures the effect on total premium revenue and total costs as a function of the number and composition of enrolled households. Households reoptimize their enrollment decisions as a function of their changed WTP for the insurer's hospital network and the changing premium. Higher price and less preferred tier assignment increase the expected out-of-pocket price to consumers, which reduces consumer WTP for the network but has an ambiguous effect on insurer costs and premiums. If the negotiated price is in a region where tier probability is changing quickly, then total spending increases due to small increases in price may be offset by consumers' greater out-of-pocket spending, so the insurer's spending and therefore premiums may fall. For larger changes in price, the effect of price increases is generally to increase spending and premiums, which will reduce the insurer's surplus if the premium changes are large enough to reduce enrollment. Finally, the third term captures the direct effect of changing premiums and prices. As price conditional on tier rises, per-household premium revenue rises but per-admission costs also rise.

The denominator of the insurer's component of the FOCs is

$$\begin{aligned}
\text{insurer denominator} & = \tag{9} \\
& \sum_{\tau} G_{\mathcal{M}}^{t_h \in \tau} (p_{\mathcal{M}h}) \mathcal{G}_{\mathcal{M}h}^{\tau} \left\{ \sum_{i \in I} \sum_{m \in \mathcal{M}} \right. \\
\text{profits, } h \in \text{network} & \left\langle s_{m\iota}^{\tau} \left[r_{m\iota}^{\tau} - \sum_{i \in \iota} (p_{\mathcal{M}h} V_{mhi}^{\tau} - C_{mhi}^{\tau}) - \sum_{i \in \iota} \sum_{j \in H \setminus h} (p_{\mathcal{M}j} V_{mji}^{\tau} - C_{mji}^{\tau}) \right] \right. \\
\text{profits, } h \notin \text{network} & \left\langle -s_{m\iota}^{\tau \setminus h} \left[r_{m\iota}^{\tau \setminus h} - \sum_{i \in \iota} \sum_{j \in H \setminus h} (p_{\mathcal{M}j} V_{mji}^{\tau \setminus h} - C_{mji}^{\tau \setminus h}) \right] \right\}
\end{aligned}$$

where $\tau \setminus h$ is the insurer's network when it is identical to tier permutation τ except that hospital h is out of network. Premiums, plan enrollments, and hospital volumes conditional on enrollment are allowed to vary based on whether the hospital is in network. Other hospitals' prices are held fixed since the FOC is a Nash equilibrium object.

The overall effect of price on the insurer's portion of the FOC as defined in Equations 8 and 9 is ambiguous. The effect of price on the insurer's surplus will depend on household WTP for the hospital, differences in out-pocket-price across copays, and enrollment response to network WTP and premiums.

Hospital's Component of FOCs: The numerator of the hospital's component of the first-order conditions (the right-hand side of Equation 7) is

$$\begin{aligned}
& \text{hospital numerator} = & (10) \\
& \text{tier likelihood effect} \left\langle \sum_{\tau} g_{\mathcal{M}}^{t_h \in \tau} (p_{\mathcal{M}h}) \mathcal{G}_{\mathcal{M}h}^{\tau} \left\{ (p_{\mathcal{M}h} - k_h) \sum_{\iota \in I} \sum_{m \in \mathcal{M}} s_{m\iota}^{\tau} \left(\sum_{i \in \iota} V_{mhi}^{\tau} \right) \right. \right. \\
& \quad \left. \left. + \sum_{\mathcal{N} \in M \setminus \mathcal{M}} (p_{\mathcal{N}h} - k_h) \sum_{\iota \in I} \sum_{n \in \mathcal{N}} s_{n\iota}^{\tau} \left(\sum_{i \in \iota} V_{nhi}^{\tau} \right) \right\} \right. \\
& \text{this insurer enrollt. effect} \left\langle + \sum_{\tau} G_{\mathcal{M}}^{t_h \in \tau} (p_{\mathcal{M}h}) \mathcal{G}_{\mathcal{M}h}^{\tau} \left\{ (p_{\mathcal{M}h} - k_h) \sum_{\iota \in I} \sum_{m \in \mathcal{M}} -\frac{\delta_1}{\lambda} s_{m\iota}^{\tau} \left(\sum_{i \in \iota} V_{mhi}^{\tau} \right) \right. \right. \\
& \quad \left. \left. \cdot \left[(1 - s_{m\iota}^{\tau}) \left(\sum_{i \in \iota} V_{mhi}^{\tau} \right) - \sum_{m' \in \mathcal{M} \setminus m} s_{m'\iota}^{\tau} \left(\sum_{i \in \iota} V_{m'hi}^{\tau} \right) \right] \right\} \right. \\
& \text{other insurers enrollt. effect} \left\langle + \sum_{\tau} G_{\mathcal{M}}^{t_h \in \tau} (p_{\mathcal{M}h}) \mathcal{G}_{\mathcal{M}h}^{\tau} \left\{ \sum_{\mathcal{N} \in M \setminus \mathcal{M}} (p_{\mathcal{N}h} - k_h) \right. \right. \\
& \quad \left. \left. \sum_{\iota \in I} \sum_{n \in \mathcal{N}} \frac{\delta_1}{\lambda} s_{n\iota}^{\tau} \left(\sum_{i \in \iota} V_{nhi}^{\tau} \right) \left(\sum_{m \in \mathcal{M}} s_{m\iota}^{\tau} \left(\sum_{i \in \iota} V_{mhi}^{\tau} \right) \right) \right\} \right. \\
& \text{direct price effect} \left\langle + \sum_{\tau} G_{\mathcal{M}}^{t_h \in \tau} (p_{\mathcal{M}h}) \mathcal{G}_{\mathcal{M}h}^{\tau} \left\{ \sum_{\iota \in I} \sum_{m \in \mathcal{M}} s_{m\iota}^{\tau} \left(\sum_{i \in \iota} V_{mhi}^{\tau} \right) \right\} \right.
\end{aligned}$$

which captures the multiple channels through which changes in the negotiated price affect hospital surplus. The first term captures the hospital tier effect, analogous to its effect on the numerator of the insurer's component. As price increases, the probability of the hospital being assigned to a preferred tier falls, which affects the hospital's surplus by reducing the hospital's volume of patients from the insurer. The second and third terms capture changes in plan enrollment in the negotiating insurer's plans and the insurer's competitors' plans, as households reoptimize their enrollment decisions due to changing WTP and premiums. Less preferred tier assignment reduces household WTP for the insurer's network, which may result in consumers re-sorting to other insurers. However, the less-preferred tier placement may also lead to lower total spending by the plan, in which case premiums may fall enough to offset the decrease in WTP. The two plan enrollment terms capture the effect on the hospital's total volume and the fraction of that volume for which the hospital is reimbursed at each of the insurers' negotiated prices. The effect on hospital revenue is ambiguous, and depends on whether the bulk of the hospital's volume comes from its highest-reimbursing insurers. The plan enrollment effect allows hospitals to recapture the negotiating insurer's patients that are lost due to network or premium changes through other

insurers. Finally, the fourth term captures the direct effect of changing prices. As price rises, so too does the hospital's per-admission reimbursement for the insurer's patients.

The denominator of the hospital's component of the FOCs is

$$\begin{aligned}
\text{hospital denominator} &= \tag{11} \\
&\sum_{\tau} G_{\mathcal{M}}^{t_h \in \tau} (p_{\mathcal{M}h}) \mathcal{G}_{\mathcal{M}h}^{\tau} \left\{ \right. \\
\text{profits from insurer } \mathcal{M}, h \in \text{network} &\left\langle \begin{aligned} &(p_{\mathcal{M}h} - k_h) \sum_{\iota \in I} \sum_{m \in \mathcal{M}} s_{m\iota}^{\tau} \left(\sum_{i \in \iota} V_{mhi}^{\tau} \right) \end{aligned} \right. \\
\text{profits from other insurers, } h \in \text{network} &\left\langle \begin{aligned} &+ \sum_{\mathcal{N} \in M \setminus \mathcal{M}} (p_{\mathcal{N}h} - k_h) \sum_{\iota \in I} \sum_{n \in \mathcal{N}} s_{n\iota}^{\tau} \left(\sum_{i \in \iota} V_{nhi}^{\tau} \right) \end{aligned} \right. \\
\text{profits from other insurers, } h \notin \text{network} &\left\langle \begin{aligned} &- \sum_{\mathcal{N} \in M \setminus \mathcal{M}} (p_{\mathcal{N}h} - k_h) \sum_{\iota \in I} \sum_{n \in \mathcal{N}} s_{n\iota}^{\tau \setminus h} \left(\sum_{i \in \iota} V_{nhi}^{\tau \setminus h} \right) \end{aligned} \right\}
\end{aligned}$$

where $\tau \setminus h$ is the insurer's network when it is identical to tier permutation τ except that hospital h is out of network. Plan enrollments and hospital volumes conditional on enrollment are allowed to vary based on whether the hospital is in network. The hospital's negotiated prices with other insurers are held fixed since the FOC is a Nash equilibrium object.

