Regulating Out-of-Network Hospital Payments: Disagreement Payoffs, Negotiated Prices, and Access*

Elena Prager
Simon Business School
University of Rochester
and NBER

Nicholas Tilipman
Carey Business School
Johns Hopkins University

April 2, 2025

Abstract

Recent policy proposals seek to regulate out-of-network hospital prices. We study how such regulation affects equilibrium prices, network formation, and hospital exit. We estimate a structural model of insurer-hospital bargaining that allows for out-of-network transactions between non-contracting parties. These transactions generate a notion of exit by rendering hospitals unprofitable under some regulations. Estimation relies on a novel measure of out-of-network prices. We find that reducing out-of-network prices would also lower negotiated prices, but potentially at the cost of narrower hospital networks. Aggressive regulation could induce substantial hospital exit, but only under the restrictive assumption that negotiators cannot anticipate the exits.

JEL Classification Codes: C78, I11, L13

*Previously circulated under the title "Disagreement Payoffs and Negotiated Prices: Evidence from Out-of-Network Hospital Payments." We thank Pierre Dubois, Chris Garmon, Kate Ho, Tom Hubbard, Robin Lee, Devesh Raval, Mark Shepard, and numerous seminar and conference participants for helpful suggestions. Any errors are our own.
Contact information: Prager: eprager@simon.rochester.edu; Tilipman: tilipman@jhu.edu

1 Introduction

Equilibrium outcomes of bilateral negotiations depend critically on disagreement payoffs—the payoffs parties receive if negotiations fail. In some contexts, failure to reach agreement will completely eliminate transactions between the disagreeing parties (Grennan 2013; Crawford and Yurukoglu 2012). However, there are contexts where some economic activity can still take place in the absence of a successfully negotiated contract, including online retail platforms, labor unions, and health care.[1] Such off-contract transactions alter the negotiating parties' disagreement payoffs, which in turn determine equilibrium prices and even whether an agreement is reached. Despite this crucial role of off-contract transactions in Nash-in-Nash bargaining, accounting for them empirically is challenging, partly because off-contract transactions are rarely observable in data.

In this paper, we study how off-contract transactions affect equilibrium outcomes in health care markets, where such transactions are both common and increasingly subject to regulation. When negotiations between an insurer and a hospital fail, patients can still seek "out-of-network" care from that hospital. Insurers then often reimburse the hospital an "out-of-network price" far above the price for equivalent in-network care. The magnitude of these payments is not trivial: in our setting, out-of-network services represent 10 percent of hospital volumes for insurers with incomplete networks, with prices 55 percent above in-network prices. Policy-makers recognize that these payments increase hospitals' bargaining leverage, and recent bipartisan proposals seek to regulate out-of-network prices.[2] Proponents argue that this would both directly reduce costs for patients receiving out-of-network care and, through the bargaining channel, reduce equilibrium in-network prices as well (Kane 2019; Chernew et al. 2019). Opponents warn of unintended consequences for network breadth and reduced access to care. Medical provider industry associations argue that some types of aggressive price regulation could even force providers to exit (Jacquis 2019; Smith and Finder 2022; Wesolowski and Finder 2022; Henry 2025).

Our aim is to evaluate empirically how regulating out-of-network payments would affect this price-access tradeoff. We make three main contributions. First, we propose a practical solution to the empirical challenge of measuring out-of-network prices actually paid to hospitals by health insurers. While the literature recognizes the importance of correctly accounting for disagreement values in the estimation of bargaining models (Ho and Lee 2019), existing papers on hospital-insurer bargaining have assumed away the presence of transactions in the case of disagreement, owing in

[1] For example, when brands leave two-sided platforms for online retail, consumers can still purchase their products through third-party sellers. Consider the case of Nike withdrawing from its contract with Amazon in fall 2019, following Nike's dissatisfaction with Amazon's handling of counterfeit and third-party merchandise (Hanbury 2019). Consumers could still purchase Nike products through third-party sellers on Amazon. In the context of labor unions, stalled negotiations do not always result in a strike that eliminates employer-worker transactions, even temporarily. Instead, the disagreement outcome may be a work-to-rule, in which workers continue to produce output, but at a substantially reduced level, and the employer continues to pay them.

[2] Many other countries directly regulate health care prices for services delivered within contracts. However, such direct price regulation is likely politically infeasible in the U.S.

part to the difficulty of measuring off-contract prices. U.S. health care markets essentially lack posted prices, especially for out-of-network care (Chernew et al. 2022). Out-of-network prices are difficult to infer because they can vary by insurer, geography, type of service, and institutional features or laws governing a particular market. To circumvent these issues, we leverage the institutional details of insurers' out-of-network payment policies. These policies typically rely on third-party standardized price benchmarks constructed from hospital charge prices in a given geographic market. We replicate these price benchmarks using widely-available claims data and then validate them using hand-collected data on insurers' networks. The result is, to our knowledge, the first data-driven measure of out-of-network prices that closely matches observed payments. Without this measure, we could not conduct policy evaluations of proposals to change out-of-network prices from their current levels.

Our second contribution is to develop and estimate a structural model of insurer-provider bargaining that permits the notion of provider exit. We achieve this by extending the standard hospital-insurer bargaining framework to allow for strictly positive patient volumes and prices even when a hospital is out of network. The bulk of existing work defines the disagreement outcome of a negotiation as severing that pair's link outright, implying zero transactions between non-contracting parties (Crawford and Yurukoglu 2012; Ho and Lee 2017; Gowrisankaran et al. 2015; Prager 2016).[3,4] This typical setup implies that a hospital can reject any contract that offers a price below marginal cost, and can therefore never be unprofitable. Exit can only occur if every insurer is unwilling to pay a price above cost, a rare event under reasonable parameter values. By contrast, our framework allows for the possibility of negative disagreement payoffs when (regulated) prices are sufficiently low. This generates a hospital shutdown condition: if a hospital's total variable profit across insurers becomes negative, it can endogenously choose to exit an unprofitable service line even if some insurers are willing to pay above-cost prices. This feature allows us to take seriously the claims of industry groups that price regulation could force closures by causing prices to fall below costs (Jacquis 2019; Smith and Finder 2022; Wesolowski and Finder 2022; Henry 2025). To our knowledge, this paper is the first to empirically evaluate claims that health care price regulation will induce exits, despite a long history of such claims. We show the conditions for closure under two sets of assumptions: "naive" and "sophisticated" negotiators, where the latter anticipate potential hospital closures and amend their disagreement payoffs accordingly during negotiation. Under reasonable parameter values, when negotiators are sophisticated, insurers are willing to pay all hospitals above-cost prices to ensure their continued operation, even at aggressive levels of regulation.

Our third contribution is to provide the first comprehensive empirical analysis of how regulating

[3]This is the Nash-in-Nash structure introduced by Horn and Wolinsky (1988).

[4]Papers that allow for more than a single deviation from the observed equilibrium, such as those using a Nash-in-Nash model with threat of replacement (Ho and Lee 2019; Ghili 2022), define the surplus from agreement more flexibly. However, those papers maintain the assumption of zero off-contract transactions.

out-of-network payments affects market outcomes. We first show that the assumption of zero disagreement payoffs significantly biases estimates of hospital marginal costs upwards (and hence estimated hospital profits downward). In our setting, this bias amounts to 20 percent on average. This has direct implications for regulation: standard models understate hospitals' ability to absorb price reductions.[5]

We operationalize our empirical analysis in the New Hampshire hospital market. This is a suitable setting for several reasons. First, New Hampshire has insurer-hospital pairs that truly lack contracts. In some markets, insurers that offer narrow-network health maintenance organization (HMO) plans still have negotiated contracts with every hospital because they use complete networks for their preferred provider organization (PPO) plans. This paper studies regulation of cases where there exists no negotiated price to fall back on, and so requires a setting with true narrow networks, where some hospitals are out of network for *all* of an insurer's plans. In the New Hampshire market, multiple insurers have true narrow networks. These are regional insurers with a large presence in neighboring states and low enrollment in New Hampshire itself. Second, we document nontrivial volumes at out-of-network New Hampshire providers. In our sample, out-of-network hospitals account for 10 percent of hospital transactions across two insurers with true narrow networks, allowing us to validate our proposed measure of out-of-network prices.

We next simulate counterfactuals that mimic proposed regulations from both major political parties.[6] Our first set of counterfactuals considers capping out-of-network reimbursements at multiples of Medicare prices. The second set of counterfactuals varies the charge price benchmarks from which most insurers in our sample determine their current out-of-network payments. We consider policies that reduce the price benchmarks and policies that raise them to the point where hospitals are paid nearly their full charge price. In counterfactuals, increasing the off-contract prices gives hospitals bargaining leverage to negotiate higher prices. Specifically, increasing the current off-contract price benchmark percent results in up to a 7 percent increase in average volume-weighted in-network prices. The increase flattens at about 1.2 times the current rate, suggesting that hospitals are paid close to their list prices at baseline. Conversely, reducing off-contract prices by half substantially reduces negotiated prices, by 14 to 29 percent depending on the insurer.

However, while capping out-of-network reimbursements reduces equilibrium prices, we show that it also imposes a trade-off against reduced access to providers. In equilibrium, lower out-

[5] Relatedly, out-of-network payments are a subject of antitrust cases against hospitals. California's antitrust complaint against Sutter Health describes Sutter's out-of-network prices as "punitively high" (Becerra et al. 2018; Ellison 2018). Another high-profile example involves insurers in New Jersey citing high out-of-network reimbursements as responsible for rapid premium growth in the state (Avalere 2015). Similarly, New Jersey's Bayonne Medical Center was accused of strategically going out of network with insurers in order to receive higher reimbursements (Creswell et al. 2013).

[6] The Republican-sponsored Lower Health Care Costs Act of 2019 proposes capping insurers' off-contract payments at median in-network prices within each market (Alexander 2019). During his campaign for the 2020 Democratic presidential nomination, Pete Buttigieg proposed setting the cap at 200 percent of Medicare (Pete For America 2019). Other third-party proposals have called for caps as low as 120 percent of Medicare (Kane 2019).

of-network prices reduce the joint surplus from agreement, a result of a reduction in the savings available to the insurer from bringing a given hospital into its network. In our counterfactuals, cutting off-contract prices to the vicinity of 280 percent of Medicare (from an observed baseline of 320 to 350 percent of Medicare, depending on the insurer) yields significant price reductions with minimal disruptions in networks for most insurers. However, more aggressive regulation leads to sharp declines in network breath. For example, reducing out-of-network reimbursements by half reduces the share of hospitals covered by about 40 percentage points for some insurers.

Finally, our counterfactual simulations suggest that industry concerns that regulation will drive hospital exit are valid, but only under strong assumptions about forward-lookingness. When negotiators are "naive" and fail to anticipate that hospitals with below-cost reimbursements will be induced to exit, aggressive price regulation leads to substantial "surprise" exits. With out-of-network reimbursements pegged to 100 percent of Medicare, the counterfactuals predict that half of hospitals exit our in-sample service lines. A disproportionate fraction of the exiters are small, rural hospitals, raising concerns about impacts on access to care. However, exits completely disappear when we model negotiators as "sophisticated." When insurers anticipate that hospitals earning below-cost prices will exit, they agree to pay higher prices that keep hospitals open. As a result, we find that regulating out-of-network reimbursements to Medicare rates not only avoids hospital exits, but in certain cases *increases* network breadth. This occurs because with sufficiently low out-of-network prices, in-network contracts with above-cost prices are the only mechanism for insurers to ensure a hospital does not exit.

Our paper relates to several strands of literature. Several recent papers have proposed approaches to relaxing the standard assumption that in case of disagreement, all other parties' contracts remain the same (Ho and Lee 2019; Ghili 2022; Liebman 2022). We view our approach as complementing these important advances by providing a computationally simple alternative for dealing with misspecified disagreement values. Another strand of the literature has recently begun investigating the prevalence and impact of out-of-network reimbursement structures and other determinants of insurer-hospital negotiated prices, especially in the context of surprise out-of-network bills (Cooper et al. 2019a; Craig et al. 2021; Cooper et al. 2019b, 2020; Fiedler 2020). We contribute to this literature by formally incorporating out-of-network reimbursements into a model designed to predict their impact on in-network prices, network breadth, and hospital exits.

The paper proceeds as follows. Section 2 describes the types of health care that are commonly received out of network, explains our algorithm for measuring out-of-network prices, and describes our data. Section 3 presents our theoretical model and empirical strategy. Section 4 presents the parameter estimates, and Section 5 presents counterfactual simulations. Finally, Section 6 concludes.

2 Out-of-Network Care, Prices, and Data

In this section, we first provide context for our empirical application: the private health insurance market in New Hampshire. We then discuss the types of care that are frequently received out of network and describe our measure of out-of-network prices. Finally, we outline our data and sample used for estimation.

2.1 Empirical Setting

Our empirical setting is insurers' negotiations with hospitals in New Hampshire. The state's health insurance market is highly concentrated, with the largest three insurers accounting for at least 85 percent of commercial enrollment. Two of the top three insurers are large national insurers, Anthem (Blue Cross Blue Shield branch) and Cigna. The third is Harvard Pilgrim Health Care, a smaller, regional insurer that draws the bulk of its enrollment from New England (Prager and Tilipman 2019). The remainder of the insurance market is divided between other national insurers, notably Aetna and United; and regional insurers, notably Tufts Health Plan and MVP Health Care. Our sample consists of the top six insurers in the state by market share (Anthem, Harvard Pilgrim, Cigna, United, Aetna, and MVP), plus an additional regional insurer (Tufts) that provides additional variation in hospital networks.

New Hampshire has 26 general acute care hospitals, including the state's premier academic hospital, Dartmouth-Hitchcock Medical Center. With more than a third of its population classified as rural, and mountainous terrain that impedes travel, fully half of New Hampshire's hospitals are designated as Critical Access Hospitals (CAH) by the Center for Medicare and Medicaid Services (CMS). Because New Hampshire is geographically small and shares a relatively densely populated border with Massachusetts, many hospitals in the southern part of the state have substantial patient volumes from Massachusetts residents or New Hampshire locals who are insured by Massachusetts insurers.

Most insurers with substantial operations in New Hampshire have complete hospital networks within the state, meaning they have negotiated contracts with each of the state's 26 acute care hospitals. This includes the state's top three insurers. This pattern is not peculiar to New Hampshire; it is common for insurers to have locally complete hospital networks for at least some of their plans. However, MVP and Tufts have incomplete hospital networks in New Hampshire. We next describe these incomplete networks and how patients use out-of-network care.

2.2 Utilization of Out-of-Network Services

Patients enrolled in plans with incomplete networks must choose between receiving care at an in-network hospital or going to an out-of-network hospital, typically at a higher out-of-pocket cost. The list of in-network hospitals is sometimes quite limited. For example, Massachusetts-based Tufts

Health Plan, which is among the smaller insurers in New Hampshire, has negotiated contracts with only eight of the state's 26 hospitals (Figure A.2b). Similarly, MVP, which operates primarily in upstate New York and Vermont, also has contracts with only eight New Hampshire hospitals (Figure A.2c). We observe these insurers' patients at least occasionally receiving care at every hospital in the state. Some of this care is emergency care or other care that was not explicitly chosen by the patient, which often results in a "surprise" out-of-network medical bill (Cooper et al. 2020; Biener et al. 2021). However, we document that a substantial portion of out-of-network care is non-emergent, voluntary care (Section 2.4).

The non-emergent out-of-network care follows some systematic patterns. Table 1 reports the most common diagnoses treated out of network. Panel A reports the top ten most frequent diagnoses, using raw counts. These are mechanically dominated by diagnoses that are common across all patients, such as routine medical exams. However, the conditional probability of receiving care out of network given a diagnosis paints a different picture. Table 1 Panel B reports the top ten diagnoses by share of care obtained out of network. This list is dominated by cancer-related diagnoses (neoplasms and dysplasias), including beyond the top ten displayed in the table; and pregnancy-related care. That is, while the raw majority of out-of-network care is routine medical care, out-of-network care is *disproportionately* used in serious and complicated cases.

The patients receiving out-of-network care also differ from other patients. Table 2 summarizes the differences in patient population among insurers with complete networks (Column 1), patients receiving in-network care among insurers with incomplete networks (Column 2), and out-of-network care among insurers with incomplete networks (Column 3). Consistent with Table 1, out-of-network care is on average more complex (a typical case is about double the severity weight) and obtained by younger patients. These patients are costlier to treat, with the cost difference being a function of observable differences in severity. Out-of-network patients are disproportionately female, driven largely by the disproportionate share of pregnancy-related care.

Patients also disproportionately choose hospitals with certain characteristics for their out-of-network care. These hospitals tend to be relatively prestigious and high-volume. For example, Dartmouth-Hitchcock, a premier academic medical center, and Concord hospital, which is affiliated with the Dartmouth Geisel School of Medicine, both command large volumes from patients enrolled in full-network insurers. These same hospitals also have a high share of volume when they are out-of-network, in spite of the fact that many patients live far away (Figure A.2).

2.3 Measuring Out-of-Network Prices

When a patient receives out-of-network care, there is no contract governing the price her insurer will pay to the provider. There is also no meaningful posted price that can act as the default off-contract price (Reinhardt 2006). As a result, insurers typically put in place explicit policies governing how much they will pay non-contracted (out-of-network) providers. While insurers could

Table 1: Most Common Out-of-Network Diagnoses

Panel A: Top 10 by Volume	Code	Volume	OON/IN Ratio
Care involving other physical therapy	V571	2,129	9.39
Other screening mammogram	V7612	1,973	0.90
Routine general medical examination	V700	1,562	1.21
Chest pain, unspecified	78650	707	1.00
Encounter for occupational therapy	V5721	499	20.51
Unspecified essential hypertension	4019	487	1.08
Routine infant or child health check	V202	486	1.99
Pain in limb	7295	435	1.01
Pain in joint, lower leg	71946	420	0.77
Other chest pain	78659	414	1.09
Panel B: Top 10 by OON/IN Ratio			
	Code	Volume	OON/IN Ratio
Malignant neoplasm of upper-inner quadrant of female breast	1742	125	184.97
Persistent disorder of initiating or maintaining sleep	30742	24	42.62
Examination of eyes and vision	V720	77	37.98
Dysplasia of prostate	6023	22	32.55
Spontaneous abortion, complicated by hemorrhage	63411	35	31.07
Unspecified antepartum hemorrhage	64190	24	30.44
Chronic ulcer of other specified sites	7078	49	29.00
Urethral diverticulum	5992	16	28.41
Nontraumatic rupture of foot and ankle tendons	72768	55	24.42
Unspecified disorder of teeth and supporting structures	5259	27	23.97

Notes: Volume refers to total number of diagnoses recorded in 2012. OON/IN ratio indicates relative frequency of out-of-network to in-network use.

in principle refuse to pay non-contracted providers at all, in practice they face demand-side pressure to provide some coverage for out-of-network care. For example, employers may want to ensure coverage for employees who need care while traveling for work or for employees or dependents who do not live near headquarters.[7] Any charge that the hospital levies on the patient in excess of what the insurer reimburses is known as a "balance bill."[8]

Most insurers' policies for determining payment amounts for out-of-network care rely on "usual and customary" rates. The definition of "usual and customary" may vary across insurers or even within an insurer's product portfolio, but typically relies on some notion of the prevailing market rate for a given service. Measuring out-of-network prices has historically presented a challenge for

[7] Insurers often pay some portion of the hospital's billed charges for out-of-network care, and these payments can be substantial. See Creswell et al. (2013) for anectodal evidence that insurers in certain markets pay substantial amounts in the form of chargemaster prices to out-of-network hospitals. Prager and Tilipman (2019) discuss this further in the context of regional Massachusetts insurers.

[8] The federal No Surprises Act, which went into effect in 2022, prohibits balance-billing consumers in certain circumstances, particularly that of emergent care. Prior to its implementation, many states had their own protections against balance bills. These regulations typically do not apply to voluntary out-of-network care.