The overall effect of price on the hospital's portion of the FOC as defined in Equations 10 and 11 is ambiguous. The effect of price on the hospital's surplus will depend on consumer preference for the hospital, as well as the hospital's effect on enrollment through overall network WTP and premiums. Generally, hospitals with very high negotiated prices, such as "star" hospitals, have little to lose from negotiating even higher prices. Such hospitals are overwhelmingly likely to be in the least preferred tier regardless of moderate changes in price, so that higher prices will result in higher per-patient revenue and no loss in volume conditional on plan enrollments. For hospitals with mid-range prices, either due to lower consumer preference or other factors, there is a clearer trade-off between gains in per-patient revenue and losses of volume as a result of higher prices.

Solving for Unobserved Parameters: Equations 8–11 are combined to solve the first-order conditions as given in Equation 7 for hospital marginal costs per patient k_h . The entry of tier probabilities $G_{\mathcal{M}}^{t_h \in \tau} (p_{\mathcal{M}h}) \mathcal{G}_{\mathcal{M}h}^{\tau}$ into the first-order conditions implies that prices are not linearly related to hospital costs, unlike in existing bargaining models that build on the inversion proposed by Berry (1994). Each hospital h 's cost parameter is identified from its first-order condition with respect to its price $p_{\mathcal{M}h}$. I calibrate the hospital cost parameter using Equation 7, focusing on a small subset of hospitals due to computational burden.

In the empirical application, the bargaining weights are assumed to be symmetric, $b_{\mathcal{M}(h)} = b_{h(\mathcal{M})} = 0.5$. The symmetry assumption implies that any observed differences in hospitals' mark-ups that are not rationalized by patient preferences through the demand model will load onto hospital marginal costs k_h . A higher bargaining weight for hospitals $b_{h(\mathcal{M})} > 0.5$ would imply market conditions closer to hospitals setting prices unilaterally (i.e. closer to insurers being price-

takers). Ho and Lee (2015) find estimates of bargaining weights that are not far from 0.5 in their primary specification, but Gowrisankaran et al. (2015) reject the hypothesis of symmetric bargaining weights for two of three insurers in their data.⁴¹

Dimensionality Reduction: The large number of possible permutations of hospitals' tiers in the insurer's network requires an approach to reducing the dimensionality of the problem. I leverage the fact that a hospital's volume is disproportionately affected by the network status of its closest competitors. The dimensionality of the expectation of insurer and hospital surplus over the set of all possible hospital tier permutations $T \ni \tau$ is exponential in the number of competing hospitals. For H hospitals in the market, there are $3^H + H \cdot 3^{H-1}$ possible permutations of hospital tiers, of which 3^H are permutations with all hospitals in the insurer's network and $H \cdot 3^{H-1}$ are the permutations excluding exactly one hospital from the network. In the Massachusetts hospital market with 72 general acute care hospitals,⁴² the number of computations required to take a full expectation over all possible network permutations is on the order of 5.6×10^{35} , far exceeding the computational capacity available to both firms and the econometrician.

I proceed by assuming that in price negotiations, firms take an expectation over the network status of only the closest $N_h < H$ competitors for the hospital in question, and hold fixed all other hospitals' tiers. The closest competitors are defined as those hospitals with which the negotiating hospital has the largest cross-price elasticities implied by the hospital demand model. In practice, the cross-price elasticity measure of closeness of competition is highly correlated with geographical closeness for the majority of hospitals.⁴³ The intuition for the assumption that negotiating firms only take an expectation over the closest competitors' tiers is that local market conditions are the most relevant information in the negotiation. A given hospital h 's negotiated price responds little to changes in the network status of a hospital j that barely competes for patients with h , as will be shown in my results by the many hospital pairs with cross-price elasticities of essentially zero (Table 12). Thus, the computational cost of taking j 's network status uncertainty into account typically exceeds the gain from the negligible adjustment in h 's price that a full expectation over j 's network status would imply.

7 Results

This section presents the results of my empirical analyses. First, I show that consumers do respond to out-of-pocket hospital price differences across providers, but the elasticity of demand is small. I then discuss the implications of the hospital demand estimates for household willingness to pay

⁴¹Future work will estimate the bargaining weights for each hospital-insurer pair by leveraging data on multiple years of price negotiations for each hospital. If hospital costs k_h are assumed to remain constant over time (a linear time trend can also be accommodated with an additional degree of freedom), then costs and bargaining weights are separately identified up to a normalization using first-order conditions from multiple years.

⁴²This includes hospitals' secondary satellite campuses.

⁴³The large, prestigious academic medical centers draw patients from a substantially broader geographic region than does the typical hospital, and their closest competitors as measured by cross-price elasticities are typically other academic medical centers.

for hospital networks. Next, I present evidence that households’ choice of health insurance plans responds to WTP for hospital networks. Finally, I present the estimates from the bargaining model. In Section 8, I use the estimates from the demand and bargaining models to conduct a counterfactual exercise of the effect of tiered networks on hospital demand and prices, relative to traditional non-tiered networks.

7.1 Hospital Demand

Estimates from the multinomial logit hospital choice model are shown in Table 11. Consistent with the literature on hospital choice, the coefficient on distance is negative and significant, implying that consumers are more likely to go to a hospital that is close by (Kessler and McClellan 2000; Town and Vistnes 2001; Capps et al. 2003; Ho 2006). Patients with cardiac or obstetric diagnoses are more likely to choose a hospital with a catheterization lab or a NICU, respectively. Older patients and patients with chronic conditions are more willing to travel to their chosen hospital. Hospital fixed effects also display a sensible pattern (not shown). The most prestigious hospitals in the state, such as Massachusetts General Hospital (MGH) and Brigham and Women’s Hospital, have among the largest estimated fixed effects, driven by their large share of patients from across the state and across diagnoses despite their high out-of-pocket prices.

The primary coefficient of interest is the coefficient on out-of-pocket price, specifically copays in this context. The negative and significant coefficient on price indicates that consumers do, indeed, respond to differences in out-of-pocket price when choosing hospitals. Since hospital demand is estimated as a multinomial logit discrete choice model, the elasticities implied by the model are a more useful quantity for interpretation than are the raw coefficients. The elasticities are reported in Table 12 and discussed below; the raw coefficients are broadly consistent with the emerging evidence on patient hospital choice under differential out-of-pocket pricing (Scanlon et al. 2008; Frank et al. 2015; Gowrisankaran et al. 2015).

The hospital demand model is robust to a number of alternative specifications. Estimates are stable when I allow price sensitivity to vary by income, as measured by the primary enrollee’s zip code’s income quintile (Table 20 in the appendix). Point estimates for price sensitivity are larger in magnitude for lower income quintiles, indicating that lower-income consumers may be more price-sensitive. Taken at face value, these results would suggest that the negative effect of price on demand is somewhat moderated for higher-income consumers, which is consistent with decreasing marginal returns over wealth or liquidity-constrained consumers. However, the differences in the price sensitivity coefficient across income groups are not statistically significant.

Furthermore, I do not find evidence that the estimates are biased by consumers sorting into plans that include their unobservably preferred hospitals at low out-of-pocket prices. Table 21 in the appendix shows the results of the control function estimation model described in Section 5.2, using the previous year’s copays to instrument for the current copays. If consumers were selecting into plans based on their unobserved preference for their preferred hospitals to be available at a low out-of-pocket price in the plan’s network, then failing to account for this endogeneity would

Table 11: Hospital choice model

	Hospital Choice	
Hospital Choice		
Copay (\$)	-0.0002***	(0.0001)
Distance (mi)	-0.1998***	(0.0019)
Distance ²	0.0007***	(0.0000)
Age \times distance	-0.0001***	(0.0000)
Male \times distance	0.0004	(0.0010)
Chronic cond \times distance	0.0205***	(0.0011)
Teaching \times distance	-0.0060***	(0.0014)
Beds \times distance	0.0000***	(0.0000)
% highly recommend \times distance	0.0047***	(0.0006)
Cardiac CCS \times cath lab	0.6617***	(0.0993)
Obstetric CCS \times NICU	0.3253***	(0.0387)
Nerv, circ, musc CCS \times MRI	-0.1055	(0.0613)
Nerv CCS \times neuro	0.2120	(0.2453)
Hospital FEs	Yes	
Observations	1689941	
Pseudo R^2	0.529	

Multinomial logit model of hospital plan choice. Consumers dislike distance and high out-of-pocket prices (copays). Hospital quality is standardized and hospital fixed-effects are included. Standard errors in parentheses. Observations = hospital-admission pairs.

Table 12: Price elasticities from hospital demand model

Statistic	Min	Median	Max
Own elast at h 's own mean copay	-0.145	-0.068	-0.022
Own elast at \$1,000 copay	-0.244	-0.210	-0.041
Cross elast at own mean copays	0.000	0.0001	0.070
Cross elast at \$1,000 copay	0.000	0.0002	0.149

Elasticities of hospital demand with respect to out-of-pocket price. Demand is fairly inelastic. Many pairs of hospitals have cross-price elasticities of essentially zero.

bias the price sensitivity coefficient away from the null. Instead, the control function estimates are consistent with the uninstrumented estimates.⁴⁴ My preferred specification for hospital demand is therefore the conservative and parsimonious model presented in Table 11.

The hospital price elasticities implied by the demand model are summarized in Table 12. The first two rows of the table show hospitals’ own-price elasticities with respect to out-of-pocket prices at each hospital’s observed mean tiered copay and a fixed \$1,000 copay, respectively. Estimates for the own-price elasticity of demand range from -0.02 to -0.24 . At \$1,000 copays, the bulk of the distribution is comparable to the RAND Health Insurance Experiment estimate of approximately -0.2 , where the maximum out-of-pocket price was \$1,000 in 1980s dollars (Manning et al. 1987). However, in the RAND study, elasticities were estimated on the extensive margin of seeking care. My results suggest that demand is also price-responsive on the intensive margin of choosing between options conditional on seeking care in the first place. This provides evidence for the importance of price transparency to controlling moral hazard on the intensive margin as well as the better-studied extensive margin (Pauly 1968). These estimates are comparable to those estimated by Gowrisankaran et al. (2015), who find own-price elasticities of -0.10 to -0.15 . While demand is somewhat responsive to price, the low magnitude of the elasticity estimates indicates that out-of-pocket price differences across tiers must be large in order to appreciably change consumer behavior.

Table 12 also reports hospitals’ pairwise cross-price elasticities. They range from essentially zero to approximately 0.15. The many hospital pairs with negligible cross-price elasticities increase confidence in my approach to reducing dimensionality in the bargaining model by allowing firms to focus primarily on the hospital’s closest competitors in price negotiations (Section 6.5). Hospital pairs that are geographically close have high cross-price elasticities, indicating that they are good substitutes for one another. Many hospitals, including those far from Boston, have a high cross-price elasticity with respect to the top Boston academic medical centers, Brigham and Women’s Hospital and Massachusetts General Hospital. That is, the model predicts that patients substituting away from a given hospital are likely to substitute either to its geographic competitors or to the top hospitals, irrespective of geographic proximity. This accords with intuition and with findings that such “star” hospitals are disproportionately attractive to patients (Ho 2009; Shepard 2014). Furthermore, the cross-price elasticity estimates suggest that these top hospitals, along with other academic medical centers such as Beth Israel Deaconess Medical Center and Tufts Medical Center, are each other’s closest substitutes. These predicted substitution patterns suggest that I am capturing real patterns in how patients choose hospitals.

⁴⁴If anything, estimates with higher-degree polynomials in the control function (degree 3 and higher, columns on the right of the table) suggest that the uninstrumented model may be underestimating the degree of price sensitivity. This is the opposite of what would be expected in the case of endogeneity due to consumer sorting into plans with low copays for their preferred hospitals.

7.2 Willingness to Pay for Hospital Networks

I now turn to the willingness to pay (WTP) for hospital networks implied by the hospital demand estimates, which enters consumers' utility from an insurance plan's hospital network. For a given plan, a household's WTP for the hospital network is defined as the dollarized expected utility of the hospitals, given the household members' risk of diagnosis and the utility from hospitals implied by the hospital demand model (Equation 3 in Section 5). Figure 6a is an illustrative map showing the geographic variation in WTP. The figure plots the median annual WTP by county, in dollars, for the Harvard Pilgrim Independence hospital network in fiscal year 2011. The plan's network includes all hospitals in the state and uses tier copays of \$250, \$500, and \$750, respectively, for hospitals in tiers 1, 2, and 3. The geographic variation in WTP is driven by the fact that some consumers are geographically closer to a larger number of hospitals or more desirable hospitals. This is apparent in the case of the Boston area, visible on the map as the dense cluster of hospitals on the eastern coast of the state. Consumers living in Central and Western Massachusetts who live far from the top hospitals, and in some cases far from any hospital, have a high disutility of distance in the hospital demand model, which implies a low WTP for the hospital network. The geographic gradient of WTP for other broad-network plans' hospital networks follows a similar pattern.⁴⁵ Using a more reduced-form approach to measuring WTP, Ericson and Starc (2014) also find large geographic heterogeneity in WTP for hospital networks in Massachusetts. Their estimates suggest that WTP is substantially higher among residents of the Boston area than those in Springfield (Western Massachusetts) or Worcester (Central Massachusetts), with the geographic variation being comparable to moving a plan's actuarial value from 100% to 0%. These findings, which use a different approach to calculating WTP from this paper, increase confidence in my estimates of geographic variation in WTP.

Figure 6a obscures the non-geographic differences in WTP driven by household size and age-sex composition (appendix Table 22). For adult men in a given geographic area, WTP for the network is monotonically increasing with age. This is driven by the increasing probability of hospital admission as men get older. For adult women, the pattern is similar except for a temporary rise in WTP for women in their 20s and 30s, which is driven by birth-related admissions. Once a woman is past childbearing age, her admission probabilities once again become comparable to those of men in the same age group. Larger households have correspondingly higher WTP.

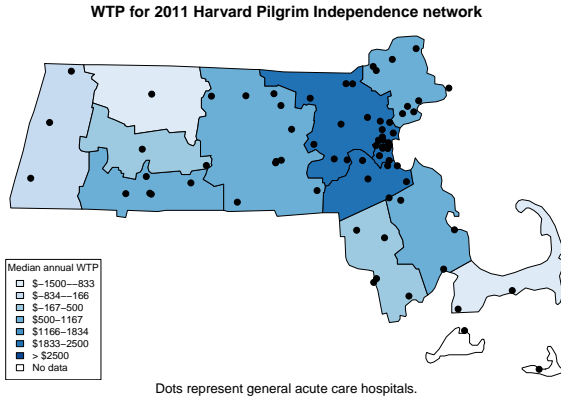
Figure 6b shows differences in WTP across plans within households.⁴⁶ The figure plots the distribution of the difference in annual WTP, in dollars, between a household's observed chosen plan and all other plans. Enrollees in the two largest plans in the GIC are plotted separately from other enrollees. For all plans, the bulk of the distribution is to the right of zero, indicating that for most enrollees, their chosen plan's WTP is greater than the mean WTP of other available plans. This figure provides suggestive evidence that consumers value hospital networks for which

⁴⁵This plot assumes that the differences in coefficient of integration in WTP across households is small in comparison to the primary component of WTP (see discussion in the appendix on page 62).


⁴⁶Comparing WTP within households relaxes the assumption made in interpreting Figure 6a that level shifts in utility across households are small (see discussion in the appendix on page 62).

Figure 6: Variation in hospital network WTP

(a) Median WTP for hospital network, by county



(b) WTP for own vs. other plans



Variation in WTP for hospital networks is shown across consumers by geography (Figure 6a) and across plans within a consumer (Figure 6b). The figures are restricted to single-person households in order to control for differences by household size. The variation in WTP helps to identify the coefficient on network WTP in the plan demand model in Section 7.3.

they have high predicted WTP when choosing plans. Much of this variation is driven by large differences in WTP across narrow-network plans and plans including all hospitals in their network, but there is some variation based on hospital tiers as well. Similar patterns are observed for households with more than one member. These stylized observations motivate the estimation of plan demand in Section 7.3. The median difference between the highest and lowest network WTP in a plan choice set is \$847 for single-person households and \$1,684 for multi-person households; the 75th percentile is \$2,057 and \$4,711, respectively. The substantial variation in WTP across plan networks within a household suggests that WTP for networks could play an important role in plan preferences. By comparison, total employee and employer contributions to premiums on the GIC range from about \$4,000 to \$9,000. The rich variation across geographic and demographic household characteristics that is captured by the WTP measure is used to identify the coefficient on hospital network valuation in the plan demand model, to which I now turn.

7.3 Plan Demand

Estimates of the multinomial logit plan demand model are reported in Table 13. I use the hospital network WTP calculated in Section 7.2 along with longitudinal data on plan availability in each county, individual and family plan premiums, and other plan characteristics from the GIC plan data (see Section 3.2) to estimate the plan demand model. As indicated by the positive and significant coefficient on hospital network WTP, consumers prefer plans whose hospital networks they value more highly. Ericson and Starc (2014) also find that plan choice is highly responsive to network

breadth, but there is no existing evidence on the responsiveness of plan choice to hospitals' tiers within an inclusive network.

The premium variable is the household's out-of-pocket premium contribution, which varies across plans, years, and individual versus family coverage. The negative point estimate on premiums suggests that consumers dislike high premiums, although the coefficient on premiums is noisy once plan fixed effects are included, likely due to noise from measurement error on the portion of the premium paid by the employee.⁴⁷ In specifications with an interaction between premium and household income, price sensitivity is muted by high income. The preferred specification, reported in the second column, includes network WTP and plan fixed effects in addition to premiums and other plan characteristics that vary across plans. The similar magnitude of the coefficients on premiums and WTP suggests that consumers trade off a dollar of premiums and a dollar of WTP roughly equally. To my knowledge, this is the first estimate of consumer preferences over a dollarized measure of WTP, allowing apples-to-apples comparisons of network valuation and other financial characteristics of plans.⁴⁸

Table 13: Plan choice model (GIC FY2009–2011)

	No FEs	+ Plan FEs	Drop Income
Plan Choice			
WTP for hospital network (\$1,000s)	0.84097*** (0.08725)	0.57742*** (0.10639)	0.61425*** (0.10857)
Annual premium (\$1,000s)	-1.92912*** (0.13836)	-0.52966 (0.28780)	-0.33815 (0.27941)
Premium (\$1,000s) \times std. income	0.00003*** (0.00000)	0.00002** (0.00001)	
PCP copay (\$; baseline)	0.48838*** (0.05115)	0.03551 (0.08478)	0.04399 (0.08473)
Uses tiered PCP copays	1.20122*** (0.27283)	0.68900 (0.51400)	0.71385 (0.51354)
Specialist copay (\$; baseline)	-0.18255*** (0.01635)	-0.17430*** (0.03506)	-0.18038*** (0.03500)
Mental health copay (\$)	-0.02409 (0.02223)	-0.15521*** (0.03658)	-0.15569*** (0.03653)
Plan FEs	No	Yes	Yes
Observations	7983	7983	7990
Pseudo R^2	0.178	0.384	0.382

Standard errors in parentheses

All specifications estimated using multinomial logit. N = plan-year-household triplets.