Table 2: Demographics of Patients Seeking OON Care

	In Network, All Payers		In Network, NN Plans		Out of Network	
	Mean	SD	Mean	SD	Mean	SD
Age	51.555	18.846	44.046	16.999	39.836	15.360
Female	0.664	0.472	0.591	0.493	0.646	0.478
Severity Weight	17.477	79.937	20.734	58.777	40.080	230.519
Distance	8.263	9.518	8.263	9.518	11.279	3.992

Notes: Summary statistics for demand estimation sample of 2011–2012 outpatient visits. Columns 1–2 summarize the demographics of patients seeking in-network care. Column 2 refers to patients enrolled in New Hampshire's two "Narrow network" insurers. Column 3 summarizes these demographics only for patients seeking out-of-network care.

health economists. It requires inferring payment policies in a market that lacks posted prices, among transacting pairs that lack contracts and typically have low transaction volumes. Driven in part by these measurement difficulties, existing work on insurer-provider negotiations has assumed away the possibility of out-of-network transactions. This assumption may be a reasonable simplification in the settings typically studied by existing hospital-insurer bargaining papers, which study health care services that are rarely obtained out of network (e.g. inpatient hospital care). In the case of the outpatient care sample studied in this paper, however, out-of-network transaction volume can be substantial. Outpatient care has grown from less than one third of hospital revenues in 1995 to nearly half (Deloitte 2018), highlighting the importance of understanding out-of-network prices for the hospital market as a whole.

Measuring out-of-network prices is integral to this paper. We begin by inferring the general structure of out-of-network payments from insurers' public documents, and then turn to the data for specific parameter values. Table 3 quotes sections describing out-of-network payment policies from several insurers' policy documents. These excerpts demonstrate that insurers have well-defined policies governing out-of-network payments. The price determinations rely on some notion of a prevailing market price.

Insurers are not always explicit about how they define the prevailing market price, but when they are explicit, their definitions often refer to FAIR Health benchmarks. FAIR Health is a private health analytics firm that sells health care data products to health insurers, providers, employers, and other entities. Its products are based on a near-universal sample of claims from fee-for-service Medicare and privately insured patients. Among its flagship products are the FH Charge Benchmarks, which many insurers use as an input to determining out-of-network prices. This product reports quantiles summarizing the distribution of charge prices at the level of a geographic area-treatment type pair. Charge prices are sticker prices that are billed by hospitals but are typically substantially higher than the negotiated in-network prices actually paid by insurers, and than the prices paid by most patients. The Charge Benchmarks product is updated twice a year using a rolling twelve-month window of claims data. Insurers that purchase the Charge

Benchmarks can then use a predetermined percentile of the charge price distribution as an input to their determination of out-of-network prices, as indicated by the quotes from Aetna's, Cigna's, and United's policies in Table 3. Although FAIR Health emphasizes that its products are not intended as suggested prices, insurers' out-of-network payment policies are often informed by these products.

Table 3: Insurer Policies on Out-of-Network Payments

Insurer	Relevant Quote From Policy
Aetna	We get information from FAIR Health [...] For most of our health plans, we use the 80th percentile to calculate how much to pay for out-of-network services
BCBS of Massachusetts	Reimbursement for out-of-network providers will be based on a usual and customary fee schedule
Cigna	Under this option, a data base compiled by FAIR Health, Inc. (an independent non-profit company) is used to determine the billed charges made by health care professionals or facilities in the same geographic area for the same procedure codes using data. The maximum reimbursable amount is then determined by applying a percentile (typically the 70th or 80th percentile) of billed charges, based upon the FAIR Health, Inc. data
Harvard Pilgrim	When using Non-Plan Providers, the Plan pays only a percentage of the cost of the care you receive up to the Usual, Customary and Reasonable Charge for the service
Tufts	Reasonable Charge is the lesser of the: amount charged; or amount that we determine to be reasonable, based upon nationally accepted means and amounts of claims payment
United	Affiliates of UnitedHealth Group frequently use the 80th percentile of the FAIR Health Benchmark Databases

We infer insurers' policies with respect to the charge benchmarks by comparing payments for services rendered by out-of-network providers to the commonly used charge benchmark percentiles. We construct the analog of the FAIR Health benchmarks from our claims data (described in Section 2.4) by closely following FAIR Health's algorithm.[9] The algorithm is public; we describe it in detail in Appendix B. Each out-of-network claim is matched to its benchmark based on procedure code (CPT code), geographic area, and date of most recent benchmark release. We then examine the distribution of the ratio of the paid amount to the benchmarks.

Figure 1a shows the distribution of the ratio of paid amounts to the 60th percentile benchmarks for out-of-network facility claims by one of the insurers in our sample, Tufts Health Plan. When the hospital bills the insurer for more than the benchmark (charge > benchmark), the insurer typically only pays the 60th percentile benchmark, as indicated by the spike in the solid line distribution at 1.[10] When the hospital bills less than the benchmark (charge < benchmark), the

[9] Prior to FAIR Health's entry in late 2011, many insurers used benchmarks produced by the United Healthcare-owned firm Ingenix, whose benchmark products were calculated using a nearly identical algorithm to FAIR Health's but are now defunct (Bernstein 2012).

[10] Although the bulk of the mass is clustered near 1, many out-of-network claims are not paid based on this multiple. This is partially attributable to noise in our measure of the benchmarks for facility-based claims. Whereas

insurer nearly always pays precisely the entire billed amount, as shown by the solid line in Figure 1b. We therefore model Tufts Health Plan as paying the minimum of the hospital's billed charges and the 60th percentile charge benchmarks for out-of-network claims. This means that although the few hospitals whose charge prices are below the benchmark may be able to raise their out-of-network payments from Tufts by increasing their charge prices, they will still be constrained by Tufts' out-of-network payment policy once charge prices equal or exceed the benchmarks.[11]

Figure 1: Out-of-Network Payments by Tufts

(a) Payment Relative to Benchmark, All Cases

(b) Payment Relative to Charge, Charge < Benchmark

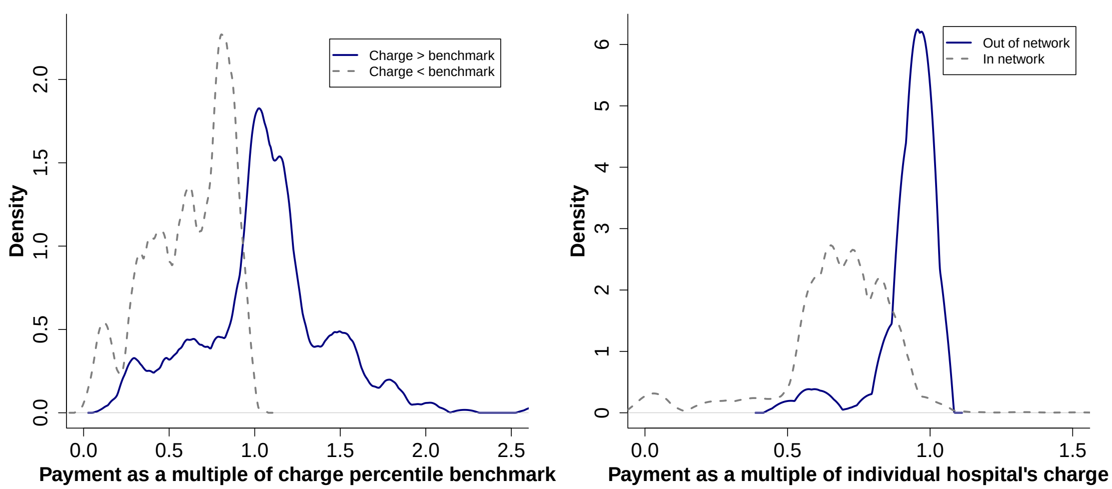

Notes: Tufts Health Plan's payment amounts for out-of-network outpatient hospital transactions in a flagship PPO plan, as a multiple of the 60th percentile charge benchmark for the corresponding procedure code. This plan typically pays out-of-network hospitals at 100 percent of the 60th percentile benchmark.

We use the procedure that underlies Figure 1 to infer insurers' policies for out-of-network payments. If an insurer has a complete provider network within New Hampshire, as is the case for Harvard Pilgrim, this requires examining its claims from other markets in which it has a narrow

FAIR Health uses the near-universe of privately insured claims and the universe of fee-for-service Medicare claims, our all-payer claims databases only capture the near-universe of privately insured claims. Our measure of the benchmark percentiles is therefore necessarily noisy. Figure A.1 shows that our constructed measure provides a good approximation of the proprietary FAIR Health data for professional claims, where we have data directly from FAIR Health.

[11] Endogenizing the charge prices themselves is beyond the scope of this paper. Anecdotally, experts on hospital pricing believe that hospitals are constrained from arbitrarily raising charge prices. Raising charge prices too high can trigger contract renegotiations with insurers with which the hospital has a percent-of-charges contract, potentially negating the gains from the higher charge prices. Charge prices perceived as excessive may also be dismissed by courts when hospitals sue patients to recoup unpaid bills. Finally, hospitals may face a reputational cost if they are seen as gouging patients. In this paper, we treat individual hospitals' charge prices as fixed. In the counterfactuals, the potential endogeneity of the charge prices is irrelevant because of the structure of the proposed policies we simulate.

network. These out-of-network policy inferences are facilitated by comprehensive data on insurers' networks, described in Section 2.4. If an insurer has a complete provider network across our entire sample, we use the insurer's stated policies, relying on the same documents as Table 3. We then use the inferred policies to construct off-contract prices for pairs of insurers and hospitals that do not necessarily have a contract.

2.4 Data and Descriptive Statistics

Health Care Claims Data: Data for estimating the hospital choice model and constructing other inputs to the bargaining model are drawn from the 2010–2012 New Hampshire All-Payer Claims Database (APCD) and the Massachusetts APCD. Private health insurers contribute data for the APCDs to the state agency that manages the data and uses it for policy-relevant analysis. In New Hampshire, this is the Comprehensive Health Care Information System (CHIS); in Massachusetts, it is the Center for Health Information and Analysis (CHIA). We use all available sample years to estimate the hospital demand model, and 2011–2012 to estimate the bargaining model.[12]

The New Hampshire and Massachusetts APCDs contain claims originating both within and outside of their respective states, as long as they are claims for services to enrollees of health insurance plans headquartered in the state. For example, a resident of New Hampshire who is insured through a Massachusetts-based employer will be included in the Massachusetts APCD. Each claim contains information on the patient's demographics, the insurance plan, the identity of the health care provider, the diagnosis, the services rendered, and prices.

There are multiple price variables in the APCDs, all of which are required for our analysis. Charge prices measure what the provider bills the insurer or the patient. As discussed in Section 2.3, charge prices are irrelevant for the vast majority of health care services obtained; importantly, they do not dictate payment amounts when an insurer and a hospital have a contract. Instead, allowed amounts and insurer paid amounts measure the insurer's contracted price with the in-network provider. In the case of an out-of-network provider, these variable measure the amount the insurer agrees to pay the provider off-contract. We use the allowed and paid amounts to construct measures of equilibrium negotiated prices used to estimate our bargaining model (Section 3.5). We also compare them to charge price benchmarks (Section 2.3) to infer insurers' out-of-network payment policies. Finally, also reported in the data are amounts for which patients are directly responsible under their insurance plan: deductibles, copays, and coinsurance.

We supplement the APCDs with hospital characteristics drawn from the American Hospital Association (AHA) Annual Survey Database and from the Centers for Medicare and Medicaid Services (CMS). Characteristics used in the analysis include teaching status, bed count, for-profit

[12] Prior to 2011, we are less confident that the geographic delineations used to calculate out-of-network price benchmarks in 2012 can be extrapolated farther back (see Appendix B for a description of how we calculate the benchmarks).

and non-profit status, system ownership, the number of licensed practical nurses per bed, and the presence of certain service lines such as neonatal intensive care units. In addition, we calculate driving distances from patient five-digit zip codes to hospitals for use in the hospital demand model.

Hospital Networks Data: To determine which hospital-insurer pairs have a negotiated contract, we use data on insurers' hospital networks. These data were hand-collected from New England insurers' current and archived plan documentation, as described in Prager (2020).[13]

In some cases, an insurer may classify a hospital as an in-network provider for its generous plans (such as PPO plans) while classifying it as an out-of-network provider for its narrow-network plans (mainly HMO plans). The analysis needs to capture whether an insurer-hospital pair has *any* negotiated price contract that an insurer can invoke if its enrollees get care at the hospital. We therefore define a hospital that is classified by an insurer as in-network in at least one plan type as having a negotiated price contract with that insurer. We define a hospital as out-of-network only if it is not classified as in-network even in the insurer's broadest-network plans.

Figure A.2 shows the hospital networks and distribution of enrollees for three illustrative insurers: Harvard Pilgrim (HPHC), Tufts, and MVP. The blue shaded areas show enrollment counts. The red solid (hollow) circles show in-network (out-of-network) hospitals, scaled by the insurer's patient volume flowing to each hospital. Tufts' and MVP's largest PPO networks cover only 8 hospitals each. Tufts' in-network hospitals tend to be clustered in the southeastern part of the state, where the majority of Tufts' New Hampshire enrollees live, and near to its primary market of Massachusetts. Similarly, MVP, which operates primarily in New York State and Vermont, has multiple pockets of enrollees in New Hampshire. The majority of MVP's in-network hospitals are also located near the enrollee masses. Some of MVP's out-of-network hospitals nevertheless attract substantial patient volumes, as shown by the large hollow circles. Notably, Tufts also covers Dartmouth-Hitchcock Medical Center (the large dot on the western border). Although Dartmouth-Hitchcock is geographically distant from the bulk of enrollees, it commands high willingness to pay due its status as the state's premier academic hospital. It also attracts a high share of out-of-network volume from MVP enrollees. Finally, Harvard Pilgrim has substantially higher enrollment in New Hampshire than either Tufts or MVP, with relatively high enrollment counts throughout the state. All New Hampshire hospitals are in Harvard Pilgrim's network.

Outpatient Hospital Sample: In the empirical analysis, we restrict our attention to health care services that are performed in an outpatient, rather than inpatient, setting. We do this for three primary reasons. First, the share of hospital expenditures attributed to outpatient care has

[13] Many claims databases, including the Massachusetts APCD, include a variable for a provider's network status. However, these variables are reported unreliably; for example, Harvard Pilgrim does not populate the field at all. We therefore view the network information collected directly from insurers' plan documentation as substantially more reliable.

been growing rapidly to the point where outpatient expenditures and inpatient expenditures are roughly equal (Deloitte 2018). Second, in our sample, out-of-network care for outpatient services is considerably more common than for inpatient hospitalizations. Third, the data we use to construct off-contract prices is based on the FAIR Health outpatient benchmark data (see Section 2.3 and Appendix B). To infer inpatient benchmarks for out-of-network reimbursements would require use of diagnosis-related groups (DRGs), which are not reliably reported in the APCDs. Reconstructing DRG classifications from the data without proprietary software would introduce additional noise into our off-contract price measures.

Appendix E describes in detail the construction of the outpatient sample. A key sample restriction is to subset to health care utilization whose primary procedure codes for non-emergent and non-inpatient services are in the top 1,000 codes by hospital revenue. These top 1,000 codes account for 96.7 percent of hospital outpatient revenue and 98.5 percent of hospital outpatient volume. Because these procedures are also often performed in physician office settings, our sample includes non-hospital visits for these procedures as an outside option. Our sample includes patients living both in and near New Hampshire. We include in the sample all enrollees who live in New Hampshire as well as those living in any Massachusetts zip code within the 75th percentile of distance traveled to a New Hampshire hospital.

Table E.1 shows the summary statistics for our demand estimation sample, subsetting just to hospital visits.[14] The share of hospital visits originating from the largest insurer, Anthem, is 62.2 percent. Tufts' share of in-sample visits is substantially larger than its New Hampshire enrollment share because Tufts is among the top three insurers in Massachusetts, from which we also draw data. Among visits to hospitals and their affiliated providers, 89.8 percent are to New Hampshire hospitals and the remaining 10.2 percent are to Massachusetts hospitals. Among narrow-network insurers, 26.4 percent of visits are to out-of-network hospitals.[15]

Patient Sample for Bargaining Model: To estimate the bargaining model described below in Section 3.5, we use a random sample of 5,000 households. For each household member, we merge in the annual prevalence of health care services in each decile of severity. Severity is measured by weights w_d, described in Appendix C. Severity quantiles and their associated prevalences are calculated from all commercially insured patients in the New Hampshire APCD, separately by sex and five-year age band. We use these prevalences to construct, for each patient, the predicted WTP and hospital volumes under each network configuration. In the bargaining estimation, each patient from the 5,000-household sample is then scaled up by the appropriate sample weight to represent the insurers' complete enrollment panel.

[14] Visits to hospitals or hospital-affiliated providers comprise 41.4 percent of the sample; the majority of the sample consists of visits to standalone physician offices and other providers not affiliated with a hospital.

[15] This is higher than the 10.1 percent of out-of-network visits in the raw data because the demand estimation sample subsets to patients living in or near New Hampshire, where the fraction of out-of-network providers is higher than in Massachusetts.

3 Model and Estimation

The goal of our model is to make inferences from the equilibrium networks and prices observed in the data. Private health insurers negotiate with hospitals to arrive at a contracted price. We model these negotiations as pairwise Nash bargaining interactions, but depart from the hospital bargaining literature by specifying strictly positive transaction volumes and prices.

The model proceeds in two stages:

1. Insurer m and hospital h engage in bilateral negotiations that, if successful, determine the in-network price p_{mh}.

2. With some probability f_{id}, patient i enrolled in insurer m's plan gets sick and requires procedure d. The patient chooses a hospital or other provider, which may or may not be in insurer m's network.

The estimation proceeds in two steps. First, we estimate a model of hospital choice corresponding to Stage 2 using maximum likelihood. Second, we estimate the insurer-hospital bargaining model corresponding to Stage 1 using a modified generalized method of moments (GMM) that nests both equality and inequality moments. This GMM step uses objects constructed from the hospital choice model estimates and our measure of off-contract prices.

In the sections that follow, we discuss the model and estimation pertaining to Stage 1. The discussion of Stage 2 is relegated to Appendix F, as we follow a well-established literature to estimate a discrete choice model of hospital demand. We make two departures from the literature. First, we assume patients incur a hassle cost of seeking care from out-of-network hospitals (e.g. prior authorization). Second, we allow hospital choices to be a function of the patient's expected balance bill, or out-of-pocket payment in the case of going to an out-of-network hospital. The balance bill is a function of the insurer's out-of-network payment policy, the hospital's charge price, and (in counterfactuals) regulation.

The remainder of this section is structured as follows. Section 3.1 derives the equilibrium prices in case of agreement, and Section 3.2 describes the conditions for an in-network agreement. Section 3.3 discusses how our model predicts hospital exit. Section 3.4 shows how the predictions of the bargaining model with out-of-network transactions depart from the predictions of a canonical model with zero volume in case of disagreement. Sections 3.5 and 3.6 discuss estimation identification, respectively.

3.1 Model Setup: Price Negotiation

In Stage 1 of the model, insurer m and hospital h negotiate over the in-network price. A successfully negotiated contract specifies a price p_{mh} that hospital h will be paid for treating insurer

m's enrollees, and assigns the hospital to be in the insurer's network.[16] We follow the literature in assuming that each hospital-insurer pair negotiates a single price index p_{mh} that is then scaled multiplicatively by a resource intensity weight w_d to determine the price for a given diagnosis or service (Gowrisankaran et al. 2015; Ho and Lee 2017). We calculate the weights relative to $w_d = 1$ for a routine blood draw, rather than using DRGs as in existing papers, because DRGs are not defined for the outpatient care we study. Appendix C describes w_d in detail. Negotiated prices at $w_d = 1$ fall primarily in the \$6 to \$13 range (Appendix Figure H.2).

In-network status grants the hospital a larger volume of the insurer's patients than out-of-network status. In the absence of a negotiated contract, the hospital remains out of network, but still in the patient's choice set. The relatively few services it does provide to insurer m's patients are paid according to the insurer's out-of-network payment policy, denoted by price p^0_m, plus any portion μ of the balance bill the hospital successfully collects from the patient.[17] The insurer's out-of-network payments depend only on the services provided, not the identity or cost structure of the hospital.[18]

Hospital Objectives

We model hospitals as profit maximizers. Hospital h's surplus from a contract with insurer m at a negotiated price p_{mh} is given by

$$S_h(m, p_{mh}) = \underbrace{(p_{mh} - c_{mh})\,\sigma^1_{mh}}_{\text{Profit from In-Network Patients}} - \underbrace{\left(p^0_m + \mu\left(p^c_h - p^0_m\right) - c_{mh}\right)\sigma^0_{mh}}_{\text{Profit from Out-of-Network Patients}} \tag{1}$$

where c_{mh} is the hospital's marginal cost of treating insurer m's typical patient, p^c_h is its charge price, μ is the average fraction of balance bills recouped by hospitals, and $\sigma^1_{mh} > \sigma^0_{mh}$ are the hospital's severity-weighted patient volumes from insurer m in the case of agreement and disagreement, respectively. The balance bill is equal to the charge price billed by the hospital when it is out of network less the out-of-network price paid by the insurer according to the insurer's policy. A hospital's volume under a given network configuration is predicted from the hospital demand model discussed in Appendix F. As discussed in Appendix C, we weight patient volumes by w_d, the resource intensity associated with the services provided.

[16] Hospital-insurer contracts are regularly updated with new prices. Throughout the paper, we omit time subscripts from the notation for brevity.

[17] In principle, a hospital should be able to balance-bill the patient the full difference between its posted charge price and the insurer's out-of-network price, p^0_m, rendering $\mu = 1$. In practice, this often does not happen for two primary reasons. First, patients incur disutility from being balanced billed and may thus choose a different hospital for voluntary out-of-network care. Second, there is a risk that a patient may not be able to pay the full bill, delaying the hospital's reimbursement and risking negative media attention.

[18] In practice, each insurer's out-of-network prices also vary across geographic markets that typically have multiple hospitals in each market (see Appendix B). We omit the geographic market subscripts from the notation for simplicity, but calculate the out-of-network prices separately within each market in the empirical application.

Insurer Objectives

We define insurers as maximizing a weighted difference of their enrollees' expected utility and their costs of paying for health care. Insurer m's enrollees' expected utility is a function of which hospitals are in its network: enrollees prefer to have more hospitals in the network. An alternative specification of insurers' objectives is profit maximization, which requires a model of health insurance plan choice. Because our data do not allow us to construct plan choice sets for the majority of patients, this is not feasible in our empirical application. We instead follow Gowrisankaran et al. (2015) in modeling the insurer as an imperfect agent for its enrollees. We note, however, that the qualitative differences that we outline in Section 3.4 between models assuming zero disagreement volumes and models accounting for positive disagreement volumes hold for both sets of insurer objectives.

Conditional on entering into negotiations, insurer m's surplus from a contract with hospital h at a negotiated price p_{mh} is given by

$$S_m(h, p_{mh}) = \underbrace{\left(\alpha W^1_{mh} - p_{mh}\sigma^1_{mh} - \psi^1_{mh}\right)}_{\text{Surplus with In-Network Status}} - \underbrace{\left(\alpha W^0_{mh} - p^0_m\sigma^0_{mh} - \psi^0_{mh}\right)}_{\text{Surplus with Out-of-Network Status}} - b_m \tag{2}$$

where α is insurers' weight on enrollee expected utility, and $W^1_{mh} > W^0_{mh}$ are the expected utilities in the case of agreement and disagreement, respectively. We do not allow α to vary by insurer because the same variation identifies α and insurer contracting costs b_m, which we allow to be insurer-specific. The contracting cost is discussed below in Section 3.2. The terms ψ^1_{mh} and ψ^0_{mh} denote the insurer's payments to other hospitals in the case of agreement and disagreement with hospital h, respectively. For example, $\psi^1_{mh} = \sum_{h' \neq h} \sigma_{mh'} p_{mh'}$, where other hospitals' volumes $\sigma_{mh'}$ are computed for the case where hospital h is in the network.

Equilibrium Negotiated Prices

In case of agreement, the negotiated price p^*_{mh} is the one that maximizes the Nash bargaining product:

$$p^*_{mh} = \arg\max_{p_{mh}} S_m(h, p_{mh})^{\gamma_h} S_h(m, p_{mh})^{1-\gamma_h}$$

where $\gamma_h \in [0, 1]$ is one less hospital h's Nash bargaining parameter vis-à-vis insurers.[19] Taking the derivative of the logged Nash product with respect to price,[20] the first-order condition describing

[19]That is, γ_h is the insurer bargaining parameter vis-à-vis hospital h. We allow the bargaining parameter to vary across hospitals, but denote it as the insurer bargaining parameter for consistency with the literature. We load heterogeneity across insurers onto other model parameters, namely hospital marginal costs c_{mh} and insurer bargaining costs b_m.

[20]The intermediate step is:

$$\gamma_h \frac{-\sigma^1_{mh}}{\alpha W^1_{mh} - p^*_{mh}\sigma^1_{mh} - \psi^1_{mh} - \left(\alpha W^0_{mh} - p^0_m\sigma^0_{mh} - \psi^0_{mh}\right) - b_m} = -\left(1 - \gamma_h\right) \frac{\sigma^1_{mh}}{\left(p^*_{mh} - c_{mh}\right)\sigma^1_{mh} - \left(p^0_m + \mu\left(p^c_h - p^0_m\right) - c_{mh}\right)\sigma^0_{mh}}.$$

p^*_{mh} becomes

$$p^*_{mh} = \frac{1}{\sigma^1_{mh}} \left[\begin{array}{l} (1-\gamma_h)\,\alpha \left(W^1_{mh} - W^0_{mh}\right) + (1-\gamma_h \mu)\, p^0_m \sigma^0_{mh} + \gamma_h \mu p^c_h \sigma^0_{mh} \\ +\gamma_h c_{mh} \left(\sigma^1_{mh} - \sigma^0_{mh}\right) - (1-\gamma_h)\left(\psi^1_{mh} - \psi^0_{mh}\right) - (1-\gamma_h)\, b_m \end{array} \right] \tag{3}$$

The first-order condition in Equation 3 contributes a set of moments used in estimation.

3.2 Model Setup: Network Formation

Insurer m and hospital h will enter into negotiations if the expected joint surplus from agreement, relative to the outside option of the hospital remaining out-of-network, is weakly positive. If the expected joint surplus is negative, then there can exist no price that would induce positive surplus for both parties individually.

We model negotiations as costly: the insurer must pay a contracting cost b_m for each pairwise negotiation. This modeling assumption is consistent with prior work (Ghili 2022) and motivated by the institutional details of the health care industry. Contract negotiations in this industry are notoriously resource-intensive, often lasting for months and requiring insurers to have a dedicated division for provider contracting.[21] We interpret the b_m parameter as a flavor of Coasian transaction cost.

The condition for agreement on an in-network contract is that insurer m's and hospital h's ex ante joint surplus from agreement is weakly positive. The ex ante joint surplus is simply the sum of the surplus available for splitting:

$$E_{mh} = \alpha W^1_{mh} - \psi^1_{mh} - \left(\alpha W^0_{mh} - p^0_m \sigma^0_{mh} - \psi^0_{mh}\right) - c_{mh}\sigma^1_{mh} - \left(p^0_m + \mu\left(p^c_h - p^0_m\right) - c_{mh}\right)\sigma^0_{mh} - b_m \tag{4}$$

The Nash-in-Nash structure of Stage 1 guarantees that, if $E_{mh} \geq 0$ and the pair enters negotiations, then an agreement will be reached. We leverage this in the estimation by inferring a weakly positive joint surplus if and only if an agreement is observed. The negotiated price p_{mh} does not enter into E_{mh} because its negative effect on the insurer's portion of the surplus, $-p_{mh}\sigma^1_{mh}$, is precisely offset by its positive effect on the hospital's surplus, $p_{mh}\sigma^1_{mh}$.

Network status is therefore entirely determined by three factors (net of bargaining costs): the ratio of willingness-to-pay for including the hospital in the network (relative to out of network) against the marginal cost of treating an in-network patient at that hospital (relative to out of network); the relative reimbursements to *other* hospitals if a hospital is removed from the network; and a hospital's ability to balance-bill patients for out-of-network care. When the parties anticipate that no agreement will be reached, the insurer will refuse to enter into negotiations in order to avoid incurring its contracting cost b_m.

[21] It is likely that contracting costs vary across hospitals and insurers. Ideally, we would estimate them as separate parameters using each insurer and hospital's network inclusion conditions separately. However, we do not have sufficient variation in our data to separately identify these parameters, along with our other parameters of interest.

3.3 Modeling Hospital Exit and Negotiator Sophistication

A key contribution of our framework is that, by allowing for positive transaction volumes in the event of disagreement, we can model hospital exit decisions, as well as how negotiator expectations of these exits affect equilibrium outcomes. Exit occurs if a hospital's total profit across insurers is negative, even if some insurers contribute positive profits in isolation.

Hospital Exit Condition: In the standard hospital-insurer bargaining framework, where disagreement volumes are zero, hospitals' disagreement payoffs are bounded below at zero:

$$\pi_h^{standard} = \sum_{m \in M_h} (p_{mh} - c_{mh})\sigma_{mh}^1$$

where M_h is the set of insurers with which hospital h has contracts. Since hospitals can always reject price offers below cost ($p_{mh} < c_{mh}$), this structure mathematically precludes the possibility of exit. A hospital's profit can never be negative as it will simply reject any contract offering below-cost prices, opting to earn zero profit from that insurer. Only if *every* insurer is unwilling to pay a price above cost does the hospital cease to treat any patients (exit). This is a rare event under reasonable parameter values; it never occurs in our counterfactuals.

In contrast, our framework with nonzero disagreement volumes allows a hospital's total variable profit to become negative when regulated out-of-network prices fall sufficiently below cost and regulation requires the hospital to treat out-of-network patients.[22]

Formally, a hospital h will exit if its total variable profit across all insurers becomes negative:

$$\pi_h = \underbrace{\sum_{m \in M_h} (p_{mh} - c_{mh})\sigma_{mh}^1}_{\pi_h^{standard}} + \underbrace{\sum_{m \notin M_h} (p_m^0 + \mu(p_h^c - p_m^0) - c_{mh})\sigma_{mh}^0}_{\text{Profit from OON Insurers}} < 0 \tag{5}$$

Unlike $\pi_h^{standard}$, this profit function includes a second term capturing potentially negative profits from treating out-of-network patients. The exit condition can be rewritten to show explicitly when out-of-network price regulation will trigger exit:

$$\sum_{m \in M_h} (p_{mh} - c_{mh})\sigma_{mh}^1 < - \sum_{m \notin M_h} (p_m^0 + \mu(p_h^c - p_m^0) - c_{mh})\sigma_{mh}^0$$

A hospital will exit the market when its profit from in-network patients is insufficient to offset losses from treating out-of-network patients at regulated prices. Taking derivatives produces the

[22] The Emergency Medical Treatment and Labor Act (EMTALA) explicitly prohibits hospitals from denying care to emergency patients on the basis of financial considerations, such as insurance. While this paper focuses on voluntary care largely performed outside the emergency room setting, there have been calls to couple out-of-network price regulation with a requirement that providers treat out-of-network patients (Fiedler and Ly 2023).

following comparative statics:

1. Exit is likelier when out-of-network prices decrease: $\partial\pi_h/\partial p_m^0 = (1-\mu)\sigma_{mh}^0 > 0$
2. Exit is likelier when out-of-network volume increases, and out-of-network price is below cost: $\partial\pi_h/\partial\sigma_{mh}^0 = (p_m^0 + \mu(p_h^c - p_m^0) - c_{mh}) < 0$
3. Exit is likelier when the rate of balance billing collection decreases: $\partial\pi_h/\partial\mu = (p_h^c - p_m^0)\sigma_{mh}^0 > 0$

Exit outcomes are also a function of negotiators' ability to anticipate exit. We introduce the notion of naive and sophisticated negotiators, where the primary distinction is different assumptions about the disagreement payoffs. We describe each in turn.

Naive Negotiators: When hospital h and insurer m negotiate with naive expectations, they incorrectly assume the disagreement outcome preserves the hospital's market participation regardless of overall profitability. The joint surplus from agreement is defined as in Equation 4, with disagreement modeled as:

$$\Pi_h^{0,naive}(m) = (p_m^0 + \mu(p_h^c - p_m^0) - c_{mh})\sigma_{mh}^0$$

$$\Pi_m^{0,naive}(h) = \alpha W_{mh}^0 - p_m^0\sigma_{mh}^0 - \psi_{mh}^0$$

The equilibrium price is identical to the one defined in Equation 3. The critical insight is that this can lead to equilibria where $p_{mh} < c_{mh}$ if insurer m has sufficient bargaining power and p_m^0 is regulated to below cost, potentially triggering a "surprise" hospital exit when Equation 5 is satisfied. This is the scenario implicitly described by opponents to regulation, including some providers: that insurers will be unwilling to pay in-network prices above the regulated out-of-network price, even when the out-of-network price fails to cover costs.

Sophisticated Negotiators: With sophisticated expectations, negotiators correctly anticipate whether disagreement would lead to hospital exit. Let us define $\pi_h^{OON}(m)$ as hospital h's total profits if it disagrees with insurer m:

$$\pi_h^{OON}(m) = \sum_{m' \neq m}(p_{m'h} - c_{m'h})\sigma_{m'h} + (p_m^0 + \mu(p_h^c - p_m^0) - c_{mh})\sigma_{mh}^0$$

This changes the disagreement payoffs relative to the naive case. In the case where the hospital would remain open even if out of network with insurer m (i.e., $\pi_h^{OON}(m) \geq 0$), the model reverts to the same equilibrium as the naive case.[23] We now discuss the case where disagreement would lead to exit ($\pi_h^{OON}(m) < 0$).

For agreement to be viable, the hospital's total profits must remain non-negative. This creates a constraint on the minimum acceptable price for the hospital. The minimum acceptable price is such that the hospital's total profits across all insurers are at least zero. The insurer will only agree to pay this minimum price if doing so yields higher surplus than its disagreement payoff (which is now based on hospital closure). This implies that the hospital will remain open if the following joint surplus condition is satisfied:

$$\alpha(W_{mh}^1 - W_m^{hc}) - (\psi_{mh}^1 - \psi_m^{hc}) - b_m \geq c_{mh}\sigma_{mh}^1 - \sum_{m' \neq m} (p_{m'h} - c_{m'h})\sigma_{m'h} \tag{6}$$

If Equation 6 holds, then an agreement is possible. The left side represents the insurer's gain from keeping the hospital open, while the right side represents the minimum subsidy required to prevent the hospital from exiting. Note that this is similar to the naive joint surplus term, E_{mh} except for two critical differences. The first is that the insurer W_{mh}^0 and ψ_{mh}^0 terms are replaced with their closure analogs, W_m^{hc} and ψ_m^{hc}. The second is that this is now a function of the hospital's profits from all other contracts (the last term on the right). The logic is that, with closure, the hospital's surplus from agreement is not just a function of the bilateral negotiation outcome with insurer m, but also a function of foregoing lost profits with all other insurers by closing. If the inequality does not hold, no feasible price exists that would make agreement preferable to both parties.

When agreement is possible, the equilibrium price is determined by Nash bargaining as usual:

$$p_{mh}^{soph} = \frac{1}{\sigma_{mh}^1}[(1-\gamma_h)[\alpha(W_{mh}^1 - W_m^{hc}) - (\psi_{mh}^1 - \psi_m^{hc}) - b_m] + \gamma_h \sigma_{mh}^1 c_{mh}] \tag{7}$$

In equilibrium, sophisticated negotiators will not allow a hospital to exit if the joint value of keeping it open exceeds the subsidy needed to cover its losses from treating out-of-network patients from other insurers. As a result, with sophisticated negotiators, hospital exit occurs less frequently than naive models would predict, even under aggressive price regulation.

3.4 Bias from Ignoring Out-of-Network Transactions

Accounting for nonzero disagreement volumes also affects our estimates of hospitals' ability to sustain price regulation before leaving networks or exiting. Empirical work on bargaining typically observes negotiated prices as an equilibrium outcome, and uses them to infer a set of structural

[23] The disagreement payoffs are as follows. For the hospital, $\Pi_h^{soph}(m) = (p_m^0 + \mu(p_h^c - p_m^0) - c_{mh})\sigma_{mh}^0$ if $\pi_h^{OON}(m) \geq 0$, and $\Pi_h^{soph}(m) = 0$ otherwise. For the insurer, $\Pi_m^{soph}(h) = \alpha W_{mh}^0 - p_m^0 \sigma_{mh}^0 - \psi_{mh}^0$ if $\pi_h^{OON}(m) \geq 0$, and $\Pi_m^{soph}(h) = \alpha W_{mh}^{hc} - \psi_{mh}^{hc}$ otherwise. Here, W_{mh}^{hc} represents enrollee welfare when hospital h exits the market, and ψ_{mh}^{hc} represents payments to other hospitals when h exits.

parameters pertaining to costs (marginal or fixed) and Nash bargaining weights. Misspecification of the disagreement payments $p^0_m \sigma^0_{mh}$ biases these structural parameters. Any of the parameters may be biased due to misspecification of the disagreement payments. To ease interpretation, it is therefore helpful to consider all but one set of parameters as being fixed at known values, and derive the bias on one set of free parameters. Here, we illustrate the bias when estimating hospital marginal costs c_{mh} and fixing all other parameters. Biases in marginal cost estimates are of interest to antitrust regulators due to their direct relationship with margins, which are used in merger evaluation.

The estimated hospital cost under the assumption of zero disagreement values will be biased upward whenever the true payments from the insurer to the hospital in the event of disagreement are "large enough." We derive this condition formally in Appendix G, and discuss it at length there. For ease of exposition, we present the underlying intuition by discussing two comparative statics rather than the inequality as a whole.

All else equal, the larger is the true out-of-network payment $p^0_m \sigma^0_{mh}$, the larger the upward bias on the hospital cost estimate. The intuition is that, when the true out-of-network payment is large due to high volume (σ^0_{mh}) or a high price (p^0_m), the hospital's true disagreement value is also large. The assumption of zero disagreement volumes therefore overstates the hospital's true surplus from agreement. As a result, the surplus implied by the observed negotiated price must instead be rationalized by a high cost estimate.

This has important implications for policy. When cost estimates are biased upward, analyses that assume zero disagreement payoffs will understate the true magnitude of price reductions from policies that restrict the bargaining game. This arises from an understatement of true hospital markups and therefore understates their ability to absorb price decreases without reaching the exit condition in Equation 5. Moreover, if policy-makers rely on economists' estimates of markups, they may craft policies that erroneously assume hospitals are capturing little producer surplus.[24] In Section 5.2, we show how the biased cost estimates affect the predicted effects of counterfactual policies.

3.5 Estimation of Model Parameters

We use the model conditions from the price first-order conditions and the network status conditions to form moments for estimation. There are five sets of parameters to estimate from the model: the insurers' weights on expected utility α, the insurers' Nash bargaining weights γ_h, the average fraction of balance bills recouped by hospitals μ, the hospital marginal costs c_{mh}, and the insurer's contracting cost b_m. All other objects in the model are predicted from the hospital demand model (see Appendix F) and treated as data. This section describes the three sets of moments that enter into our generalized method of moments estimation. We defer the discussion of identification to

[24] Berry et al. (2019) makes a forceful argument for more careful estimation of markups.

Section 3.6.

Moments from Price FOCs: The first set of moments comes from the equilibrium price first-order conditions (FOCs). In our equilibrium price condition (Equation 3), prices are observed, whereas hospital marginal costs, c_{mh}, are parameters to be identified. We express hospital h's marginal cost for treating a patient with resource intensity $w_d = 1$ as a function of observables g_{mh}:

$$c_{mh} = \lambda g_{mh} + \nu_h \tag{8}$$

where λ is a parameter vector and ν_h is the unobservable component of hospital costs. The observable characteristics in g_{mh} on which we project costs include observable hospital characteristics that determine a hospital's cost structure; insurer fixed effects, which allow heterogeneous patient populations across insurers; and year fixed effects, which allow for flexible statewide trends in cost growth.[25]

For hospital-insurer pairs that have a negotiated contract, we define the econometric error for the first set of moments as the difference between the projected cost from Equation 8 and the cost implied by the first-order conditions on equilibrium prices:

$$\xi_{mh} = \lambda g_{mh} - \frac{1}{\gamma_h \left(\sigma^1_{mh} - \sigma^0_{mh}\right)} \left[\begin{array}{l} p^*_{mh}\sigma^1_{mh} - (1-\gamma_h)\,\alpha\left(W^1_{mh} - W^0_{mh}\right) - (1-\gamma_h\mu)\,p^0_m\sigma^0_{mh} \\ -\gamma_h\mu p^c_h\sigma^0_{mh} + (1-\gamma_h)\left(\psi^1_{mh} - \psi^0_{mh}\right) + (1-\gamma_h)\,b_m \end{array} \right]$$

Moments from Observed Costs: In Nash bargaining models, bargaining weights, γ_h, are notoriously difficult to separately identify from hospital marginal costs, c_{mh}. We therefore include an additional set of moments that match model predictions to observed data on average hospital costs.