Only first-time GIC enrollees are included. In tiered plans, baseline = tier 1.

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

⁴⁷Some municipalities and other agencies whose employees purchase insurance through the GIC but which are not direct state government employers may charge their employees a portion of the premium different from the standard 20% charged to state employees, and this distinction is not observed in the data.

⁴⁸Ericson and Starc (2014) and Ho and Lee (2015) both estimate plan demand with an ordinal measure of network WTP.

Table 14: Price elasticities from plan demand model, by coverage type

Plan	Individual	Family
Fallon Direct	-0.59	-1.40
Fallon Select	-0.69	-1.65
Harvard Pilgrim Independence	-0.44	-0.97
Harvard Pilgrim Primary Choice	-0.59	-1.41
Health New England	-0.23	-0.50
Neighborhood Health Plan	-0.57	-1.53
Tufts Navigator	-0.60	-1.45
Tufts Spirit	-0.62	-1.51

Elasticities of plan demand with respect to the employee’s premium contribution (usually 25% of the total premium). Insurer-facing elasticities with respect to the total premium are larger in magnitude.

Own-price elasticities of demand with respect to employee premium contributions are reported in Table 14. Plan demand elasticities for high-enrollment plans are in the -0.4 to -0.6 range for individual coverage, and the -1.0 to -1.5 range for family coverage. Family premiums on the GIC are 2.4 times larger than individual premiums for the same plans (see Table 2). These elasticities are in line with previous estimates from the literature. Cutler and Reber (1998) and Royalty and Solomon (1999) find elasticities of -1 and in the range of -0.2 to -0.8 , respectively, with respect to employees’ out-of-pocket premium contributions. Ho and Lee (2015) find elasticities of -1.2 to -1.6 for individual coverage with respect to the total employee and employer premium contribution, which are similar in magnitude to the employee contribution for family premiums in my data on which I find comparable elasticities. The point estimates suggest that consumers respond to premiums when choosing health insurance plans, even in settings such as the GIC where consumers pay a small fraction of the total premium cost.

7.4 Bargaining and Hospital Costs

I use the bargaining model described in Section 6 to infer hospitals’ marginal costs of treating a patient, which are unobserved in the data, from the hospitals’ first-order conditions when negotiating prices with Harvard Pilgrim. The model is solved for the year 2011 using several different parameterizations of the effects of prices on Harvard Pilgrim’s non-GIC plans. Non-GIC plans are assumed to have traditional non-tiered hospital networks, motivated by the fact that Harvard Pilgrim’s GIC plans accounted for 99% of all its tiered enrollment in 2011 (Boros et al. 2014).

The bargaining model is solved for a cluster of hospitals that are each other’s closest competitors but are relatively isolated from competition with other hospitals in Massachusetts. The hospitals—Baystate Medical Center, Cooley Dickinson Hospital, Mercy Medical Center, and Noble Hospital—are clustered in the southwestern part of Massachusetts, in Hampshire and Hampden counties. These hospitals are geographic neighbors and have high cross-price elasticities implied by

the hospital demand model. Intuitively, these four hospitals can be thought of as constituting a distinct submarket of the state’s full hospital market.

Table 15 reports the hospital costs implied by the bargaining model. Costs k_h are the hospital’s marginal cost of treating a patient with a disease severity weight of $l_d = 1$, so that they adjust for differences in patient mix across hospitals. Due to the high computational cost of calculating hospital choice, hospital network WTP, and plan demand for a large number of households, I use a random sample of 1,000 households that are first-time GIC enrollees. Implied hospital costs range from \$5,000 for Noble Hospital, a small community hospital serving a population with a large share of Medicaid and uninsured patients; to \$17,000 for Baystate Medical Center, a large teaching hospital affiliated with Tufts University that houses a physician training (residency) program.

Table 15: Hospital marginal costs k_h from bargaining model (\$)

	k_h , GIC only	k_h , all plans	k_h , all pass thru
Baystate Medical Center	17,171	13,744	15,409
Cooley Dickinson Hospital	17,369	16,035	16,566
Mercy Medical Center	9,800	10,185	10,186
Noble Hospital	4,943	5,098	4,992

Hospital marginal cost per patient k_h implied by the bargaining model, in dollars. Costs reported are per patient with a disease weight of one. The first column reports costs assuming that newly negotiated prices affect only Harvard Pilgrim’s GIC plans; the second column allows prices to affect non-GIC plans, holding those plans’ premiums constant; and the third column allows prices to affect non-GIC plans, assuming that price changes are fully passed through to their enrollees via premiums.

Each column of the table calculates hospital costs with different assumptions on the response of Harvard Pilgrim’s non-GIC plans and enrollees to changes in price. The first column reports costs assuming that the negotiated prices apply only to the GIC enrollees, which is equivalent to assuming that Harvard Pilgrim negotiates a separate price vector for its GIC plans. The final two columns instead allow negotiated prices to apply to all of Harvard Pilgrim’s commercial plans, which is a more accurate approximation. The middle column holds non-GIC plans’ premiums fixed while prices are allowed to adjust; the last column assumes that price changes are fully passed through to non-GIC consumers via changes in premiums, so that price changes affect hospital surplus but not insurer surplus. All reported costs hold non-GIC plan enrollments fixed, since I do not observe plan choice sets or premiums that would allow me to model enrollment responses outside the GIC.

The variation in implied costs across columns in Table 15 highlights the importance of accounting for the effect of price negotiations on the full portfolio of plans to which those prices apply. Due to data limitations, previous papers in the literature have often assumed that negotiated prices would only apply to the subset of the insurer’s population that they could observe, such as narrow-network managed care plans or plans for a single large employer group (Ho and Lee 2015). Comprehensive data on insurance plan choice sets for a broader segment of the privately insured market than is currently available would allow researchers to more reliably estimate price setting in health care. Unfortunately, since I do not observe plan choice sets or premiums for enrollees

outside the GIC, it is difficult to take a stand on which of these sets of assumptions is the most accurate approximation of the non-GIC market. The range of hospital costs produced by the parameterizations should therefore be interpreted as plausible bounds on the true parameter values. In the counterfactual analyses, I use the specification from the middle column that accounts for the effect of prices on non-GIC plans but holds non-GIC premiums fixed.

8 Counterfactuals

The primary question of this paper is how health care demand, prices, and spending respond to tiered provider networks. I therefore conduct counterfactual exercises that examine the impact of tiered networks on the demand side alone, holding prices fixed; and using the full model that allows hospital prices to adjust in new equilibria.

Demand Effects of Tiered Networks: To evaluate the effect of tiered networks on the demand side alone, I compare predicted inpatient hospital spending in Harvard Pilgrim’s largest tiered plan in three scenarios: a non-tiered network using flat hospital copays; a tiered network using copays of \$250, \$500, and \$750 for its three tiers (the actual network); and the same tiered network but with a copay of \$1,500 instead of \$750 for the least preferred tier. The comparison of spending under the observed tiered network to a non-tiered network allows me to evaluate the demand-side effect of moving from a traditional health plan to a tiered hospital network. The third scenario with the \$1,500 least preferred tier (tier 3) copay is motivated by the actual increase of the Harvard Pilgrim GIC plan’s tier 3 copay starting in fiscal year 2016, which was implemented due to a sense that the previous copay differences of \$250 between tiers were not sufficient for steering demand away from the highest-priced hospitals. In these analyses, I hold Harvard Pilgrim’s negotiated hospital prices, hospital tiers, and plan premiums fixed in order to isolate the demand effect of tiered networks from the supply-side response.

Table 16 presents the results of the demand-side counterfactuals. The table presents the distribution of hospital volumes and mean spending per admission across the three scenarios: the baseline non-tiered network, a tiered network with copays of \$250, \$500, and \$750, and the same tiered network but with copays of \$250, \$500, and \$1,500. From left to right, the spread in out-of-pocket price across hospital tiers rises from \$0 to \$1,250.

Hospitals in the more preferred tiers, 1 and 2, gain volume as patients are faced with higher out-of-pocket price spreads with respect to tier 3 hospitals. Total spending per hospital admission falls by 0.7% going from a flat network to a tiered network with a small copay differential across tiers, and by an additional 1.1% moving to the tiered network with the large copay differential across tiers. The total savings gained from moving from a flat network to a tiered network are small: \$45 per patient per year for the network with copays of \$250, \$500, and \$750; and \$112 per patient per year for the network with copays of \$250, \$500, and \$1,500. By comparison, the total annual premium for individual coverage in Harvard Pilgrim’s largest tiered network plan ranges between

Table 16: Demand counterfactuals: demand-side effects of tiered networks

	Flat copay \$250	\$250/\$500/\$750	\$250/\$500/\$1,500
Tier 1 hospitals % of volume	26.36	27.75	29.15
Tier 2 hospitals % of volume	37.49	37.66	39.42
Tier 3 hospitals % of volume	36.14	34.59	31.44
Patient spending per admission (\$)	250.00	517.09	741.51
Δ patient spending over flat copay (%)	0.00	106.84	196.60
Insurer spending per admission (\$)	36125.64	35595.06	34974.11
Δ insurer spending over flat copay (%)	0.00	-1.47	-3.19
Total spending per admission (\$)	36370.88	36101.69	35698.52
Δ total spending over flat copay (%)	0.00	-0.74	-1.85

Demand-side effects of tiered networks, holding prices and enrollments fixed. The first column is the baseline scenario: a traditional hospital network with no tiering and a flat copay across all hospitals. The second column is Harvard Pilgrim’s largest tiered network plan in 2011, with tier copays of \$250, \$500, and \$750 across its three tiers, respectively. The third column is the same as the second except that the copay for tier 3 hospitals is \$1,500.

\$6,000 and \$8,000 throughout the sample period. The incidence of these spending differences is not symmetric across consumers and the insurer. Consumers’ mean out-of-pocket spending rises as copay differentials increase, because demand for the non-preferred tier hospitals remains positive. The welfare consequences therefore depend on how consumers trade off lower premiums against higher and more variable out-of-pocket costs.

Price Effects of Tiered Networks: Tiered networks affect not only the sorting of consumers across hospitals but also the prices insurers negotiate with those hospitals. I now evaluate the effect of tiered networks taking both the demand side and price-setting into account. I consider the supply-side equivalents of the counterfactual scenarios above: non-tiered plans with a flat copay across hospitals (the baseline); and some tiered plans using copays of \$250, \$500, and \$750. I allow hospital demand, prices, tiers, and GIC plan premiums and enrollment to adjust to the new market conditions. In the counterfactual exercise, I fix Harvard Pilgrim’s non-GIC enrollment to be approximately equal to its observed GIC enrollment. Due to the high computational cost of each iteration, I use a random subsample of 100 households from Section 7.4 for the demand portion of the counterfactual calculation.

Table 17 presents the counterfactual prices. I report changes from baseline prices due to data use agreement restrictions on reporting actual prices. Prices are allowed to adjust for the subset of hospitals whose costs are solved from the bargaining model in Section 7.4. The use of tiered hospital networks has a sizable impact on negotiated prices. In a tiered network, hospitals stand to gain volume by agreeing to a lower price that gives them a greater probability of preferred tier status and therefore lower consumer out-of-pocket prices. Under the market conditions in my data, this positive volume effect of lower prices outweighs the negative effect of lower per-patient revenues, resulting in lower average prices when a tiered network is implemented. In moving from a traditional, flat-copay plan to tiered networks with copays of \$250/\$500/\$750, prices are reduced

Table 17: Counterfactuals: price effects of tiered networks

	Baystate	Cooley	Mercy	Noble
No tiers to \$250/\$500/\$750				
% Δ price (hospital)	-8.93	-15.84	-10.21	-17.29
% Δ price (wtd. avg.)	-11.47	-11.47	-11.47	-11.47
Tier (observed)	2	3	1	1
Demand FE	2.06	1.58	0.36	0.13
Hospital share (%)	0.07	0.04	0.01	0.01

Price effects of tiered networks for Harvard Pilgrim, allowing prices, tiers, premiums, and enrollments to adjust. The top panel reports how prices change when moving from the baseline scenario of a non-tiered network with a flat copay to a scenario consistent with Harvard Pilgrim’s largest tiered network plan in 2011, with tier copays of \$250, \$500, and \$750 across its three tiers, respectively. Baystate = Baystate Medical Center. Cooley = Cooley Dickinson Hospital. Mercy = Mercy Medical Center. Noble = Noble Hospital.

by 9% to 17%, with a volume-weighted average reduction of 11% across hospitals.

The magnitudes of these price changes should be interpreted with caution, since this counterfactual exercise approximates the full equilibrium by holding fixed the prices of those hospitals for which I do not have cost estimates from Section 7.4. Allowing all hospitals’ prices to adjust would shift the price-to-tier mapping, potentially affecting the negotiations. The likely effect of allowing full adjustment will be to partially attenuate the price reduction effects implied by my analysis. In addition, “star” hospitals commanding a high preference among consumers may have little to lose in volume from remaining in the least preferred tier, and may therefore be less affected by the downward pressure on prices exerted by tiered networks. The cluster of hospitals for which price changes are calculated in Table 17 form a relatively isolated submarket that responds much more strongly to prices of hospitals within the cluster than to those in the rest of the state. Thus, the directions of the price changes are likely capturing meaningful market dynamics.

The effects of tiered hospital networks on prices are far larger than the effects on demand steering alone, whose small magnitude is driven by the comparatively low price elasticity of demand for hospital care. These results suggest that tiered networks may have large downward effects on hospital prices, even when the demand side effect on individual consumers is small. By aggregating demand responses across many consumers, the effects of demand-side incentives on insurers and hospitals can be substantial.

9 Conclusion

This paper argues that the recent shift toward demand-side incentives in health insurance can indeed reduce health care spending. My findings highlight the importance of considering not only the immediate demand-side effects of these insurance designs, but also their impacts on the supply side. Although tiered networks are moderately successful in steering consumers toward lower-priced care, this steering comes at the expense of higher consumer out-of-pocket spending and muted risk-smoothing. More importantly, price negotiations between insurers and hospitals are responsive to

the aggregate effects of small demand responses by individual consumers. It is these supply-side effects that constitute the bulk of the projected savings from tiered networks. Taken together, the results in this paper provide a cautiously optimistic prognosis for demand-side incentives in health insurance.

To my knowledge, this paper is the first to model and empirically analyze the construction of complex provider networks that go beyond the simple inclusion or exclusion of providers. I build on approaches from the bargaining estimation literature in industrial organization to shed light on issues of interest in health insurance design and the demand for health care. I show that although demand response to price differences across providers is small in an absolute sense, it is sufficiently large to have a material impact on health care spending and prices. Consumers respond to hospital prices both on the margin of hospital choice when they are sick, and on the margin of ex ante health insurance plan choice. These demand responses change the incentive structure on the supply side by giving insurers the additional bargaining lever of tier status in their price negotiations with hospitals. In total, tiered networks have a material effect on the demand for hospital care and the prices paid for that care.

My findings build toward a more complete understanding of the effects of demand-side incentives on health care spending. Extending this work to account for the broader equilibrium impacts of tiered networks on multiple insurers and a larger set of hospitals is a natural direction for future research. More broadly, I provide a new framework that can be used to analyze hospital pricing under various types of complex insurance designs. As the market penetration of such insurance innovations continues to rise, so too will the importance of explicitly accounting for them in analyses of the health care market, especially in applied areas such as antitrust evaluations. The effectiveness of demand-side incentives for reducing health care spending can be improved by careful policy design that accounts for the upstream effects of demand incentives on health care prices.

References

- Abaluck, Jason, Jonathan Gruber, and Ashley Swanson (2015) "Prescription Drug Use under Medicare Part D: A Linear Model of Nonlinear Budget Sets," Working Paper 20976, National Bureau of Economic Research.
- AHRQ, (Agency for Healthcare Research and Quality) (2015) "Clinical Classifications Software (CCS) for ICD-9-CM."
- Bebinger, Martha (2014) "Mass. Patients Can 'Shop' For Health Care – At Least In Theory," *Kaiser Health News*.
- Berry, Steven T. (1994) "Estimating Discrete-Choice Models of Product Differentiation," *The RAND Journal of Economics*, Vol. 25, No. 2, p. 242.
- Boros, Aron, Michael Chernew, Alan Weil, Office of the Massachusetts Attorney General, and Health Policy Commission (2014) "2014 Health Care Cost Trends Hearing," October.
- Buntin, Melinda Beeuwkes, Amelia M. Haviland, Roland McDevitt, and Neeraj Sood (2011) "Healthcare spending and preventive care in high-deductible and consumer-directed health plans," *The American Journal of Managed Care*, Vol. 17, No. 3, pp. 222–230.
- Capps, Cory, David Dranove, and Mark Satterthwaite (2003) "Competition and Market Power in Option Demand Markets," *The RAND Journal of Economics*, Vol. 34, No. 4, p. 737.
- Chandra, Amitabh, Jonathan Gruber, and Robin McKnight (2010) "Patient Cost-Sharing and Hospitalization Offsets in the Elderly," *The American Economic Review*, Vol. 100, No. 1, pp. 193–213.
- CHIA, (Massachusetts Center for Health Information and Analysis) (2013) "Annual Report on the Massachusetts Health Care Market," Technical report, Commonwealth of Massachusetts.
- (2014) "All-Payer Claims Database."
- (2015a) "Massachusetts Hospital Profiles Technical Appendix: Data Through Fiscal Year 2013," Technical report, Commonwealth of Massachusetts.
- (2015b) "2015 Annual Report on the Performance of the Massachusetts Health Care System," Annual Report 15-245-CHIA-01, Commonwealth of Massachusetts.
- Christensen, Hans B., Eric Floyd, and Mark Maffett (2013) "The Effects of Price Transparency Regulation on Prices in the Healthcare Industry," working Paper.
- CMS, (Centers for Medicare & Medicaid Services) (2014a) "Medicare Provider Utilization and Payment Data," April.
- (2014b) "HCAHPS: Patients' Perspectives of Care Survey," September.
- Collard-Wexler, Allan, Gautam Gowrisankaran, and Robin S. Lee (2014) "Bargaining in Bilateral Oligopoly: An Alternating Offers Representation of the "Nash-in-Nash" Solution," Technical report, National Bureau of Economic Research.
- Corlette, Sabrina, Kevin Lucia, and Sandy Ahn (2014) "Implementation of the Affordable Care Act: Cross-Cutting Issues," *Robert Wood Johnson Foundation/Urban Institute*.