If hospital marginal costs were readily measurable, we would simply be able to treat c_{mh} as data. However, as with prior work, we observe data on total rather than marginal costs. In addition, our sample of services provided in a hospital *outpatient* setting poses a further challenge: existing methods of scaling total costs to per-patient costs scale to inpatients rather than outpatients. We therefore supplement those methods with additional data to calculate average variable costs per outpatient, as described in Appendix D.

For this set of moments, we construct hospital-year level predictions of average costs and match these to our per-outpatient cost measure from observed data (Cuesta et al. 2024). The econometric

[25] The hospital characteristics we include are: logged bed count, to capture returns to scale; teaching status, to capture the additional costs and subsidies arising from training medical students, residents and fellows; non-profit and for-profit system status, to capture potential cost advantages from professional management or non-profit objective functions resulting in effective marginal costs diverging from standard marginal costs (Lakdawalla and Philipson 2006); and the number of licensed practical nurses per bed, to capture lower costs due to substituting from physician labor to lower-priced inputs.

error for these moments is:

$$\xi_h^c = \frac{\sum_m \lambda g_{mh}(\sigma_{mh}^1 + \sigma_{mh}^0)}{\sum_m(\sigma_{mh}^1 + \sigma_{mh}^0)} - \bar{c}_h^o$$

where $\bar{c}_h^o$ represents observed average costs from hospital cost reports.

We then search for parameters to set both sets of marginal cost-related errors orthogonal to instruments z_{mh} and z_h respectively:

$$\mathbb{E}\left[\xi_{mh} z_{mh} \ \xi_h^c z_h\right] = 0$$

The instruments z include the hospital characteristics and year fixed effects included in the cost moments; the potential balance-bill for each hospital-insurer pair; and a plausibly exogenous shifter of hospital demand, but not costs: the number of the insurer's enrollees who live within a 20-mile radius of the hospital. The logic of this is similar to our demand model: the mass of an insurer's consumers who live near a particular hospital should determine hospital prices through volume, but should be uncorrelated with idiosyncratic shocks to hospitals' marginal costs of treatment.

Moments from Network Formation: In addition to the equality moments contributed by hospital-insurer-year cells in which we observe an agreement, each hospital-insurer-year contributes an inequality from the network formation conditions discussed in Section 3.2. We require these conditions for two reasons. First, a primary goal of the paper is to examine how negotiated prices change with different assumptions about the magnitudes of out-of-network transaction volumes and out-of-network price benchmarks. However, varying the level of out-of-network payments may result in insurers or hospitals deciding it is more profitable to enter into a formal contract (and negotiate an in-network price) rather than remain out-of-network under counterfactual policies. Our model therefore needs to incorporate firm decisions about the network status of hospitals observed to be out of network, as several recent papers have done (Ghili 2022; Liebman 2022; Ho and Lee 2019). Second, the estimation procedure must account for the fact that in our setting, network status is endogenously determined. Since some networks are incomplete, using only the first-order conditions from in-network hospitals would lead to biased parameter estimates.

To construct these network moments, we follow the spirit of the literature on moment inequalities (Ho 2009; Pakes 2010; Pakes et al. 2015). Formally, we define insurer m's and hospital h's ex ante joint surplus from agreement as:

$$\begin{aligned} E_{mh}(\boldsymbol{\theta}) &= \alpha W_{mh}^1 - \psi_{mh}^1 - \left(\alpha W_{mh}^0 - p_m^0 \sigma_{mh}^0 - \psi_{mh}^0\right) - c_{mh}\sigma_{mh}^1 - \left(p_m^0 + \mu\left(p_h^c - p_m^0\right) - c_{mh}\right)\sigma_{mh}^0 - b_m \\ &= \alpha\left(W_{mh}^1 - W_{mh}^0\right) - \psi_{mh}^1 + \psi_{mh}^0 + \left(-\sigma_{mh}^1 + \sigma_{mh}^0\right)c_{mh} + \mu\left(p_h^c - p_m^0\right)\sigma_{mh}^0 - b_m \end{aligned}$$

where c_{mh} is projected from Equation 8, and $\sigma_{mh}^1, \sigma_{mh}^0, W_{mh}^1, W_{mh}^0, \psi_{mh}^1, \psi_{mh}^0$ are predicted from the demand model. The penultimate term captures the effect of balance billing: when a hospital is out of network, it may be able to bill its patients for the difference between its charge price p_h^c and the insurer's out-of-network payment p_m^0, but it only successfully collects a fraction μ of these attempted balance bills.

If hospital h is in insurer m's network, then both parties must have positive gains from trade at the observed negotiated price and at the current parameter guesses $\hat{\boldsymbol{\theta}}$, relative to the outside option of the hospital remaining out-of-network.

We assume that insurers and hospitals have expectations over their surplus from a contract and that they predict these gains with error.[26] Let ω_{mh} be the difference between the parties' expected total surplus from agreement and the realized surplus, and let $\mathbb{E}[\omega_{mh}|\mathcal{J}] = 0$, where $\mathcal{J}$ is the insurer's and hospital's information set at the time of contracting decision.[27] Then:

$$\begin{aligned}\mathbb{E}[E_{mh}(\boldsymbol{\theta})|\mathcal{J}] &= E_{mh}(\boldsymbol{\theta}) - \omega_{mh} \\ &= \left[\alpha\left(W_{mh}^1 - W_{mh}^0\right) - \psi_{mh}^1 + \psi_{mh}^0 + \left(-\sigma_{mh}^1 + \sigma_{mh}^0\right)c_{mh} + \mu\left(p_h^c - p_m^0\right)\sigma_{mh}^0 - b_m\right] - \omega_{mh}\end{aligned}$$

Each hospital-insurer pair that is observed to have a negotiated contract therefore contributes one inequality that imposes a lower bound on the total available surplus from agreement:

$$0 \leq \mathbb{E}[E_{mh}(\boldsymbol{\theta})|\mathcal{J}] = E_{mh}(\boldsymbol{\theta}) - \omega_{mh} \tag{9}$$

We refer to these inequalities as network inclusion moments.

For hospital-insurer pairs that are observed not to have a contract, our model requires that there exists no price that would make both the hospital and the insurer better off than if they do not have a negotiated contract. Thus, in the estimation, we impose that at the current parameter guesses $\hat{\boldsymbol{\theta}}$, there exists no price that would make both parties' surpluses positive.[28] The resulting inequality for estimation is given by:

$$0 > \mathbb{E}[E_{mh}(\boldsymbol{\theta})|\mathcal{J}] = E_{mh}(\boldsymbol{\theta}) - \omega_{mh} \tag{10}$$

Each hospital-insurer pair that is observed not to have a negotiated contract therefore contributes a single inequality, defined by Equation 10, that imposes upper bounds on the surpluses from agreement. We refer to these inequalities as network exclusion moments.

Collectively, the network inclusion and exclusion conditions are what Ghili (2022) calls network

[26] For example, they may be uncertain as to how other insurers or hospitals might react to any contracting decision, which would impact the ultimate negotiated prices and estimates of gains from trade.

[27] Recall that the Nash-in-Nash setup assumes that all bargaining parties have the same information set.

[28] Equivalently, if h is observed not to be in m's network, then we assume that the highest price that m would be willing to pay while still maintaining a positive surplus is less than the lowest price that h would be willing to accept while still maintaining a positive surplus. This is because the insurer's surplus is monotonically decreasing in price and the hospital's surplus is monotonically increasing in price.

stability conditions. Because of the mean-zero assumptions on ω and v conditional on insurer and hospital information sets, when the sample of inequalities grows large, the errors tend to zero in the limit. Given instruments $z \in \boldsymbol{J}$, our estimating equations for the network inclusion conditions become:

$$0 \leq E_{mh}(\boldsymbol{\theta})(z)$$

In the estimation, we incorporate these inequality moments into the GMM estimation by penalizing violations of these conditions proportional to their magnitude. For instance, if insurer m and hospital h are observed not to have a network but Equation 10 is violated—that is, if the implied total surplus at the current parameter values is positive—we penalize the objective function by a multiple of the magnitude of the violation. We then stack these moments together with the equality moments from the bargaining first-order conditions (Section 3.1) and search for parameters $\boldsymbol{\theta}$ that minimize the weighted sum of the network inclusion, network exclusion, and bargaining first-order condition moments.

3.6 Identification of Bargaining Model Parameters

This section discusses identification of a subset of bargaining model parameters. Due to length, a detailed discussion of the identifying variation for the remaining parameters is relegated to Appendix H. Here, Table 4 briefly summarizes the key sources of variation identifying each set of parameters.

We discuss identification for one set of parameters here to give a flavor of the argument. Identification of hospital marginal cost parameters, λ, and bargaining weights, γ_h, relies primarily on the two sets of equality moments presented in Section 3.5, leveraging variation in observed negotiated prices and hospital average costs. Each set of moments helps discipline parameters implied by the other set, and this interaction between model-implied costs from the first-order conditions and average cost moments is essential for separating estimates of marginal costs from bargaining power. In particular, without the average cost moments, high observed prices could be rationalized either by high marginal costs or high ability of a hospital to extract surplus.

Marginal cost parameters, λ, are identified through two channels. First, the marginal cost projections in Equation 8 must align with observed cost data, allowing hospital characteristics to explain patterns in average costs. Second, conditional on bargaining weights, these cost projections must match implied costs from the first-order conditions. While the relationship between average costs and hospital characteristics helps pin down the majority of the cost parameters, this second channel serves two roles: it provides additional variation in the relationship between prices and hospital characteristics across hospitals, and it helps identify insurer fixed effects in λ through systematic variation in implied costs within-hospital across insurers.

Table 4: Summary of Variation Identifying Model Parameters

Symbol	Parameter name	Primary source of variation
λ	Hospital cost parameters	Cross-hospital variation in average costs; correlation between hospital characteristics and implied costs from FOCs
γ_h	Insurers' bargaining weight for each hospital h	Price variation after controlling for costs; relationship between WTP and price differences
μ	Fraction of balance bill recouped	Correlation between potential balance bills and negotiated prices; network formation patterns for high-balance-bill hospitals
α	Insurers' weight on enrollee WTP	Network inclusion patterns and price sensitivity for high-WTP hospitals
b_m	Contracting costs	Network incompleteness patterns for subset of insurers

4 Model Estimates

4.1 Hospital Demand Estimates

Appendix Table A.1 shows the results of the hospital demand model for outpatient care. Our preferred specification is Column 2, in which we instrument for the in-network hospital indicator using enrollee geography as described in Appendix F. Consistent with the literature on hospital and physician demand, distance enters negatively and significantly into the utility function. Patients dislike dislike exposure to higher balance bills.

Most of the interactions between patient and hospital characteristics follow the expected signs. Older patients are less willing to travel. Patients are more willing to travel for hospitals with a cardiac catheterization lab, larger hospitals, and teaching hospitals. Patients requiring more resource-intensive procedures are also more willing to travel to larger hospitals and hospitals with cardiac catheterization labs and also, again, critical access hospitals.

We include two key terms in the model to capture patient disutility from seeking out-of-network care. The first is an in-network indicator and the second is the potential "balance bill" a consumer might incur from the hospital.[29] We interpret the coefficient on the in-network indicator as measuring a hassle cost that patients pay for seeking out-of-network care that does not vary by hospital. Such costs can include, for instance, receiving prior authorization from the insurer or the cost of acquiring price information from the hospital.

The potential balance bill is calculated as the difference between a hospital's posted list price

[29] These assumptions are motivated by institutional details. First, given that we focus on voluntary care—as opposed to emergent care—consumers are more likely to make a price assessment before agreeing to seek care out-of-network. Second, while hospitals technically do have the ability to fully balance-bill patients for out-of-network care, patients rarely receive bills for the full charge amounts. In fact, the portion of a balance bill that hospitals typically collect is not well documented (Duffy et al. 2020). As such, while there may be a correlation between a patient's expectation of their balance bill and their actual ultimate out-of-pocket price burden, the correlation is not necessarily strong.

(i.e. its charge price, p_h^c) and the amount that insurer reimburses for out-of-network care (p_m^0). It is meant to capture the fact that patients may have expectations over their out-of-pocket burden at any particular hospital, which may dissuade them from seeking care at that hospital. The key assumption we make here is that patients are broadly aware of their *potential* out-of-network bill, but may be uncertain as to the balance bill the hospital will actually levy. The coefficient on the potential balance bill captures patient disutility from higher out-of-pocket costs, but may be diluted by patients' lack of knowledge about those costs. In the counterfactuals, the balance bill coefficient becomes moot because the proposed regulations we study include bans on balance billing.

The coefficient on the hospital's in-network indicator is positive and significant, confirming that patients receive significant disutility from getting outpatient care out-of-network. The estimate translates to an average patient willing to travel about six additional miles to receive care from an in-network facility as opposed to an out-of-network facility, or a little more than 50 percent farther than the average distance traveled in our sample (10.5 miles). The result is stronger for non-hospital health care providers such as physician office: patients are willing to travel almost double the average distance traveled to a hospital to access a an in-network physician as opposed to an out-of-network one. The combination of both of these forces results in positive consumer willingness-to-pay for hospitals to be part of an insurer's network, which then enters into the insurer objective function (Equation 2).

4.2 Hospital Costs and Bargaining Parameters

Table 5 presents the results of our bargaining estimation for an abridged set of hospitals (the five smallest and five largest in terms of volume), with bootstrapped standard errors. The full set of parameters is reported in Appendix Table A.2. The first two columns show the results for our full model, including out-of-network transactions and network inclusion and exclusion moments, while the right panel displays estimates from the standard model without out-of-network transactions.

The estimated hospital cost parameters are all positive and exhibit substantial variation across both hospitals and insurers, with a mean estimated marginal cost of about \$8.60.[30] We find evidence of economies of scale, with a 1 percent increase in beds associated with a \$1.58 decrease in marginal costs. Teaching hospitals such as Dartmouth-Hitchcock have higher marginal costs (+\$1.23), suggesting the presence of more specialized or resource-intensive treatment options. Hospitals owned by systems have substantially lower costs (−\$3.17), potentially reflecting operational efficiencies or enhanced bargaining leverage with suppliers. Similarly, hospitals employing more licensed practical nurses (LPNs) per bed have lower costs (−\$1.13), suggesting substitution of less expensive nursing labor for more costly physician labor.

We also estimate considerable variation in hospitals' marginal costs across insurers. Aetna (the

[30]This is a sensible magnitude. Recall that costs are reported at severity weight $w_d = 1$, which corresponds to a simple blood draw (Appendix C).

omitted insurer) and MVP have higher estimated marginal costs are higher than other insurers. Both insurers have small market shares, which may imply hospitals incur a higher administrative burden in claims processing or other administrative inefficiencies. In contrast, hospitals exhibit lower marginal costs when treating patients from insurers that have have a strong prescence, such as Harvard Pilgrim and Anthem.

The hospital bargaining weights, γ_h, exhibit substantial variation, ranging from 0.57 (Exeter Hospital) to the corner solution of 1. Many hospitals have an estimated bargaining weight of 1, indicating that insurers capture all surplus gains from these contracts. This pattern appears most frequently with Critical Access Hospitals (CAHs), suggesting either that these facilities struggle to negotiate effectively with insurers or that they operate under alternative negotiation frameworks. Nevertheless, even for non-CAH facilities, the model indicates insurers maintain considerable bargaining power, including with "star" hospitals that extract the most surplus (e.g., Dartmouth Hitchcock, Elliot Hospital, Exeter Hospital).

This insurer advantage persists despite parameter estimates showing that insurers place slightly higher weight on enrollee welfare from CAH hospitals (α). This differential weighting may reflect regulatory incentives to including CAHs in networks or other considerations related to ensuring sufficient coverage in rural areas. Indeed, most insurers in our data cover all CAHs despite the small enrollee populations they serve.

We estimate that hospitals recoup approximately 38.2 percent of potential balance bills from out-of-network patients ($\mu = 0.38$). The confidence interval rules out hospitals recouping 100 percent of balance bills ($\mu = 1$) but is consistent with recouping 0 percent ($\mu = 0$). This finding aligns with previous reports suggesting that, even before balance bill regulations were enacted, hospitals could often recoup only a portion of the potential balance bill (Fiedler 2020).[31]

Finally, we estimate a negligible bargaining cost for most insurers. As discussed in Section 3.6, these insurers include all hospitals in their networks, leaving us no way to plausibly identify a lower bound on costs. The narrow-network insurers, however, show heterogeneity in b_m. Tufts, in particular, has a contracting cost of $2.9 million, suggesting substantial frictions in their contracting process. In the counterfactuals, the level of bargaining costs will be relatively unimportant when evaluating comparisons across counterfactual scenarios.

4.3 Estimates Under Zero Disagreement Volumes

We now turn to the impact that allowing for out-of-network transactions has on the estimated cost parameters of the model. To do so, we hold fixed the estimated insurer weight on enrollee surplus (α), the bargaining weights (γ_h), the share of balance bill collected (μ), and contracting costs (b_m), and re-estimate the hospital marginal costs (c_{mh}) under the standard model with no out-of-

[31] We were unable to find reliable estimates of the balance bill recoupment fraction in the academic or industry literatures.

Table 5: Bargaining Model Estimates

	Full Model		Standard Model	
	Estimate	Std. Error	Estimate	Std. Error
Hospital marginal costs (c_{mh})				
Constant	11.1133	(0.4108)	12.3694	(0.4543)
Log(Beds)	-1.5820	(0.1918)	-1.5565	(0.2098)
Teaching	1.2320	(0.3908)	0.5203	(0.5215)
For-Profit	0.2175	(0.6885)	-0.9891	(0.8466)
System	-3.1657	(0.7481)	-2.8441	(0.7139)
LPNsPerBed	-1.1307	(0.2325)	-0.8115	(0.2415)
Cigna	-2.8454	(0.5514)	-2.5659	(0.5476)
MVP	1.2998	(0.7048)	-0.0470	(1.7066)
United	-5.4012	(0.5184)	-3.8330	(0.7425)
Anthem	-3.4430	(0.4913)	-3.0518	(0.6747)
Harvard Pilgrim	-3.8401	(0.6688)	-3.0772	(0.6180)
Tufts	-3.4968	(1.2960)	-1.6393	(1.3241)
Year 2012	0.3168	(0.3508)	0.9220	(0.5673)
Nash bargaining weight (γ)				
Alice Peck Day Memorial Hospital	0.9896	(0.0354)		
Concord Hospital	0.9996	(0.0016)		
Cottage Hospital	0.9994	(0.0029)		
Dartmouth Hitchcock Medical Center	0.6590	(0.1748)		
Elliot Hospital	0.6804	(0.0626)		
Littleton Regional Hospital	0.8731	(0.0941)		
Southern New Hampshire Medical Center	0.9962	(0.0127)		
St Joseph Hospital	0.7059	(0.0590)		
Valley Regional Hospital	0.8177	(0.1368)		
Weeks Medical Center	1.0000	(0.0001)		
Insurer weight on WTP (α, Hundreds)				
Non-CAH	1.3017	(0.1261)		
CAH	1.6695	(0.2485)		
Balance bill recoupment fraction (μ)				
μ	0.3822	(0.2123)		
Bargaining cost (b_m)				
Aetna	0.0021	(0.0045)		
Cigna	0.0026	(0.0171)		
MVP	1,682	(1,265)		
United	0.0002	(0.0007)		
Anthem	0.0035	(0.0099)		
Harvard Pilgrim	0.0001	(0.0003)		
Tufts	2,908,943	(684,294)		

Notes: Bootstrapped standard errors in parentheses. The first two columns report results from the full model, which estimates all bargaining parameters. The second two columns reports results from the restricted specification, which only estimates marginal costs under a simplified bargaining structure that assumes zero disagreement payoffs and volumes. These results only show the five smallest and five largest hospitals by volume. The full set of results is available in Table A.2.

network transactions. To implement this version, we remove out-of-network hospitals from patients' choice sets, and use the demand model from Table A.1 to recompute predicted hospital shares and WTP from the demand model parameters. The predicted demand quantities are then used to generate new predictions for total spending under the assumption that volumes and payments to out-of-network hospitals are zero. Finally, we re-estimate hospital marginal costs using a version of the bargaining estimation without network inclusion and exclusion moments.