- Crawford, Gregory S and Ali Yurukoglu (2012) “The Welfare Effects of Bundling in Multichannel Television Markets,” *American Economic Review*, Vol. 102, No. 2, pp. 643–685.
- Cutler, David M., Mark McClellan, and Joseph P. Newhouse (2000) “How Does Managed Care Do It?” *The RAND Journal of Economics*, Vol. 31, No. 3, p. 526.
- Cutler, David M. and Sarah J. Reber (1998) “Paying for Health Insurance: The Trade-Off between Competition and Adverse Selection,” *The Quarterly Journal of Economics*, Vol. 113, No. 2, pp. 433–466.
- Davis, Steve (2013) “Blues Increasingly Turn to Tiered Networks To Compete in Transparent 2014 Market,” *AISHealth*.
- DHCFP, (Division of Health Care Finance and Policy) (2010) “Massachusetts Health Care Cost Trends 2010 Final Report,” Technical report, Commonwealth of Massachusetts.
- Einav, Liran, Amy Finkelstein, and Paul Schrimpf (2013) “The Response of Drug Expenditures to Non-Linear Contract Design: Evidence from Medicare Part D,” Working Paper 19393, National Bureau of Economic Research.
- Enthoven, A. C. (2014) *Theory and Practice of Managed Competition in Health Care Finance*: Elsevier.
- Ericson, Keith M. Marzilli (2014) “Consumer Inertia and Firm Pricing in the Medicare Part D Prescription Drug Insurance Exchanges,” *American Economic Journal: Economic Policy*, Vol. 6, No. 1, pp. 38–64.
- Ericson, Keith M. Marzilli and Amanda Starc (2015) “Pricing Regulation and Imperfect Competition on the Massachusetts Health Insurance Exchange,” *Review of Economics and Statistics*.
- Ericson, Keith Marzilli and Amanda Starc (2014) “Measuring Consumer Valuation of Limited Provider Networks,” Working Paper 20812, National Bureau of Economic Research.
- Finkelstein, Amy (2002) “The effect of tax subsidies to employer-provided supplementary health insurance: evidence from Canada,” *Journal of Public Economics*, Vol. 84, No. 3, pp. 305–339.
- Frank, Matthew B., John Hsu, Mary Beth Landrum, and Michael E. Chernew (2015) “The Impact of a Tiered Network on Hospital Choice,” *Health Services Research*.
- Fronstin, Paul (2003) “The Impact of a Tiered Network on Hospital Choice,” Issue Brief 260, Employee Benefit Research Institute.
- Gaynor, Martin (2006) “What do we know about competition and quality in health care markets?” Technical report, National Bureau of Economic Research.
- Gaynor, Martin, Kate Ho, and Robert J. Town (2015) “The Industrial Organization of Health-Care Markets,” *Journal of Economic Literature*, Vol. 53, No. 2.
- GIC, (Group Insurance Commission) (2008) “Fiscal Year 2007 Annual Report,” annual Report, Commonwealth of Massachusetts Group Insurance Commission.
- (2009) “Fiscal Year 2008 Annual Report,” annual Report, Commonwealth of Massachusetts Group Insurance Commission.

- (2011) “Fiscal Year 2010 Annual Report,” annual Report, Commonwealth of Massachusetts Group Insurance Commission.
- Gowrisankaran, Gautam, Aviv Nevo, and Robert Town (2015) “Mergers When Prices Are Negotiated: Evidence from the Hospital Industry,” *American Economic Review*, Vol. 105, No. 1, pp. 172–203.
- Grennan, Matthew (2013) “Price Discrimination and Bargaining: Empirical Evidence from Medical Devices,” *American Economic Review*, Vol. 103, No. 1, pp. 145–177.
- Gruber, Jonathan and Robin McKnight (2014) “Controlling Health Care Costs Through Limited Network Insurance Plans: Evidence from Massachusetts State Employees,” Technical report, National Bureau of Economic Research.
- Gruber, Jonathan and James Poterba (1994) “Tax Incentives and the Decision to Purchase Health Insurance: Evidence from the Self-Employed,” *The Quarterly Journal of Economics*, Vol. 109, No. 3, pp. 701–733.
- Handel, Benjamin R (2013) “Adverse Selection and Inertia in Health Insurance Markets: When Nudging Hurts,” *American Economic Review*, Vol. 103, No. 7, pp. 2643–2682.
- Ho, Kate and Robin S. Lee (2013) “Insurer competition and negotiated hospital prices,” Technical report, National Bureau of Economic Research.
- Ho, Kate and Ariel Pakes (2013) “Hospital choices, hospital prices and financial incentives to physicians,” Technical report, National Bureau of Economic Research.
- Ho, Katherine (2006) “The welfare effects of restricted hospital choice in the US medical care market,” *Journal of Applied Econometrics*, Vol. 21, No. 7, pp. 1039–1079.
- (2009) “Insurer-Provider Networks in the Medical Care Market,” *The American Economic Review*, Vol. 99, No. 1, pp. 393–430.
- Ho, Katherine and Robin S. Lee (2015) “Insurer Competition in Health Care Markets,” working Paper.
- Horn, Henrik and Asher Wolinsky (1988) “Bilateral Monopolies and Incentives for Merger,” *The RAND Journal of Economics*, Vol. 19, No. 3, p. 408.
- Kessler, Daniel P. and Mark B. McClellan (2000) “Is Hospital Competition Socially Wasteful?” *The Quarterly Journal of Economics*, Vol. 115, No. 2, pp. 577–615.
- KFF, (Kaiser Family Foundation) (2014) “Employer Health Benefits: 2014 Annual Survey,” Technical report, Kaiser Family Foundation.
- (2015) “2015 Employer Health Benefits Survey,” Technical report, Henry J. Kaiser Family Foundation.
- Kolstad, J. T. and M. E. Chernew (2009) “Quality and Consumer Decision Making in the Market for Health Insurance and Health Care Services,” *Medical Care Research and Review*, Vol. 66, No. 1 suppl, pp. 28S–52S.

- Kolstad, Jonathan T. and Amanda E. Kowalski (2012) “The impact of health care reform on hospital and preventive care: Evidence from Massachusetts,” *Journal of Public Economics*, Vol. 96, No. 11-12, pp. 909–929.
- Kowalski, Amanda E. (2012) “Estimating the Tradeoff Between Risk Protection and Moral Hazard with a Nonlinear Budget Set Model of Health Insurance,” Working Paper 18108, National Bureau of Economic Research.
- Lewis, Matthew and Kevin Pflum (2013) “Hospital systems and bargaining power: evidence from out-of-market acquisitions.”
- Manning, Willard J., Joseph P. Newhouse, Naihua Duan, Emmett B. Keeler, and Arleen Leibowitz (1987) “Health Insurance and the Demand for Medical Care: Evidence from a Randomized Experiment,” *The American Economic Review*, Vol. 77, No. 3, pp. 251–277.
- Martin, Shawn (2014) “We’re Not Gonna Take It: Network Optimization Disrupts Continuity of Care,” *American Academy of Family Physicians*.
- Massachusetts, Commonwealth of (2010) “An Act To Promote Cost Containment, Transparency And Efficiency In The Provision Of Quality Health Insurance For Individuals And Small Businesses.”
- (2012a) “An Act Improving The Quality Of Health Care And Reducing Costs Through Increased Transparency, Efficiency And Innovation.”
- (2012b) “An Act Relative to Tiered and Selective Network Health Plans.”
- McCanne, Don (2013) “Exchange plans have sharply limited networks,” *Physicians for a National Health Program*.
- McKinsey, and Company (2015) “Hospital networks: Evolution of the configurations on the 2015 exchanges,” Technical report, McKinsey Center for U.S. Health System Reform.
- Motheral, Brenda and Kathleen A. Fairman (2001) “Effect of a Three-Tier Prescription Copay on Pharmaceutical and Other Medical Utilization,” *Medical Care*, Vol. 39, No. 12, pp. 1293–1304.
- Pauly, Mark V. (1968) “The economics of moral hazard: comment,” *The American Economic Review*, pp. 531–537.
- Petrin, Amil and Kenneth Train (2010) “A control function approach to endogeneity in consumer choice models,” *Journal of Marketing Research*, Vol. 47, No. 1, pp. 3–13.
- Reinhardt, Uwe E. (2006) “The Pricing of U.S. Hospital Services: Chaos Behind a Veil,” *Health Affairs*, Vol. 25, No. 1, pp. 57–69.
- Robinson, J. C. (2003) “Hospital Tiers In Health Insurance: Balancing Consumer Choice With Financial Motives,” *Health Affairs*.
- Robinson, James C. and Timothy T. Brown (2013) “Increases In Consumer Cost Sharing Redirect Patient Volumes And Reduce Hospital Prices For Orthopedic Surgery,” *Health Affairs*, Vol. 32, No. 8, pp. 1392–1397.
- Royalty, Anne Beeson and Neil Solomon (1999) “Health Plan Choice: Price Elasticities in a Managed Competition Setting,” *The Journal of Human Resources*, Vol. 34, No. 1, p. 1.

- Scanlon, Dennis P., Richard C. Lindrooth, and Jon B. Christianson (2008) "Steering Patients to Safer Hospitals? The Effect of a Tiered Hospital Network on Hospital Admissions," *Health Services Research*, Vol. 43, No. 5p2, pp. 1849–1868.
- Shepard, Mark (2014) "Hospital Network Competition and Adverse Selection: Evidence from the Massachusetts Health Insurance Exchange," working paper.
- Sinaiko, Anna D. (2012) "Tiered networks as strategy to improve health care quality and efficiency," Technical report, National Institute for Health Care Management.
- Sinaiko, Anna D. and Meredith B. Rosenthal (2014) "The Impact of Tiered Physician Networks on Patient Choices," *Health Services Research*, Vol. 49, No. 4, pp. 1348–1363.
- Song, Z., D. G. Safran, B. E. Landon, M. B. Landrum, Y. He, R. E. Mechanic, M. P. Day, and M. E. Chernew (2012) "The 'Alternative Quality Contract,' Based On A Global Budget, Lowered Medical Spending And Improved Quality," *Health Affairs*, Vol. 31, No. 8, pp. 1885–1894.
- Sorensen, Alan T. (2003) "Insurer-Hospital Bargaining: Negotiated Discounts in Post-Deregulation Connecticut," *The Journal of Industrial Economics*, Vol. 51, No. 4, pp. 469–490.
- Starc, Amanda (2014) "Insurer pricing and consumer welfare: Evidence from medigap," *The RAND Journal of Economics*, Vol. 45, No. 1, pp. 198–220.
- Stremikis, Kristof, Karen Davis, and Stuart Guterman (2010) "Health Care Opinion Leaders' Views on Transparency and Pricing," Data Brief 102, The Commonwealth Fund.
- Terza, Joseph V., Anirban Basu, and Paul J. Rathouz (2008) "Two-stage residual inclusion estimation: Addressing endogeneity in health econometric modeling," *Journal of Health Economics*, Vol. 27, No. 3, pp. 531–543.
- Town, Robert and Gregory Vistnes (2001) "Hospital competition in HMO networks," *Journal of Health Economics*, Vol. 20, No. 5, pp. 733–753.
- Train, Kenneth (2002) *Discrete Choice Methods with Simulation*: Cambridge University Press.
- Trish, Erin E. and Bradley J. Herring (2015) "How Do Health Insurer Market Concentration and Bargaining Power with Hospitals Affect Health Insurance Premiums?" *Journal of Health Economics*.
- Trivedi, Amal N., Husein Moloo, and Vincent Mor (2010) "Increased ambulatory care copayments and hospitalizations among the elderly," *New England Journal of Medicine*, Vol. 362, No. 4, pp. 320–328.
- Weisman, Robert (2011) "Blue Cross and Tufts Medical set contract," *The Boston Globe*.
- Weisman, Robert and Liz Kowalczyk (2011) "Tufts, Blue Cross contract row threatens members."
- White, Chapin, Paul B. Ginsburg, Ha T. Tu, James D. Reschovsky, Joseph M. Smith, and Kristie Liao (2014) "Healthcare Price Transparency: Policy Approaches and Estimated Impacts on Spending," policy Analysis, West Health Policy Center.
- Wooldridge, Jeffrey M. (2010) *Econometric Analysis of Cross Section and Panel Data*: MIT Press.

- Wrobel, Marian V., David Auerbach, and Sara Sadownik (2014) “2014 Cost Trends Report,” annual Report, Commonwealth of Massachusetts Health Policy Commission.
- Wu, Sze-jung, Gosia Sylwestrzak, Christiane Shah, and Andrea DeVries (2014) “Price Transparency For MRIs Increased Use Of Less Costly Providers And Triggered Provider Competition,” *Health Affairs*, Vol. 33, No. 8, pp. 1391–1398.
- Yegian, J. M. (2003) “Tiered Hospital Networks,” *Health Affairs*, pp. W3–147–153.
- Yong, Pierre L., LeighAnne Olsen, and J. Michael McGinnis (2010) “Approaches to Improving Value—Consumer Incentives,” in *Institute of Medicine (US) Roundtable on Value & Science-Driven Health*.

Appendices

Appendix A: List of Notation

Symbols used throughout the paper are listed in the following table in alphabetical order. Greek letters are included at their analogous place in the English alphabet, e.g. δ is included with entries for the letter D. The table continues on the next page.

Symbol	Description
α	Demand coefficient on hospital out-of-pocket price
β	Demand coefficient vector on hospital and patient characteristics
$b_{\mathcal{M}(h)}, b_{h(\mathcal{M})}$	Insurer \mathcal{M} 's and hospital h 's Nash bargaining weights w.r.t each other
c_{mh}	Plan m 's copay for hospital h
C_{mhi}	Consumer i 's total expected out-of-pocket payments to hospital h under plan m
d	Index of diagnosis categories, $d \in D$
δ_1	Demand coefficient on plan premium
δ_2	Demand coefficient on WTP for plan's hospital network
ε_{mhid}	Error term in hospital demand model
f_{id}	Consumer i 's probability of contracting diagnosis d
$G_{\mathcal{M}}^t(p_{\mathcal{M}h})$	Probability that hospital with price $p_{\mathcal{M}h}$ is in tier t of insurer \mathcal{M} 's network
$g_{\mathcal{M}}^t(p_{\mathcal{M}h})$	Derivative of $G_{\mathcal{M}}^t(p_{\mathcal{M}h})$ with respect to price
$\mathcal{G}_{\mathcal{M}h}^\tau$	Product of of probabilities of network tier permutation τ for all hospitals except h
γ	Demand coefficient vector on plan characteristics
i	Index of individual consumers
ι	Index of households, $\iota \in I$
h, j	Index of hospitals, $h, j \in H$
k_h	Hospital h 's baseline marginal cost for a patient with diagnosis weight $l_d = 1$
l_d	Diagnosis d 's multiplier for scaling price and hospital cost
λ	Medical loss ratio (MLR)
\mathcal{M}, \mathcal{N}	Index of insurers, $\mathcal{M}, \mathcal{N} \in M$
m, n	Index of insurance plans, $m \in \mathcal{M}, n \in \mathcal{N}$
$p_{\mathcal{M}h}$	Base price negotiated between insurer \mathcal{M} and hospital h
$\Pi_{\mathcal{M}}$	Insurer \mathcal{M} 's expected profit
Π_h	Hospital h 's expected profit
$r_{m\iota}$	Plan m 's premium for household of size ι
$s_{m\iota}$	Household ι 's probability of enrolling in plan m
$S_{\mathcal{M}(h)}$	Insurer \mathcal{M} 's expected surplus in negotiations with hospital h
$S_{h(\mathcal{M})}$	Hospital h 's expected surplus in negotiations with insurer \mathcal{M}
σ_{mhid}	Consumer i 's probability of choosing hospital h in plan m with diagnosis d

t	Indexes individual hospitals' network tiers, $t \in \{1, 2, 3\}$
τ	Indexes permutations of all hospitals' network tiers, $\tau \in T$
u_{mhid}	Consumer i 's utility from hospital h in plan m with diagnosis d
$U_{m\iota}$	Household ι 's utility from plan m
V_{mhi}	Hospital h 's total expected casemix-adjusted volume for consumer i in plan m
$W_{mi}, W_{m\iota}$	Consumer i 's and household ι 's WTP for plan m 's hospital network
x_{hid}	Hospital and patient characteristics in hospital demand model
X_m	Plan m 's characteristics in plan demand model
$\zeta_{m\iota}$	Error in plan demand model

Appendix B: Additional Tables and Figures

Figure 7: Massachusetts' insurers hospital tiers (2012)