The resulting estimates are reported in the last two columns of Table 5. In most cases, the full model yields substantially lower hospital marginal cost estimates than the model that shuts down the out-of-network channel. The magnitude of the bias is large: on average, marginal costs are estimated to be approximately 20 percent lower under the full model with nonzero disagreement volumes. Moreover, the direction and magnitude of the bias are consistent with the comparative statics discussed in Section 3.4 and Appendix G.

In the appendix, Figure A.3 plots the empirical analog of the comparative static on prices. It shows that hospitals that have greater effective price differentials between in-network and out-of-network care also tend to have greater bias in their model-predicted marginal cost estimates when disagreement payoffs are assumed to be zero. Though this is not the case for all of the insurer-hospital pairs, the majority of standard model estimates of hospital costs are biased upward.

Returning to the example of Dartmouth-Hitchcock, the full model predicts marginal costs that are 11 percent lower than the predictions of a model with no out-of-network volumes. This result makes sense: since Dartmouth-Hitchcock is a high-demand hospital, the full model predicts it would retain considerable out-of-network volume should it leave an insurer's network. This outside option allows it to negotiate a higher price. The model without out-of-network transactions rationalizes Dartmouth-Hitchcock's high price by loading it onto a high estimate of its marginal cost. In other words, allowing for out-of-network transactions implies a considerably higher hospital mark-up than would be estimated in their absence.

5 Regulation, Prices, and Exit

We conduct a series of policy counterfactual simulations using our bargaining model estimates to simulate equilibrium networks and in-network negotiated prices for various out-of-network price policies.

An influential set of policy proposals call for fixing out-of-network reimbursements to multiples of Medicare prices. A high-profile candidate for the 2020 Democratic presidential nomination proposed setting the cap at 200 percent of Medicare (Pete For America 2019). Other proposals have called for prices as low as 120 percent of Medicare (Kane 2019). These prices are substantially lower than the current standard based on FAIR Health benchmarks, which are greater than three times Medicare prices.

These proposals have consequently drawn considerable scrutiny from hospital and physician groups, with some warning that reducing out-of-network payments would jeopardize their long-run financial viability. Some groups have proposed requiring insurers and providers to settle disputes over out-of-network reimbursement through binding arbitration. Others have proposed *increasing* the standard by which providers are reimbursed to the full charge price (Luthi 2019). As such, we also simulate policies that decrease or increase the current benchmark prices.

Finally, we conduct counterfactuals mirroring proposed legislation for a subset of out-of-network care (surprise bills), but applying the legislation to all out-of-network payments. The Lower Health Care Costs Act of 2019 proposed to regulate surprise out-of-network billing by capping insurers' off-contract payments at median in-network prices in a given market (Alexander 2019).

To predict the impacts of these policies, we use our full model estimates from Table 5. Under the standard model, the procedure would involve using our estimated parameters and computing in-network prices, p_{mh}, for every hospital-insurer pair under the different out-of-network reimbursement structures. However, our analysis is complicated by the fact that imposing alternate disagreement payoffs may result in different networks being formed in equilibrium.

To incorporate this feature, our iterative simulation proceeds in a series of steps described in Appendix I. Because network links are allowed to change, finding an equilibrium is not guaranteed. As with other Nash-in-Nash bargaining settings, our setting also permits multiple equilibria. We discuss both of these issues in Appendix I.

In all counterfactuals, we restrict hospitals' ability to balance-bill patients, consistent with proposed regulations. We also hold fixed the hospitals' cost structures. The counterfactual results should therefore be interpreted as static, as they do not account for hospitals' ability to reduce their marginal costs by changing their production technology or merging.

5.1 Medicare-Based Out-of-Network Payment Caps

We begin by evaluating proposals to peg out-of-network prices to a multiple of Medicare prices. Because these prices are set by the government, we view this as the most likely form of out-of-network price regulation to be implemented by policy-makers. Pegging to any quantity derived from hospitals' charge prices or median in-network prices is vulnerable to gaming by hospitals.

Medicare reimbursements are substantially lower than observed in-network prices, and often lower than estimated hospital marginal costs. It is therefore not surprising that most proposals benchmark out-of-network prices to a multiple of Medicare price above one. We simulate the counterfactual equilibrium in-network prices and networks for a range of multipliers. In these simulations, we assume a hospital can refuse to treat out-of-network patients if those patients are unprofitable.

Figure 2 plots the results of this simulation for our seven insurers. The solid line is the volume-weighted average price among in-network hospitals. The set of hospitals contributing to the average price changes if hospitals' network status changes. Using our full model, reducing out-of-network

Figure 2: Predicted Negotiated Prices Vs. Multiples of Medicare Reimbursements

(a) Aetna

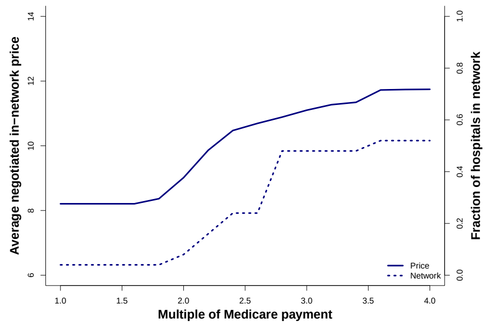

(b) Anthem

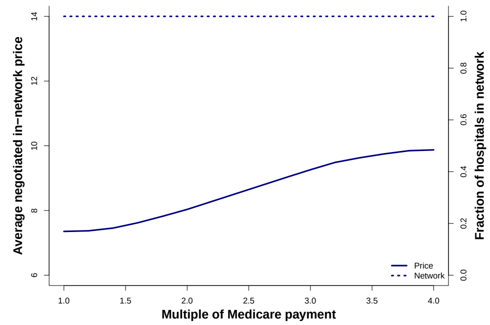

(c) Cigna

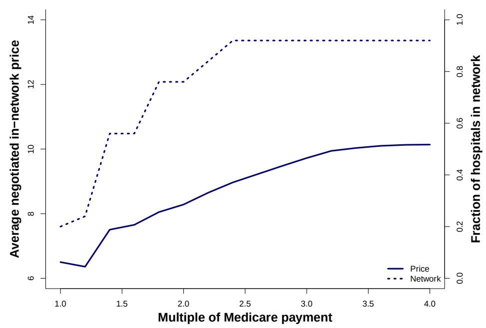

(d) Harvard Pilgrim

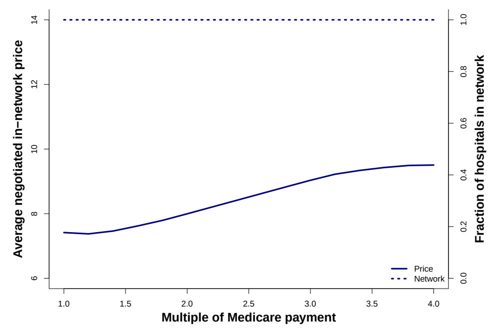

(e) MVP

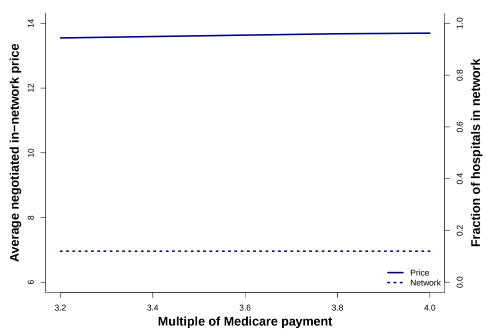

(f) Tufts

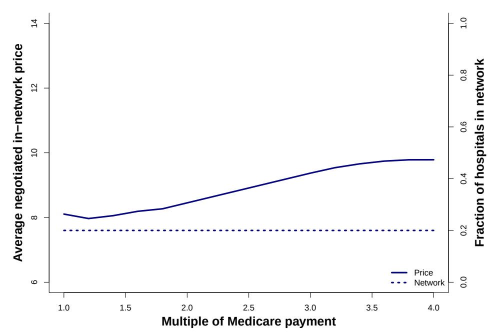

(g) United

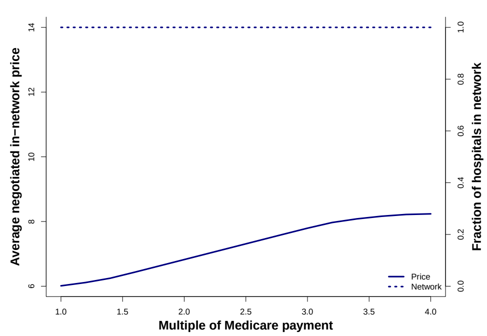

Notes: This figure plots results of counterfactual simulations varying out-of-network prices to be various multiples of Medicare reimbursements. The left vertical axis plots the counterfactual negotiated price for hospitals predicted to be in-network (volume-weighted average prices). The right vertical axis plots the fraction of hospitals predicted to be in-network. Gaps represent counterfactuals for which no equilibrium was found.

prices by half is predicted to result in a within-insurer average price reduction of 14 to 27 percent.[32] Large national insurers that have a small presence in New Hampshire (e.g. Aetna and United) see substantial price decreases (27 and 19 percent, respectively). Meanwhile, insurers that have full networks and strong presences in the state (e.g. Harvard Pilgrim, Cigna) see smaller, but still meaningful decreases (23 and 17 percent, respectively). Narrow-network insurers (MVP, Tufts) also see relatively smaller price decreases (below 15 percent).

Appendix Figure A.6 plots the distribution of predicted hospital prices within-insurer at each level of Medicare benchmarks. It shows that the most dramatic reductions in average in-network price are driven by compositional changes in in-network hospitals. For example, the hospitals with the highest in-network prices are dropped from Aetna's network at the lowest levels of out-of-network price.

While equilibrium price reductions are desirable to policy-makers, our results also demonstrate that capping out-of-network prices impacts hospital access. As shown by the dotted line in Figure 2, which plots equilibrium networks, price reductions and access appear to be in direct competition. As negotiated prices fall, the fraction of hospitals that are in the equilibrium network also falls. This is because, all else equal, lower prices paid to other hospitals reduce the insurer's available savings from bringing a hospital into its network, which reduces the joint surplus from agreement. Reducing the benchmark from 320 to 160 percent of Medicare is predicted to reduce network breadth by approximately 40 percentage points for some insurers (Aetna and Cigna). Across all insurers, proposals that peg out-of-network prices to values above 280 percent of Medicare would likely cause price reductions of 3–5 percent with minimal network disruptions.

The effect of regulating out-of-network prices is moderated by hospitals' ability to turn away patients. Indeed, full-network insurers such as Harvard Pilgrim and Anthem see price reductions with no impact on networks. This occurs because once the regulated OON price falls below a hospital's predicted marginal cost, hospitals can simply turn unprofitable patients away. This allows hospitals to have a lower-bound outside option of zero, at which point the bargaining devolves into the standard model without out-of-network transactions. Some have argued that out-of-network price regulation will not effectively reduce negotiated prices unless it is paired with a requirement for providers to accept all patients (Fiedler 2020). This set of counterfactuals shows that is not the case. As long as the regulated out-of-network price covers costs, regulating the out-of-network price will have substantial effects on in-network price. In Section 5.3, we consider the case where hospitals can have negative disagreement payoffs by mandating that they accept all patients, regardless of profitability.

[32] The insurers' average out-of-network prices range from 320 to 350 percent of Medicare. We report the price reduction from regulating out-of-network prices down from 320 percent of Medicare prices to 160 percent of Medicare prices.

5.2 Additional Counterfactuals without Hospital Closures

Bias from Assuming No Out-of-Network Transactions: Appendix Figure A.4 illustrates how conclusions about the counterfactual policies would differ under estimates from the standard model that assumes zero disagreement volumes. The light gray lines show the results of the same simulation, but using our bargaining model estimates from the second two columns of Table 5. Due to the higher estimated hospital marginal costs, the counterfactual in-network prices are always higher than those using our full model. Moreover, despite the higher prices, the equilibrium networks are (weakly) narrower. This illustrates the importance of accurately estimating hospital costs when conducting policy simulations whose goal is to reduce equilibrium prices. The standard model both overstates equilibrium prices and understates network breadth. In evaluating a policy proposal, this would cause overly pessimistic predictions about spending and, under some parameter values, about access to care.

Reducing Barriers to Out-of-Network Care: When regulation reduces out-of-network prices, insurers may have weaker incentives to steer patients away from out-of-network hospitals and may therefore impose lower hassle costs for out-of-network care. Although modeling endogenous hassle cost changes is outside the scope of this paper, we bound the effects of such a reduction by repeating the counterfactuals with patients' out-of-network disutility set to 10 percent of the actual estimate from our demand model. Appendix Figure A.5 shows the results. The slope of negotiated prices with respect to regulated out-of-network prices is steeper, because out-of-network volumes are larger, resulting in larger surplus changes from a given change in out-of-network prices. For instance, Aetna's reduction in prices rises from less than 30 percent in the baseline scenario to more than 50 percent. Insurers that previously saw no movement in networks (e.g. MVP), now see reductions in breadth at aggressive levels of regulation.[33]

Alternate Multiples of Charge Price Benchmarks: In Appendix I.2, we consider rescaling the out-of-network prices to alternate multiples of current benchmarks. This is meant to approximate the impact on in-network hospital prices of proposals to set out-of-network reimbursements closer to hospitals' charge prices. We find that doing so yields substantial negotiated price increases up until approximately 1.2 times the current benchmark (Figure I.1). Above this value, most hospitals are paid their charge prices because the benchmark begins to exceed the charge price.

Dynamic Payment Caps: Appendix I.3 considers the Alexander (2019) proposal to peg out-of-network prices to median negotiated in-network prices, which are typically substantially lower

[33] These effects are all expected, though we note that network breadth is a less meaningful metric when the distinction between in-network and out-of-network providers is erased. If prices for out-of-network care are similar to in-network care, and the hassle cost is negligible, then narrow networks pose no meaningful access issues.

than current out-of-network prices. In the first year of implementation, this would reduce in-network prices due to the worsening of hospitals' bargaining leverage. The following year, the median will be calculated from a negotiated price distribution that has shifted to the left, further reducing negotiated in-network prices. We find that both negotiated and out-of-network prices drop substantially, along with several insurers' network breadth (Figure I.2). Nearly the entire reduction in in-network prices obtains within the first two years, without the continued dynamic effects.

5.3 Forecasting Hospital Closures

The counterfactual simulations discussed thus far allow hospitals not only to leave insurers' networks, but also to turn away out-of-network patients if they are unprofitable ($p^0_m < c_{mh}$). Regulation of out-of-network prices could include a provision requiring hospitals to accept out-of-network patients at the regulated price, even if price is below cost. This section evaluates the impact of such regulation on hospital exit. In Section 3.4, we showed that existing models of hospital-insurer bargaining cannot accommodate an exit response, necessitating our richer model.[34] Appendix I describes how we amend our counterfactual simulation algorithm to accommodate closures.

Assumptions for Modeling Service Line Closures: Modeling hospital exit in the counterfactuals requires several assumptions. First, we assume that the hospital service lines used in our empirical analyses (Section 2.4) are separable from hospitals' other service lines. Prices for these service lines are therefore sufficient to determine whether a hospital will close them. We therefore interpret our hospital closure results as pertaining only to the service lines included in our outpatient sample (Appendix E), which accounts for approximately half of hospital revenues. Hereafter, we use the terms "exit" and "closure" interchangeably, to mean closures of these specific lines.

Second, we assume for this set of counterfactuals that hospitals cannot turn away patients, regardless of the hospital's network status with a patient's insurer. This assumption is what generates a notion of exit. If hospitals can turn away out-of-network patients, then they need never accept prices below marginal cost. Requiring hospitals to treat even unprofitable out-of-network patients, as suggested by some researchers (Fiedler and Ly 2023), produces scenarios in which a hospital can make strictly negative profits by remaining open. This requirement resembles existing regulation of emergency health care: the Emergency Medical Treatment and Active Labor Act (EMTALA) requires hospitals to accept and treat patients until they are stable, regardless of insurance status or ability to pay.

Finally, we assume that if the hospital's total variable profit summed across insurers is negative, that induces the hospital to exit. This assumption is a simplification, in that we check profits

[34] The assumption of zero out-of-network volumes precludes the possibility of care being reimbursed at below marginal cost. In these models, a hospital will only enter into a contract if the in-network price exceeds its marginal cost; and remaining out-of-network means no marginal costs are incurred.

only among the insurers for which have bargaining model estimates, not accounting for cross-subsidization from public payers or other service lines.[35] Our notion of exit allows a hospital to close a service line even if, in equilibrium, its contracts with a subset of insurers would be profitable in isolation. By contrast, standard models of hospital-insurer bargaining can only generate a notion of exit if the hospital's contracts with each insurer are individually unprofitable.

Closure Results: Figure 3 plots the results of the counterfactuals from Section 5.1, now permitting hospital exit while assuming negotiators are naive. Appendix A.7 shows the corresponding in-network prices. Consistent with the earlier results, the fraction of hospitals that remain in the insurer's network (dark green in the figure) drops slightly as out-of-network prices p_m^0 drop and in-network negotiated prices p_{mh} follow. Beyond the narrowing networks, however, Figure 3 also makes clear that severe price reductions will also induce some hospitals to close service lines when p_m^0 is set to a very low level.

While regulating out-of-network prices from their current value of approximately 320 percent of Medicare prices to 220 percent of Medicare prices would not affect the incidence of exits (or network breadth of some insurers), we find that regulating beyond this point could substantially affect access. Indeed, capping out-of-network reimbursements at Medicare prices is predicted to induce fifty percent of all hospitals to exit the market for in-sample service lines. While evaluating the relative welfare impacts of large price reductions against hospital closures is beyond the scope of this paper, these counterfactual simulations support concerns about providers exiting in response to various payment-reducing policy proposals.

With sophisticated negotiators, however, *no* hospital exits the market. Figure 4 shows the resulting networks and closures, and Appendix Figure A.8 plots the corresponding in-network prices. When a hospital would be induced to exit if it remained out-of-network at the regulated price, insurers correctly forecast this exit and willingly pay a higher in-network price to prevent closure. This occurs because enrollees value insurance more in a world that retains some access to the hospital (even if that access is outside the network) than a world without that hospital altogether. Given our estimated parameters, every hospital creates sufficient value that insurers are willing to pay to keep it open.

Moreover, some insurers' networks broaden as the out-of-network price is regulated lower. For instance, regulating to 260 percent of Medicare induces some hospitals to leave Aetna's network (as in the counterfactuals without hospital closures). Regulating them below this level induces hospitals to join the network, creating a broader network than at current levels of reimbursement. This seemingly counterintuitive result occurs because at sufficiently low out-of-network prices, there

[35] In reality, hospitals derive revenues from public payers in addition to private insurers. Even if private insurers' prices drop below cost, a hospital may be able to stay open profitably if Medicare or Medicaid profits exceed the losses from private patients. Since Medicare prices are generally lower than private insurers' prices (see Section 5.1) and Medicaid is less generous than Medicare in most states, we do not view this potential cross-subsidization as a serious threat to our assumptions.

Figure 3: Predicted Networks and Hospital Service Line Closures with Naive Negotiators

(a) Aetna

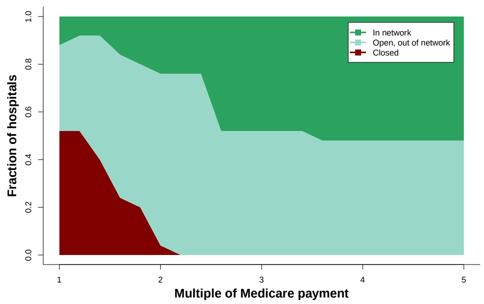

(b) Anthem

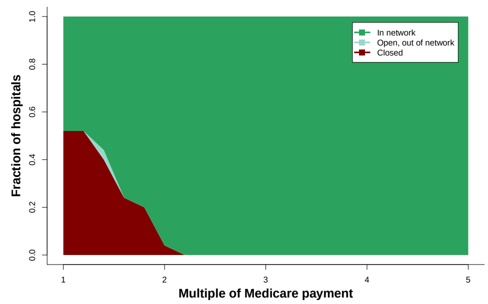

(c) Cigna

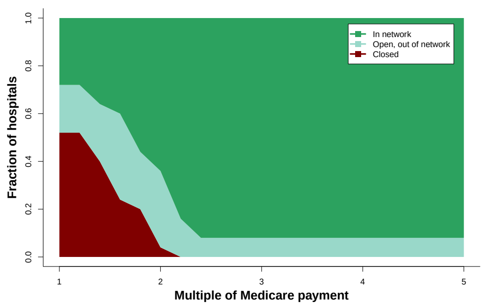

(d) Harvard Pilgrim

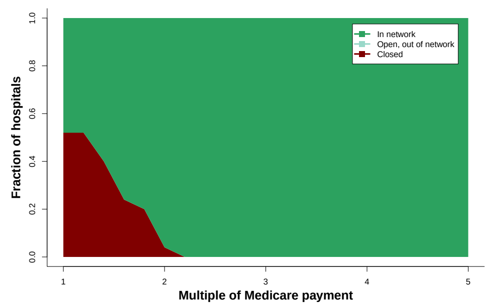

(e) MVP

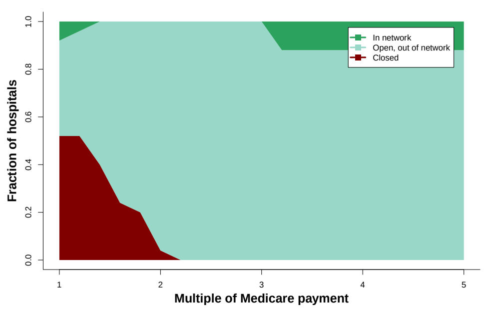

(f) Tufts

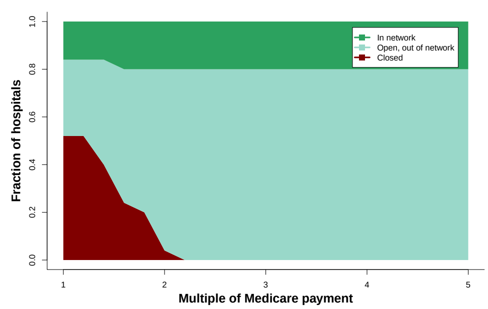

(g) United

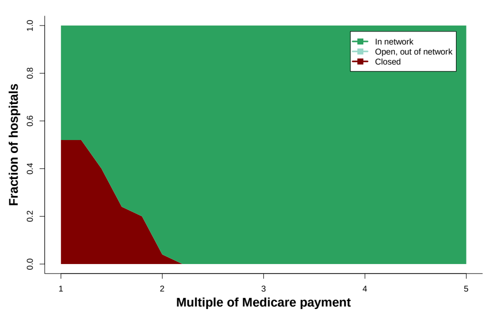

Notes: This figure plots results of counterfactual simulations varying out-of-network prices to be various multiples of Medicare reimbursements. The vertical axis plots the fraction of hospitals whose service lines are open and that are in network (dark green), open but out of network (light green), or exited from the market (red). The corresponding prices are plotted in Figure A.7.

Figure 4: Predicted Networks and Hospital Service Line Closures with Sophisticated Negotiators

(a) Aetna

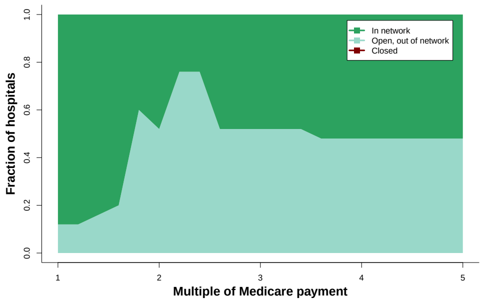

(b) Anthem

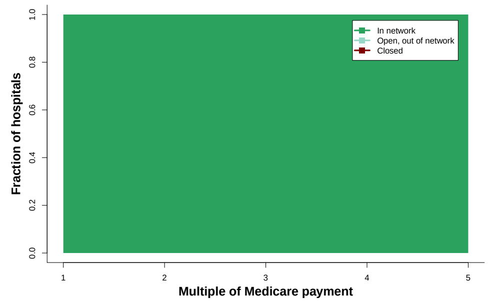

(c) Cigna

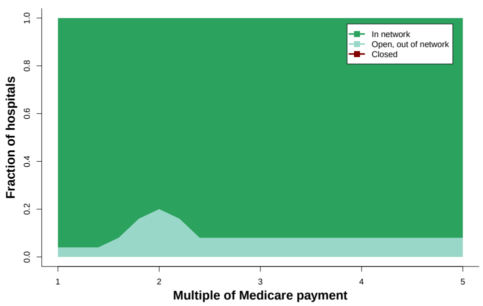

(d) Harvard Pilgrim

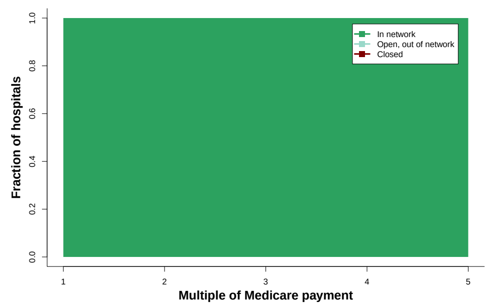

(e) MVP

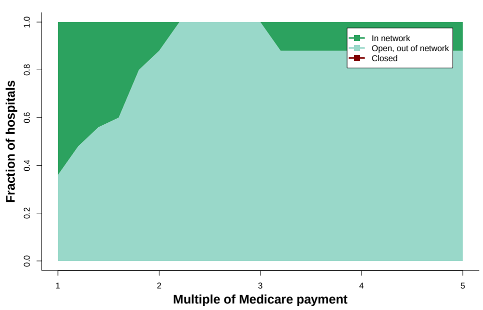

(f) Tufts

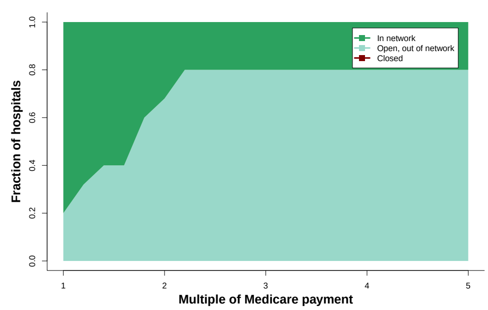

(g) United

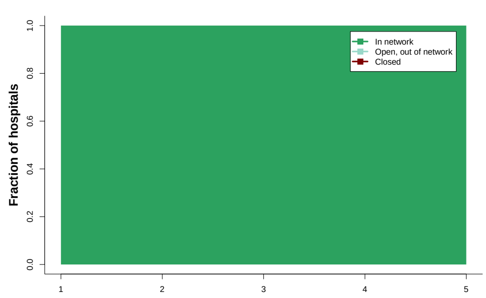

Notes: This figure plots results of counterfactual simulations varying out-of-network prices to be various multiples of Medicare reimbursements. The vertical axis plots the fraction of hospitals whose service lines are open and that are in network (dark green), open but out of network (light green), or exited from the market (red). The corresponding prices are plotted in Figure A.8.

are two possible outcomes from a negotiation: either there is disagreement and the hospital must close; or there is agreement at a price that keeps the hospital open. In-network status is the only mechanism for the insurer to pay the hospital a price that covers costs. Thus, with sufficiently low out-of-network prices—and hence threat of exit—hospitals that were previously out of network can enter insurers' networks. Consistent with this mechanism, the equilibrium negotiated prices at low regulated out-of-network prices are higher in the sophisticated case Appendix Figure A.8) than the naive case (Appendix Figure A.7).

6 Conclusion

Nash-in-Nash bargaining models are a workhorse tool of empirical work studying markets with negotiated prices. This paper demonstrates the importance of disagreement payoffs when Nash-in-Nash models are used in health care markets. Given the non-trivial volume of off-contract transactions, we show that these payoffs are crucial determinants of negotiated prices, equilibrium networks, and even hospital exit. The approach we develop should be straightforward to implement in future research. It allows out-of-network volumes and prices to enter the model without significant computational burden. Moreover, now that we have validated our measure of out-of-network prices, it can easily be calculated in the types of datasets commonly available to researchers.

Our empirical findings have direct implications for ongoing policy debates about the price-access tradeoff in health care. Our counterfactuals reveal that regulating out-of-network reimbursements would indeed reduce negotiated prices with in-network hospitals. Reducing out-of-network prices from current levels (approximately 320 percent of Medicare) to 280 percent of Medicare yields price reductions of 3–5 percent with minimal network disruption. However, more aggressive regulation can reduce network breadth. At out-of-network prices regulated to 160 percent of Medicare, the share of hospitals covered by insurers' networks drops by up to 40 percentage points.

Our model allows us to conduct the first (to our knowledge) empirical evaluation of claims that health care price regulation will lead to exit. We show that regulation's predicted effects on hospital exit depend critically on whether agents anticipate impending exits. When insurers and hospitals are naive about the possibility of exit, aggressive out-of-network price regulation can trigger "surprise" hospital closures as prices fall below costs. However, with sophisticated negotiators who anticipate that below-cost prices will force exits, insurers willingly pay prices that ensure hospitals' continued operation. This suggests that fears of widespread hospital closures from price regulation may be overstated, even for rural hospitals.

As policy-makers consider interventions in health care markets, our results highlight the importance of carefully considering how regulated disagreement payoffs affect hospital-insurer negotiations, network formation, and ultimately patient access.

References

Alexander, Lamar (2019) "Lower Health Care Costs Act," July.

Avalere (2015) "An Analysis of Policy Options for Involuntary Out-of-Network Charges in New Jersey."

Becerra, Xavier, Kathleen Foote, Michael Jorgenson, Cheryl Lee Johnson, Esther La, and Emilio Varanini (2018) "Complaint for Vilations of the Cartwright Act."

Bernstein, Nina (2012) "Insurers Alter Cost Formula, and Patients Pay More (Published 2012)," *The New York Times*.

Berry, Steven, Martin Gaynor, and Fiona Scott Morton (2019) "Do Increasing Markups Matter? Lessons from Empirical Industrial Organization," *Journal of Economic Perspectives*, Vol. 33, No. 3, pp. 44–68.

Biener, Adam I, Benjamin L Chartock, Christopher Garmon, and Erin Trish (2021) "Emergency Physicians Recover A Higher Share Of Charges From Out-Of-Network Care Than From In-Network Care: Study examines how much privately insured emergency department patients paid when they likely received a surprise bill and how much physicians received in these situations.," *Health Affairs*, Vol. 40, No. 4, pp. 622–628.

Chernew, Michael E, Maximilian J Pany, and Leemore S Dafny (2022) "Two Approaches To Capping Health Care Prices," *Health Affairs Forefront*.

Chernew, Michael, Maximilian Pany, and Richard Frank (2019) "The Case For Market-Based Price Caps," September.

Cooper, Zach, Stuart Craig, Martin Gaynor, Nir Harish, Harlan Krumholz, and john Van Reehen (2019a) "The Price Ain't Right? Hospital Prices and Health Spending on the Privately Insured," *Quarterly Journal of Economics*, Vol. 134.

Cooper, Zack, Hao Nguyen, Nathan Shekita, and Fiona Scott Morton (2019b) "Out-Of-Network Billing And Negotiated Payments For Hospital-Based Physicians," *Health Affairs*, Vol. 39, No. 1, p. 10.1377/hlthaff.

Cooper, Zack, Fiona Scott Morton, and Nathan Shekita (2020) "Surprise! Out-of-network billing for emergency care in the United States," *Journal of Political Economy*, Vol. 128, No. 9, pp. 3626–3677.

Craig, Stuart V, Keith Marzilli Ericson, and Amanda Starc (2021) "How important is price variation between health insurers?" *Journal of Health Economics*, Vol. 77, p. 102423.

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.

Creswell, Julie, Barry Meier, and Jo Craven McGinty (2013) "New Jersey Hospital Has Highest Billing Rates in the Nation," *The New York Times*.

Cuesta, José Ignacio, Carlos Noton, and Benjamin Vatter (2024) "Vertical Integration and Plan Design in Healthcare Markets," *NBER Working Paper*.

Deloitte (2018) "Growth in outpatient care: The role of quality and value incentives," article, Deloitte Insights.

Duffy, Erin L., Christopher Whaley, and Chapin White (2020) "The Price and Spending Impacts of Limits on Payments to Hospitals for Out-of-Network Care,"Technical report, RAND Corporation.

Ellison, Alia (2018) "California files antitrust suit against Sutter Health: 5 things to know," *Becker's Hospital Review*.

Fiedler, Matthew (2020) "Capping Prices or Creating a Public Option: How Would They Change What We Pay for Health Care?" white Paper, Brookings Institution.

Fiedler, Matthew and Bich Ly (2023) "How much can an out-of-network cap reduce in-network prices," *USC Schaeffer Brookings Initiative for Health Policy*.

Ghili, Soheil (2022) "Network formation and bargaining in vertical markets: The case of narrow networks in health insurance," *Marketing Science*, Vol. 41, No. 3, pp. 501–527.

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.

Hanbury, Mark (2019) "Nike stops selling on Amazon, kills deal after 2 years," *Business Insider*.

Henry, Tanya Albert (2025) "Medicare pay cuts: How they threaten patients' access to care," AMA News Wire article, American Medical Association.

Ho, Kate (2009) "Insurer-Provider Networks in the Medical Care Market," *American Economic Review*, Vol. 99, No. 1, pp. 393–430.

Ho, Kate and Robin Lee (2017) "Insurer Competition in Health Care Markets," *Econometrica*, Vol. 85, No. 2, pp. 379–417.

——— (2019) "Equilibrium Provider Networks: Bargaining and Exclusion in Health Care Markets," *American Economic Review*, Vol. 109, No. 2.

Horn, Henrick and Asher Wolinsky (1988) "Bilateral Monopolies and Incentives for Merger," *The RAND Journal of Economics*, Vol. 19, No. 3, p. 408.

Jacquis, William P. (2019) "Letter from the American College of Emergency Physicians," letter, American College of Emergency Physicians.

Kane, Arielle (2019) "Affordable Health Care for All: An American Solution to High Costs and Coverage Gaps," September.

Lakdawalla, Darius and Tomas Philipson (2006) "The nonprofit sector and industry performance," *Journal of Public Economics*, Vol. 90, No. 8-9, pp. 1681–1698.

Liebman, Eli (2022) "Bargaining in Markets with Exclusion: An Analysis of Health Insurance Networks," *Working Paper*.

Luthi, Susannah (2019) "Doctors, hospital groups gear up for surprise medical bill fight," *Modern Healthcare*.

Pakes, Ariel (2010) "Alternative Models for Moment Inequalities," *Econometrica*, Vol. 78, No. 6.

Pakes, Ariel, Jack Porter, Kate Ho, and Joy Ishii (2015) "Moment Inequalities and their Applications," *Econometrica*, Vol. 83, No. 1.

Pete For America, Campaign (2019) "Medicare for All Who Want It," September.

Prager, Elena (2016) "Tiered Hospital Networks, Health Care Demand, and Prices," *Working Paper*.

——— (2020) "Healthcare demand under simple prices: evidence from tiered hospital networks," *American Economic Journal: Applied Economics*, Vol. 12, No. 4, pp. 196–223.

Prager, Elena and Nicholas Tilipman (2019) "Hospitals Are Paid Their Charemaster Prices By Some Health Insurers," *Working Paper*.

Reinhardt, Uwe E. (2006) "The Pricing of U.S. Hospital Services: Chaos Behind a Veil," *Health Affairs*, Vol. 25, No. 1, pp. 57–69.

Schmitt, Matt (2017) "Do hospital mergers reduce costs?" *Journal of Health Economics*, Vol. 52, pp. 74–94.

Shepard, Mark (2022) "Hospital network competition and adverse selection: evidence from the Massachusetts health insurance exchange," *American Economic Review*, Vol. 112, No. 2, pp. 578–615.

Smith, Molly and Benjamin Finder (2022) "CBO's Proposals Do Not Address The Real Causes of Rising Commercial Health Insurance Premiums," AHA Stat blog, American Hospital Association.

Tilipman, Nicholas (2022) "Employer incentives and distortions in health insurance design: Implications for welfare and costs," *American Economic Review*, Vol. 112, No. 3, pp. 998–1037.

Wesolowski, Aaron and Benjamin Finder (2022) "RAND 4.0 Still Riddled with Methodology Flaws and Incomplete Data," AHA Stat blog, American Hospital Association.

Wikler, Elizabeth, Peter Basch, and David Cutler (2012) "Paper Cuts: Reducing Health Care Administrative Costs,"Technical report, Center for American Progress.

A Additional Tables and Figures

Figure A.1: Comparison of Benchmarks Reconstructed From Data Against Actual FAIR Health Benchmarks

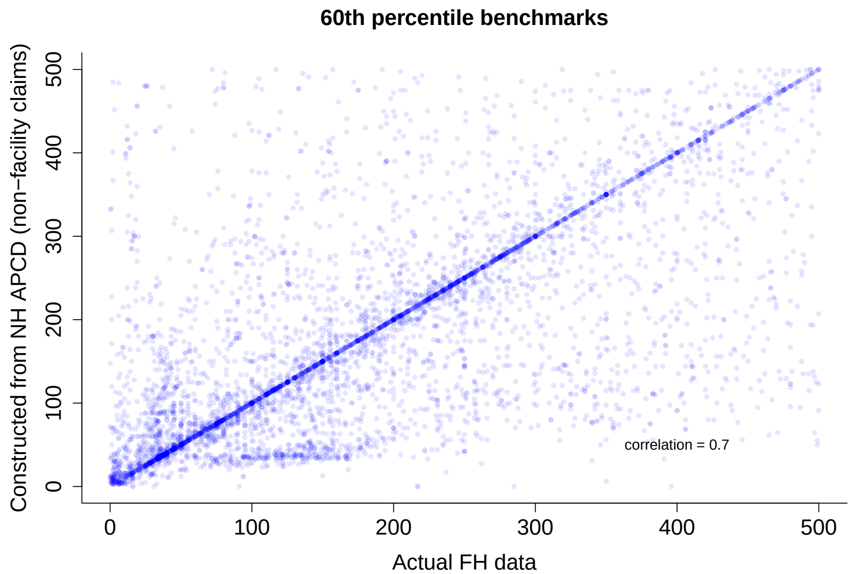

Notes: This figure plots our approximation of the FAIR Health Charge Benchmarks, calculated from the New Hampshire APCD, against the actual benchmark product purchased from FAIR Health. Each point is a CPT code-geographic market pair. Plot is for outpatient claims.