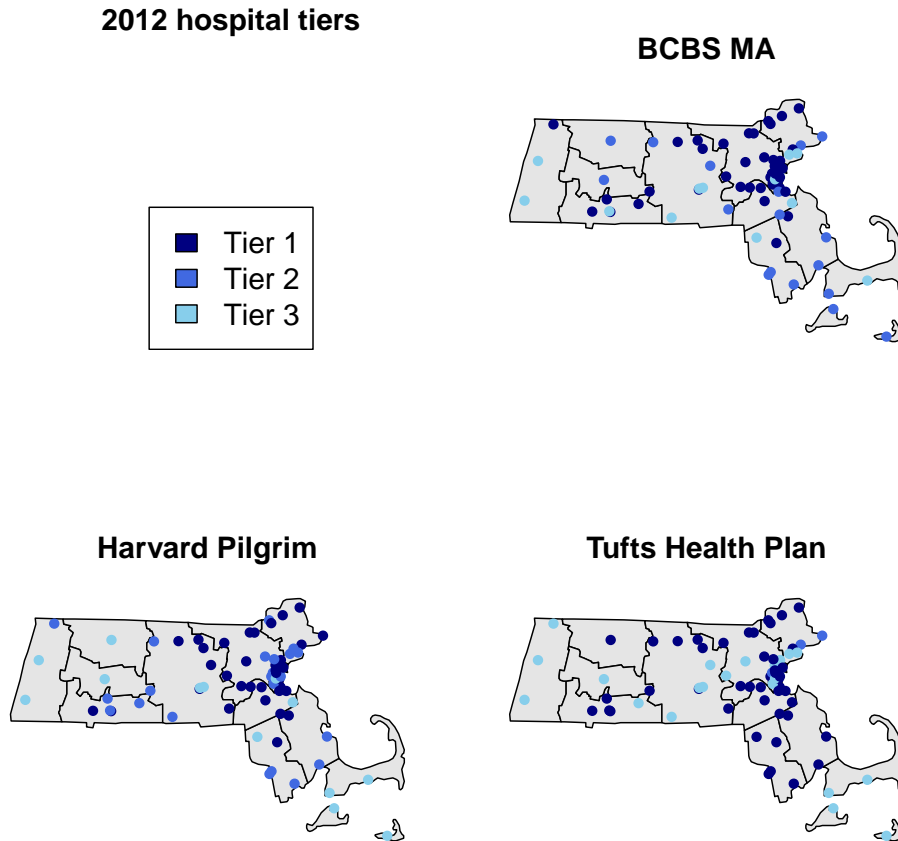


Table 19: Enrollment in GIC plans

Plan	Share (%)	New policies	New enrollees	2009-2012 enrollt.
Fallon Direct	1.52	891	1,543	7,177
Fallon Select	3.78	1,286	2,684	11,167
Harvard Pilgrim Independence	36.42	16,358	36,444	96,103
Harvard Pilgrim Primary Choice	3.04	2,079	4,472	22,208
Health New England	9.54	3,443	6,451	29,312
Neighborhood Health Plan	1.71	924	1,645	7,552
Tufts Navigator	41.97	10,137	20,438	120,519
Tufts Spirit	1.16	1,228	2,577	13,775
UniCare Basic				
UniCare Community Choice				
UniCare PLUS				

GIC plan enrollment for employees and their dependents, excluding UniCare plans.

Share is market share is at the end of fiscal year 2011 (June 2011).

Enrollee and policy holder counts are for first-time GIC enrollees in 2009–June 2011.

Final column is total number of unique enrollees in 2009–2012.

Appendix C: Level Shifts in WTP and Identification

This section discusses the identification of WTP for hospital networks from Section 5.3 within versus across individuals. The consumer surplus in discrete choice models is generally identified only up to a level shift, denoted here by C (Train 2002). Thus, differences in willingness to pay across networks or plans are meaningful while the absolute level is not. In the context of willingness-to-pay for a product whose probabilities of any consumption vary across individuals, as with hospital care, this property has important implications for comparing consumer surplus across individuals. Consider a constant level shift v affecting all exponentiated terms for all individuals. This implies a shifted willingness-to-pay of

$$\begin{aligned}
W_{mi}(v) &= \frac{1}{\alpha} \sum_{d \in D} f_{id} \ln \left(\sum_{h \in H} \exp(\alpha c_{mh} + \beta x_{hid} + v) \right) + C \\
&= \frac{1}{\alpha} \sum_{d \in D} f_{id} \ln \left(H \exp(v) \cdot \sum_{h \in H} \exp(\alpha c_{mh} + \beta x_{hid}) \right) + C \\
&= W_{mi} + C_i(v)
\end{aligned}$$

where $C_i(v) = \frac{1}{\alpha} \sum_{d \in D} f_{id} (\ln(H) + v)$ is a shift that may be heterogeneous across individuals. Individuals or households with lower price sensitivity α and those with higher probabilities of hospital admission $\sum_{d \in D} f_{id}$ will experience a greater shift in willingness-to-pay. This has the effect that the level shift of willingness-to-pay between individuals will vary depending on, say, which hospital is designated as the baseline in a model with hospital fixed effects. However, the

Table 20: Hospital choice model

	(1) No FEs	(2) Hospital FEs
HospitalChoice		
Copay (income Q1)	0.00068*** (0.00013)	-0.00032* (0.00015)
Copay (income Q2-4)	0.00082*** (0.00005)	-0.00025*** (0.00006)
Copay (income Q5)	0.00070*** (0.00008)	-0.00022* (0.00009)
Distance (mi)	-0.20096*** (0.00166)	-0.20003*** (0.00187)
Distance ²	0.00053*** (0.00001)	0.00071*** (0.00001)
Age \times dist	-0.00006* (0.00003)	-0.00010*** (0.00003)
Male \times dist	0.00284** (0.00099)	0.00066 (0.00102)
Chronic cond \times dist	0.01979*** (0.00103)	0.02047*** (0.00108)
Teaching \times dist	0.01754*** (0.00107)	-0.00588*** (0.00138)
Beds \times dist	0.00006*** (0.00000)	0.00004*** (0.00000)
Satellite hosp campus	-0.29251*** (0.02148)	18.35662 (1753.65494)
Cardiac CCS \times cath lab	1.15067*** (0.09551)	0.67128*** (0.09953)
Obstetric CCS \times NICU	0.85982*** (0.03463)	0.32700*** (0.03874)
Nerv, circ, musc CCS \times MRI	0.03747 (0.04932)	-0.10803 (0.06137)
Nerv CCS \times neuro	1.80914*** (0.22686)	0.20747 (0.24528)
% good pain control \times dist	-0.00227*** (0.00054)	-0.00536*** (0.00060)
% highly recommend \times dist	0.01028*** (0.00045)	0.00471*** (0.00057)
Hospital FEs	No	Yes
Observations	1687820	1687820
Pseudo R^2	0.457	0.529

Standard errors in parentheses

N = number of admission-hospital pairs. All specifications estimated using multinomial logit.

Hospital quality variables are standardized.

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table 21: Hospital choice model (with control function)

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	Pref. spec	+IV sample	IV deg1	IV deg2	IV deg3	IV deg4	IV deg5
Hospital Choice							
Copay (dollars)	-0.0002*** (0.0001)	-0.0002* (0.0001)	-0.0002* (0.0001)	-0.0002* (0.0001)	-0.0009*** (0.0002)	-0.0010*** (0.0002)	-0.0009** (0.0003)
Distance (mi)	-0.1998*** (0.0019)	-0.1986*** (0.0023)	-0.1993*** (0.0026)	-0.1993*** (0.0027)	-0.2007*** (0.0027)	-0.2009*** (0.0027)	-0.2008*** (0.0027)
Distance ²	0.0007*** (0.0000)	0.0007*** (0.0000)	0.0007*** (0.0000)	0.0007*** (0.0000)	0.0007*** (0.0000)	0.0007*** (0.0000)	0.0007*** (0.0000)
Hospital FEs	Yes	Yes	Yes	Yes	Yes	Yes	Yes
CF degree 1	No	No	Yes	Yes	Yes	Yes	Yes
CF degree 2	No	No	No	Yes	Yes	Yes	Yes
CF degree 3	No	No	No	No	Yes	Yes	Yes
CF degree 4	No	No	No	No	No	Yes	Yes
CF degree 5	No	No	No	No	No	No	Yes
Observations	1689941	1107964	1107467	1107467	1107467	1107467	1107467
Pseudo R^2	0.529	0.538	0.538	0.538	0.538	0.538	0.538

Standard errors in parentheses

N = number of admission-hospital pairs.

All specifications estimated using multinomial logit.

IV columns reported with bootstrapped standard errors with 100 replications.

Hospital quality variables are standardized.

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table 22: Household WTP for plan networks

	(1)
	WTP
Household size	703.26*** (126.07)
Youngest household member age	8.77 (4.97)
Oldest member age, if household size > 1	6.73 (5.12)
Narrow network \times household size	-397.72*** (29.31)
Tiered network \times household size	89.92*** (17.49)
Narrow tiered network \times household size	-203.79** (65.16)
Observations	8100
Pseudo R^2	

Estimated using OLS. N = plan-year-household triplets.

Standard errors (in parentheses) are clustered by household.

shift $C_i(v)$ is constant for all networks within an individual or household. The variation generated by predicted willingness-to-pay therefore remains informative for estimating plan choice models, which are identified from differences across choice alternatives within rather than across households. While the heterogeneous shifts across individuals are intuitive in the case of a hospital network, this property will also hold for any other discrete choice model where demand coefficients or probabilities of product purchase differ across individuals.