Table A.1: Hospital Demand Estimates

	(1) No IV		(2) Ctrl. Func.	
In-Network Hospital	0.0858	(0.1190)	1.9439***	(0.2032)
Other In-Network Provider	2.8057***	(0.2541)	2.9066***	(0.3565)
Balance Bill ($)	−0.0005	(0.0005)	−0.0003	(0.0015)
Distance (miles)	−0.3023***	(0.0134)	−0.3026***	(0.0126)
Distance^2	0.0006***	(0.0000)	0.0006***	(0.0000)
Distance × Age	−0.0001**	(0.0000)	−0.0001**	(0.0000)
Distance × Intensity Weight	−0.0000*	(0.0000)	−0.0000**	(0.0000)
Beds × Age	0.0000***	(0.0000)	0.0000***	(0.0000)
Beds × Intensity Weight	0.0000***	(0.0000)	0.0000***	(0.0000)
Beds × Distance	−0.0000	(0.0000)	−0.0000	(0.0000)
Beds × Female	0.0000	(.)	0.0000	(0.0000)
Critical Access × Age	−0.0014	(0.0025)	−0.0024	(0.0023)
Critical Access × Intensity Weight	−0.0008	(0.0004)	−0.0007	(0.0005)
Critical Access × Distance	0.0967***	(0.0136)	0.0993***	(0.0134)
Critical Access × Female	0.0000	(.)	0.0000	(0.0000)
Teaching × Age	−0.0101***	(0.0024)	−0.0096***	(0.0022)
Teaching × Intensity Weight	−0.0009***	(0.0003)	−0.0009***	(0.0002)
Teaching × Distance	0.0727***	(0.0043)	0.0727***	(0.0051)
Teaching × Female	0.0000	(.)	0.0000	(0.0000)
Cath Lab × Age	0.0024	(0.0024)	0.0018	(0.0023)
Cath Lab × Intensity Weight	0.0004	(0.0003)	0.0004	(0.0003)
Cath Lab × Distance	0.0910***	(0.0134)	0.0920***	(0.0128)
Cath Lab × Female	0.0000	(.)	0.0000	(0.0000)
MRI × Age	−0.0014	(0.0015)	−0.0014	(0.0015)
MRI × Intensity Weight	−0.0000	(0.0001)	−0.0000	(0.0001)
MRI × Distance	0.0312***	(0.0053)	0.0315***	(0.0055)
MRI × Female	0.0000	(.)	0.0000	(0.0000)
NICU × Age	0.0000	(.)	0.0000	(0.0000)
NICU × Intensity Weight	0.0000	(.)	0.0000	(0.0000)
NICU × Distance	0.0066	(0.0043)	0.0070	(0.0048)
NICU × Female	0.1057**	(0.0360)	0.1036***	(0.0313)
Neuro × Age	−0.0003	(0.0021)	−0.0002	(0.0020)
Neuro × Intensity Weight	−0.0002	(0.0002)	−0.0002	(0.0002)
Neuro × Distance	−0.0014	(0.0059)	−0.0014	(0.0056)
Neuro × Female	0.0000	(.)	0.0000	(0.0000)
Hospital FEs	Yes		Yes	
Up to Deg 1 of 1st Stage Resid	No		Yes	
Pseudo R^2	0.557		0.558	
Choices	18616		18616	

Notes: ***p<0.01, **p<0.05, *p<0.10. Results from multinomial logit provider choice model from years 2011–2012. IV columns estimated using a control function with boostrapped standard errors with 100 replications. *Choices* is the number of choice sets used in estimation. *In-Network Hospital* and *Other In-Network Provider* are indicators for whether the provider is in the patient's insurer's network and is a hospital or non-hospital provider, respectively. *Balance Bill* is the potential (maximum) dollar amount a patient would be charged out-of- pocket in cases where the hospital is out of network.

Table A.2: Bargaining Model Estimates

	Full Model		Standard Model	
	Estimate	Std. Error	Estimate	Std. Error
Hospital marginal costs (c_{mh})				
Constant	11.1133	(0.4108)	12.3694	(0.4543)
Log(Beds)	-1.5820	(0.1918)	-1.5565	(0.2098)
Teaching	1.2320	(0.3908)	0.5203	(0.5215)
For-Profit	0.2175	(0.6885)	-0.9891	(0.8466)
System	-3.1657	(0.7481)	-2.8441	(0.7139)
LPNsPerBed	-1.1307	(0.2325)	-0.8115	(0.2415)
Cigna	-2.8454	(0.5514)	-2.5659	(0.5476)
MVP	1.2998	(0.7048)	-0.0470	(1.7066)
United	-5.4012	(0.5184)	-3.8330	(0.7425)
Anthem	-3.4430	(0.4913)	-3.0518	(0.6747)
Harvard Pilgrim	-3.8401	(0.6688)	-3.0772	(0.6180)
Tufts	-3.4968	(1.2960)	-1.6393	(1.3241)
Year 2012	0.3168	(0.3508)	0.9220	(0.5673)
Nash bargaining weight (γ)				
Alice Peck Day Memorial Hospital	0.9896	(0.0354)		
Androscoggin Valley Hospital	0.7887	(0.1765)		
Catholic Medical Center	0.8299	(0.1060)		
Cheshire Medical Center	0.9692	(0.0661)		
Concord Hospital	0.9996	(0.0016)		
Cottage Hospital	0.9994	(0.0029)		
Dartmouth Hitchcock Medical Center	0.6590	(0.1748)		
Elliot Hospital	0.6804	(0.0626)		
Exeter Hospital	0.5667	(0.1551)		
Franklin Regional Hospital	1.0000	(0.0001)		
Frisbie Memorial Hospital	0.9743	(0.0587)		
Huggins Hospital	0.9846	(0.0472)		
Lakes Region General Hospital	0.7653	(0.1240)		
Littleton Regional Hospital	0.8731	(0.0941)		
Memorial Hospital	0.6336	(0.1637)		
Monadnock Community Hospital	0.5810	(0.1117)		
New London Hospital	0.8801	(0.0930)		
Parkland Medical Center	0.8909	(0.1670)		
Portsmouth Regional Hospital	0.7808	(0.1332)		
Southern New Hampshire Medical Center	0.9962	(0.0127)		
Speare Memorial Hospital	0.9711	(0.0759)		
St Joseph Hospital	0.7059	(0.0590)		
Valley Regional Hospital	0.8177	(0.1368)		
Wentworth Douglass Hospital	0.9993	(0.0028)		
Weeks Medical Center	1.0000	(0.0001)		
Insurer weight on WTP (α, Hundreds)				
Non-CAH	1.3017	(0.1261)		
CAH	1.6695	(0.2485)		
Balance bill recoupment fraction (μ)				
μ	0.3822	(0.2123)		
Bargaining cost (b_m)				
Aetna	0.0021	(0.0045)		
Cigna	0.0026	(0.0171)		
MVP	1,682	(1,265)		
United	0.0002	(0.0007)		
Anthem	0.0035	(0.0099)		
Harvard Pilgrim	0.0001	(0.0003)		
Tufts	2,908,943	(684,294)		

Notes: Bootstrapped standard errors in parentheses. The first two columns report results from the full model, which estimates all bargaining parameters. The second two columns reports results from the restricted specification, which only estimate marginal costs under a simplified bargaining structure that assumes zero disagreement payoffs and volumes.

Figure A.2: Hospital Networks and Enrollees by Health Plan in New Hampshire

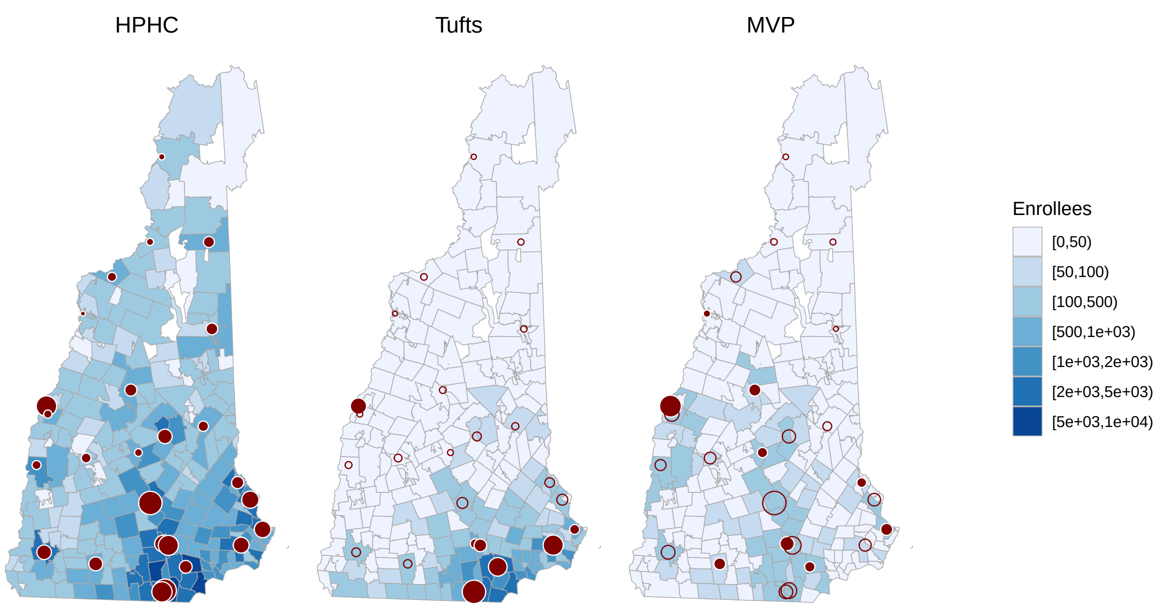

Notes: This figure plots the New Hampshire enrollments of three key insurers at the 5-digit zip code level. Solid red circles represent in-network hospitals. Hollow circles represent out-of-network hospitals. Hospital circles are scaled by the insurer's patient volume flowing to each hospital. For the two narrow-network insurers, Tufts and MVP, in-network hospitals are located near large masses of enrollees.

Figure A.3: Bias in Hospital Cost Estimates Under Zero Disagreement Volumes

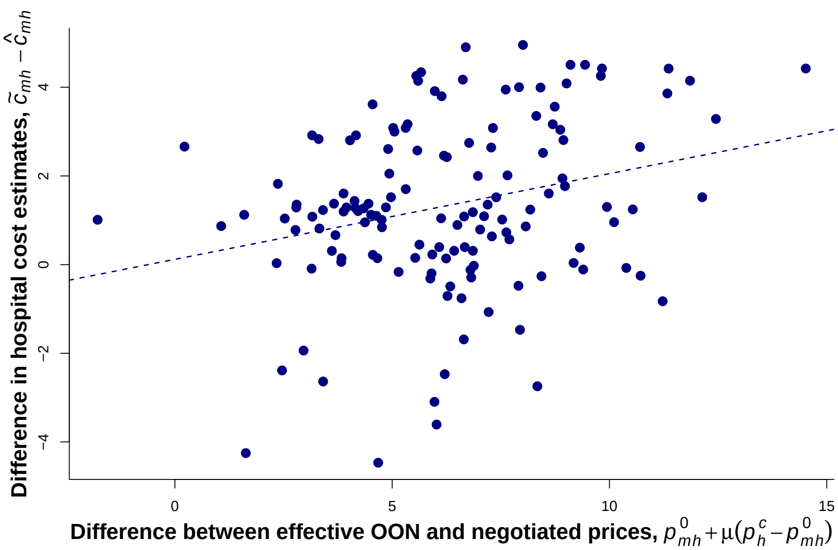

Notes: This figure illustrates the direction of the bias arising from assuming zero disagreement volumes using our estimates for 2011. The horizontal axis is the effective price differential between out-of-newtork and in-network care for a given insurer-hospital pair ($(1-\mu)p^0_{mh} + \mu p^c_h - p_{mh}$). The vertical axis is the difference between the hospital cost estimates from the standard model (assuming zero disagreement volumes) and the hospital cost estimates from the full model.

Figure A.4: Predicted Negotiated Prices Vs. Multiples of Medicare Reimbursements, with Standard Model Estimates

(a) Aetna

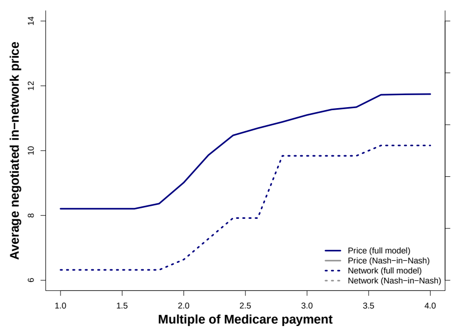

(b) Anthem

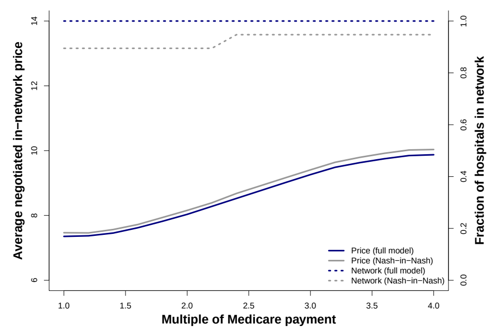

(c) Cigna

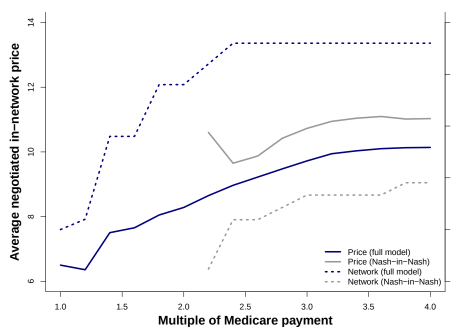

(d) Harvard Pilgrim

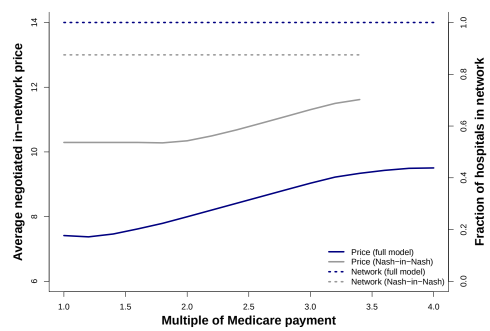

(e) MVP

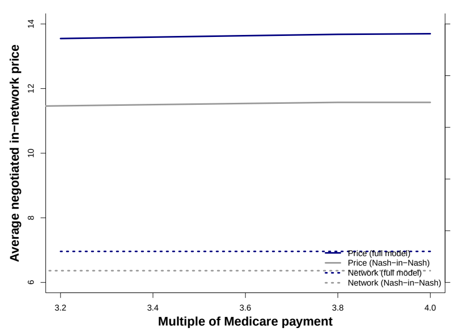

(f) Tufts

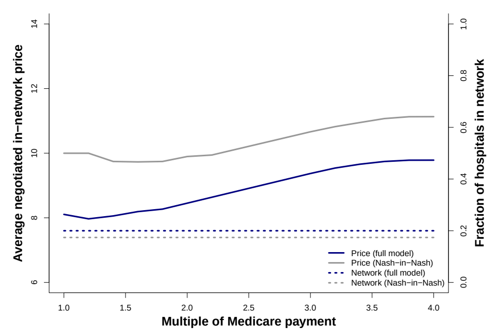

(g) United

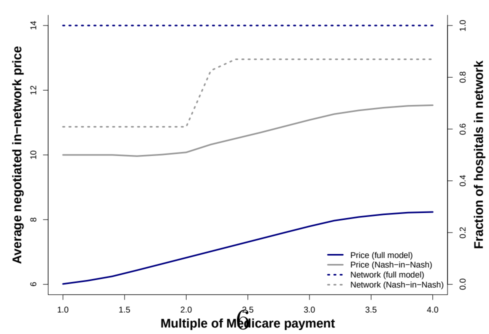

Notes: This figure plots results of counterfactual simulations varying out-of-network prices to be various multiples of Medicare reimbursements. The left vertical axis plots the counterfactual negotiated price for hospitals predicted to be in-network (volume-weighted average prices). The right vertical axis plots the fraction of hospitals predicted to be in-network. Gaps represent counterfactuals for which no equilibrium was found.

Figure A.5: Predicted Negotiated Prices Vs. Multiples of Medicare Reimbursements with Reduced OON Disutility

(a) Aetna

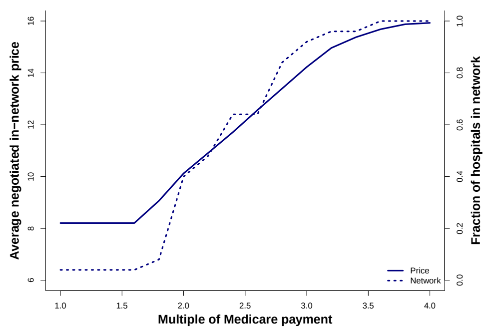

(b) Anthem

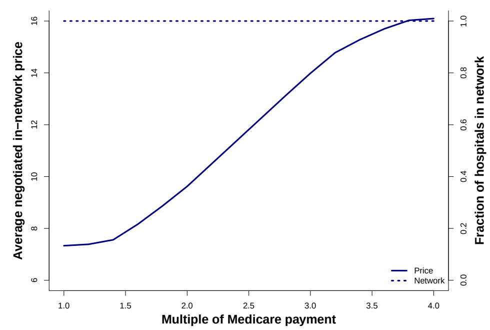

(c) Cigna

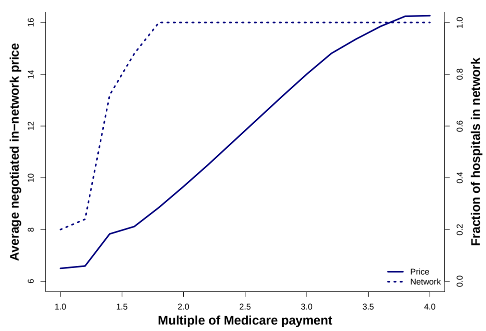

(d) Harvard Pilgrim

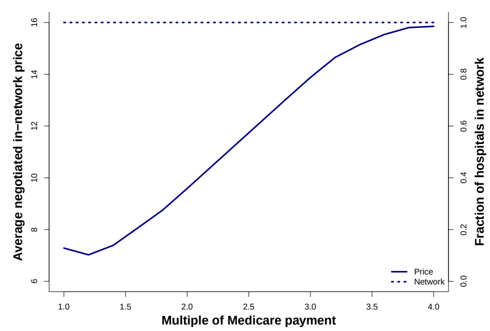

(e) MVP

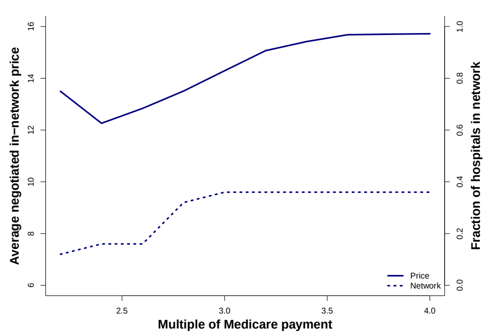

(f) Tufts

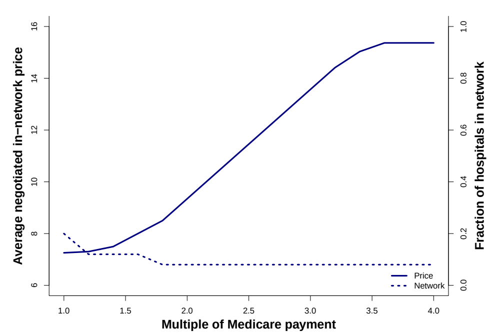

(g) United

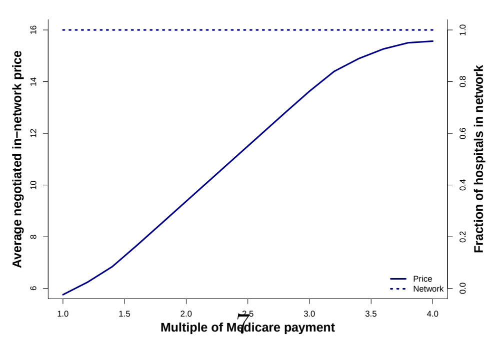

Notes: This figure plots results of counterfactual simulations varying out-of-network prices to be various multiples of Medicare reimbursements, setting the disutility of out-of-network providers to 10% of the estimated value. The left vertical axis plots the counterfactual negotiated price for hospitals predicted to be in-network (volume-weighted average prices). The right vertical axis plots the fraction of hospitals predicted to be in-network. Gaps represent counterfactuals for which no equilibrium was found.

Figure A.6: Distribution of Negotiated Prices Vs. Multiples of Medicare Reimbursements

(a) Aetna

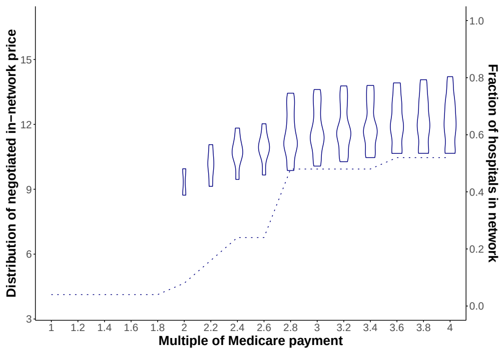

(b) Anthem

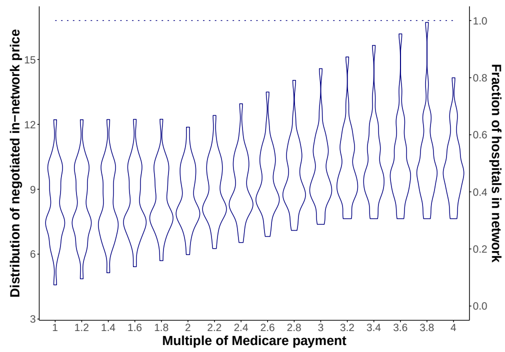

(c) Cigna

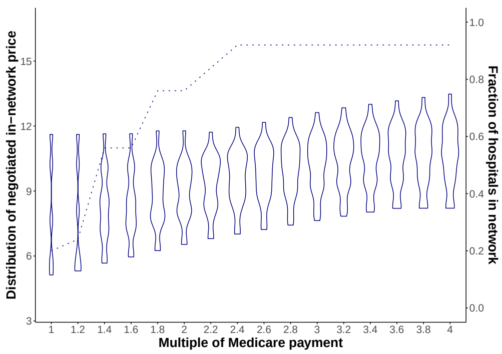

(d) Harvard Pilgrim

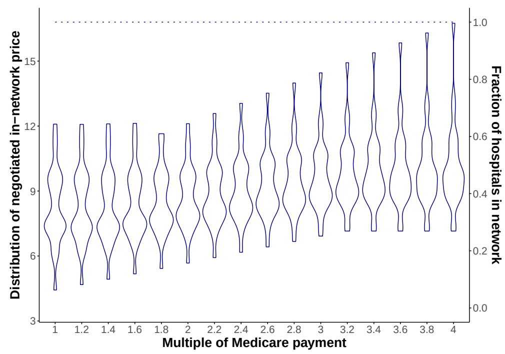

(e) MVP

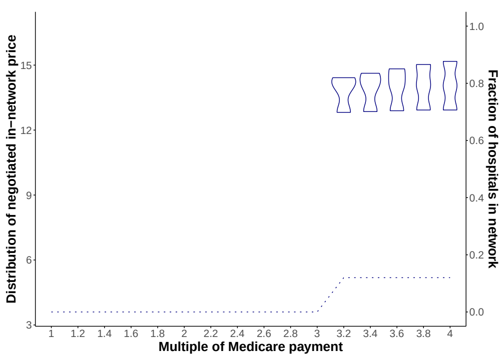

(f) Tufts

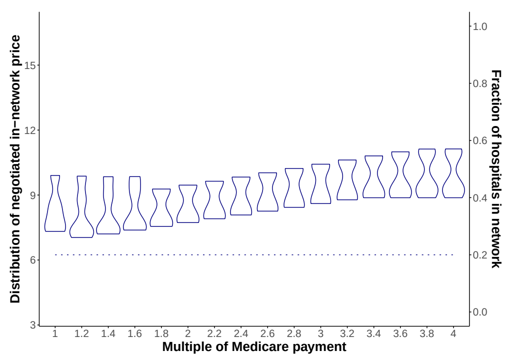

(g) United

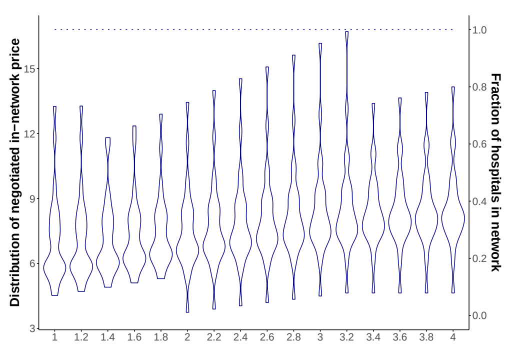

Notes: This figure plots results of counterfactual simulations varying out-of-network prices to be various multiples of Medicare reimbursements. The left vertical axis plots the distribution of in-network prices (corresponding to average in-network prices in Figure 2). The size of each distribution is scaled by the relative number of in-network hospitals within that insurer. Gaps represent counterfactuals for which no equilibrium was found, or too few hospitals were in-network to generate a violin. The dashed line (right vertical axis) plots network breadth.

Figure A.7: Predicted Negotiated Prices in Counterfactuals with Hospital Service Line Closures and Naive Negotiators

(a) Aetna

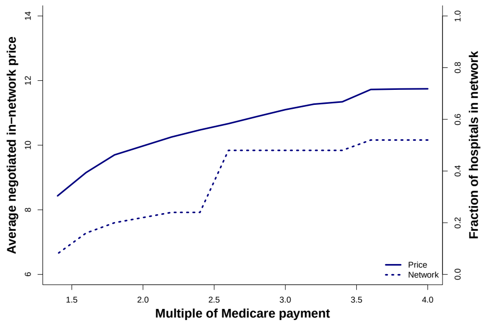

(b) Anthem

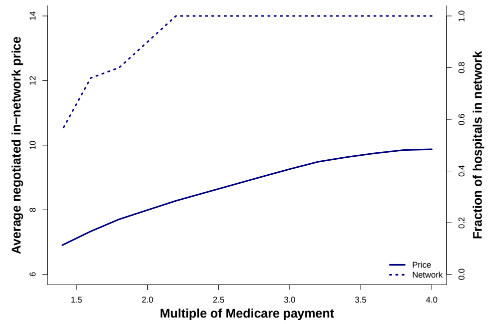

(c) Cigna

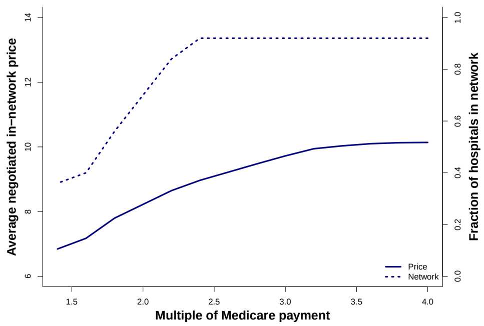

(d) Harvard Pilgrim

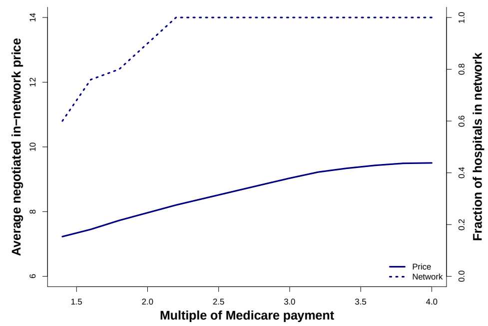

(e) MVP

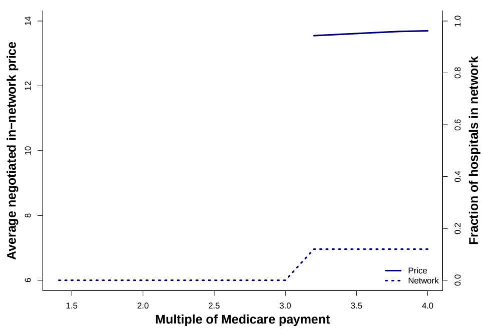

(f) Tufts

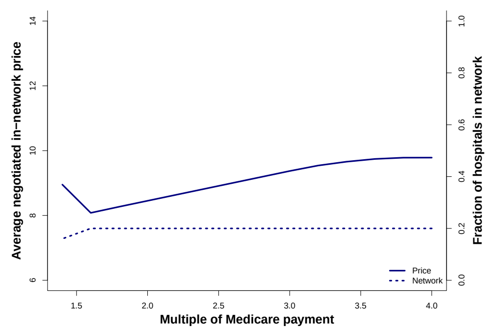

(g) United

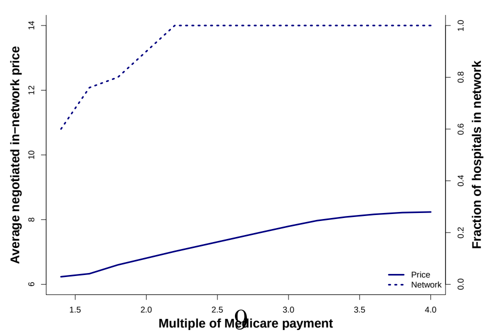

Notes: This figure plots results of counterfactual simulations varying out-of-network prices to be various multiples of Medicare reimbursements, assuming negotiators are naive (corresponding to 3). The left vertical axis plots the counterfactual negotiated price for hospitals predicted to be in-network (volume-weighted average prices). The right vertical axis plots the fraction of hospitals predicted to be in-network. Gaps represent counterfactuals for which no equilibrium was found.

Figure A.8: Predicted Negotiated Prices in Counterfactuals with Hospital Service Line Closures and Sophisticated Negotiators

(a) Aetna

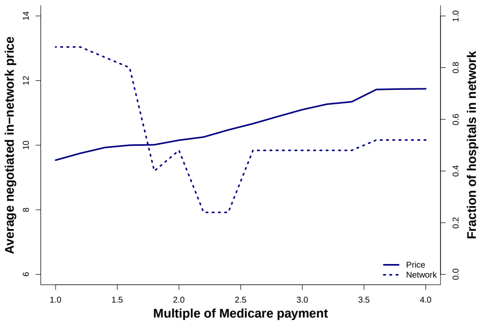

(b) Anthem

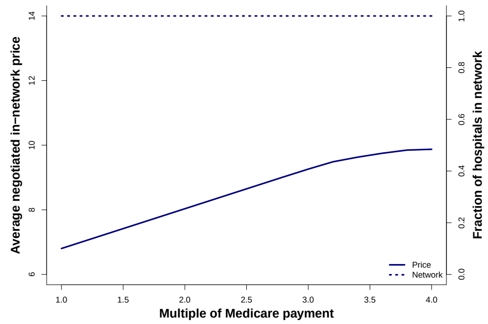

(c) Cigna

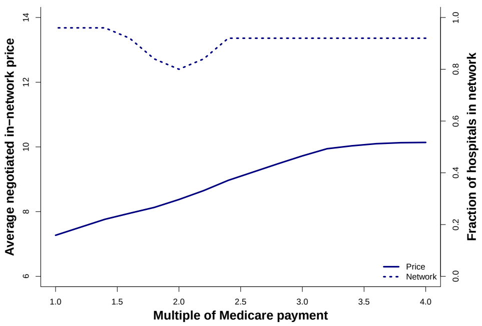

(d) Harvard Pilgrim

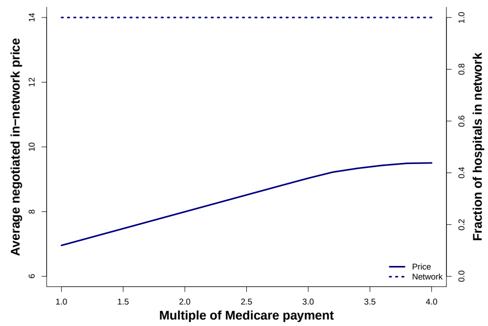

(e) MVP

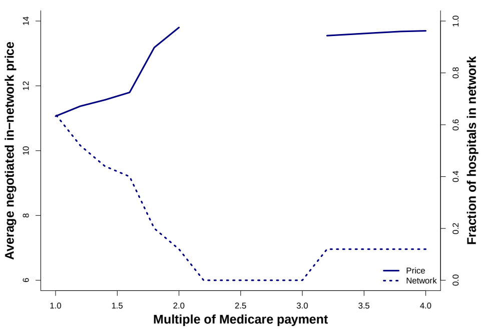

(f) Tufts

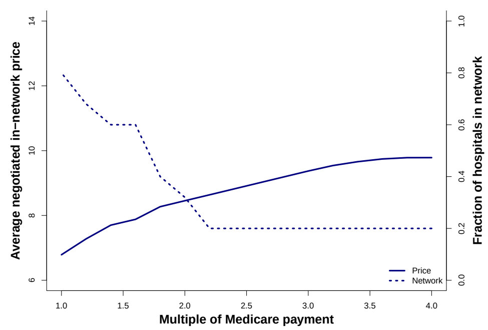

(g) United

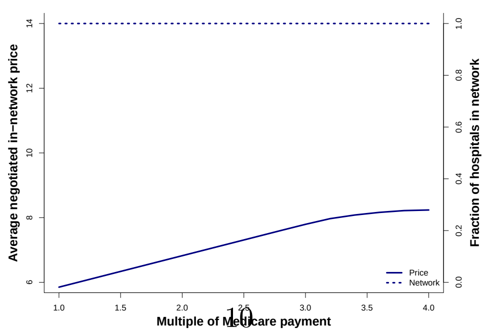

Notes: This figure plots results of counterfactual simulations varying out-of-network prices to be various multiples of Medicare reimbursements, assuming negotiators are sophisticate (corresponding to 4). The left vertical axis plots the counterfactual negotiated price for hospitals predicted to be in-network (volume-weighted average prices). The right vertical axis plots the fraction of hospitals predicted to be in-network. Gaps represent counterfactuals for which no equilibrium was found.

B Constructing Price Benchmarks

This appendix section describes in detail the price benchmarks used to construct the off-contract prices first described in Section 2.3.

B.1 The FAIR Health Algorithm

FAIR Health is the source of charge price benchmarks for many insurers (see Table 3). For each type of health care service, FAIR Health calculates the distribution of charge prices within a geographic region over the course of one year. The geographic regions chiefly correspond to three-digit zip codes, although in low-density areas a handful of three-digit zips might be aggregated into one geographic unit of analysis (typically no more than three, but up to a maximum of twelve). The country is partitioned into 493 such geographic regions. Four of these are in New Hampshire.

FAIR Health has multiple benchmark price products: hospital inpatient benchmarks, based on ICD diagnosis codes or bundled DRG diagnosis codes; hospital outpatient benchmarks, based on CPT procedure codes; anesthesia benchmarks, based on CPT procedure codes; professional services benchmarks, based on HCPCS/CPT codes; and others. As our empirical exercise is limited to outpatient hospital demand, we are interested in the CPT-based benchmarks.

For each CPT code in each geographic unit, FAIR Health starts with all health care claims in that CPT-geography pair. This includes both claims from their large sample of private insurers and the universe of fee-for-service Medicare claims. It then calculates for each claim the absolute distance from the median charge price for that CPT-geography pair. The median of those distances is then computed. Next, extreme outliers are dropped: any claim whose distance from the median charge price is more than 5.92 times the median distance (in either direction) is dropped from the sample. Finally, the remaining claims are used to calculate charge price percentiles within each CPT-geography pair.

The standard FAIR Health benchmark products report the 50th, 60th, 70th, 75th, 80th, 85th, 90th, and 95th percentiles, but insurers can also purchase custom products reporting other quantiles of the distribution. The benchmarks are updated every six months based on a rolling one-year sample of claims. There is a May release based on data from the prior March through the most recent February, and a November release based on data from the prior September through the most recent August.

B.2 Approximating FAIR Health Benchmarks

We show that researchers can approximate the FAIR Health benchmarks well using all-payer claims data. We approximate the outpatient price benchmarks using the near-universe of private insurance claims in the New Hampshire APCD by following FAIR Health's algorithm as possible within. We match the geographic units exactly using FAIR Health's crosswalk between three-digit zip codes and their definition of the four geographic units in New Hampshire. We also match the level of the procedure code by using CPT codes (without modifiers). Figure A.1 shows that the approximation performs well, despite FAIR Health's inclusion of fee-for-service Medicare claims that are absent from our APCD data.

C Constructing Price and Cost Indices

To operationalize the bargaining model from Section 3.1, we adopt from the literature a key simplifying assumption about how prices and marginal costs are scaled. Following Gowrisankaran et al. (2015) and Ho and Lee (2017), we assume that each hospital-insurer pair negotiates a single price index p_{mh} that is then scaled multiplicatively to determine the price for a given diagnosis or service.[36] The multiplicative scaling w_d is based on the resource intensity of the diagnosis or service, so that the price that insurer m pays to hospital h for service d is given by $w_d p_{mh}$. In our empirical application, this becomes a weaker assumption, requiring that prices are scaled in this manner only for the relatively narrow range of services we consider. We make the same scaling assumption about hospital marginal costs c_{mh}, as in those papers. This makes the Nash bargaining first-order conditions in Equation 3 linear in hospital marginal costs.

Existing work on hospital-insurer bargaining has generally restricted the analysis to inpatient hospital care. In an inpatient setting, a natural choice for the resource weights w_d is DRG weights, which are weights specifically designed to measure the relative resource intensity of various types of inpatient care. Since our analysis focuses instead on outpatient hospital care, we turn to a different measure of w_d. We select a measure that achieves internal consistency with our algorithm for measuring off-contract prices, described in Section 2.3: the FAIR Health charge benchmark percentiles. For each procedure in the sample, we assign a measure of resource intensity constructed from FAIR Health weights. We normalize the weights such that $w_d = 1$ for venipuncture (i.e., a blood draw; CPT code 36415), chosen because it is both common and a fairly uniform procedure. Thus, the prices and costs we report should be scaled by the resource intensity of a given type of care relative to the resource intensity of venipuncture. We have validated our price measure against DRG-deflated inpatient prices for the same hospital-insurer pairs, and found similar patterns over time across the two price measures.

In the demand model, the weights act as a measure of severity, allowing patients seeking care for more complex bundles of care to choose different providers than otherwise similar patients who require simpler care. In the bargaining model, the weights serve as a multiplier for scaling hospital volumes. A higher weight w_d will increase both the hospital's marginal cost of treating a patient and the price paid to the hospital by the same multiplicative factor. Figure H.1, described in Section 3.6, plots the price indices computed for the five insurers with complete hospital networks

[36] Other papers making analogous assumptions include Shepard (2022), Ghili (2022) and our own work in Prager (2020) and Tilipman (2022).

in New Hampshire: Anthem, Harvard Pilgrim, Cigna, United, and Aetna. The price indices reflect negotiated prices for procedures with the resource intensity of a routine venipuncture. Negotiated prices fall primarily in the \$6 to \$13 range.

D Constructing Average Variable Costs

Identification of the bargaining model relies on moments that are a function of $\bar{c}_h^o$, the average variable cost per outpatient. We follow existing work on measuring hospital costs to construct a suitable measure. For each hospital-year, we follow Schmitt (2017) and Ghili (2022) to calculate total variable costs using data from the Centers for Medicare and Medicaid's (CMS) Healthcare Cost Report Information System (HCRIS). Hospitals are required to report the elements of their cost structure to CMS.

To obtain a per-outpatient measure of average variable costs, we need the volume of commercially insured outpatients per hospital-year (scaled by w_d) and the fraction of total variable costs that is attributable to outpatient services for commercially insured patients. We measure the outpatient volumes from the APCD data. This is a departure from the approach in the literature, which measures volumes directly from HCRIS reports. However, unlike for the inpatient volumes targeted in the literature, for outpatient volumes HCRIS reports include only the count of patients without allowing scaling by severity.

Our measure of variable cost attribution requires several steps. We follow the literature (Schmitt (2017), Ghili (2022), and others) in assuming that the fraction of total costs attributable to either overall inpatient volume or overall outpatient volume is equal to the revenue share of inpatient or outpatient revenues, respectively. This gives us a measure of the total variable cost attributable to all outpatient volume. We further assume that the share of the outpatient cost attributable to commercially insured outpatients is equal to commercially insured patients' share of all inpatient discharges. (We use inpatient rather than outpatient volume shares due to data constraints. Neither the HCRIS data nor alternative data sources such as the American Hospital Association Annual Survey report the fraction of outpatient volume coming from each payer type.)

Finally, we divide the total variable cost attributed to commercially insured outpatients by the total volume of commercially insured outpatients to arrive at $\bar{c}_h^o$, our measure of average variable cost per outpatient.

E Outpatient Health Care Services Sample

Our analytical sample consists of health care services that are performed in an outpatient, rather than inpatient, setting. We restrict our sample to outpatient procedures that plausibly constitute the primary reason for a patient's choice of provider. This requires dropping procedure codes that are incidental to the main treatment or procedure. We drop the following classes CPT codes:

pathology and laboratory services (codes beginning with 8 or P); codes specific to the emergency department (codes 99281–99288); anesthesia (codes 00100–01999, 99100–99150); modifier codes for visits or services that are already reported separately (Category III CPT codes); temporary codes for emerging technologies (Category III CPT codes); ambulance and other transportation (codes beginning with A); durable medical equipment (codes beginning with E or K); dental procedures (codes beginning with D); and other temporary and miscellaneous codes (codes beginning with Q or S). The vast majority of volume among the dropped categories belongs to pathology and laboratory services. We refer to the remaining CPT codes as "primary" procedure codes.

We subset the primary procedure codes to the top 1,000 codes by hospital revenue. These top 1,000 codes account for 96.7 percent of hospital outpatient revenue and 98.5 percent of hospital outpatient volume among primary codes. The top ten of these codes, which account for 17.2 percent of revenue and 65.4 percent of volume, are dominated by generic visit codes and diagnostic procedures. Two of them are outpatient or physician office visits by established patients; six are the diagnostic procedures of diagnostic colonoscopies, head MRI scans, mammograms, echocardiograms, abdominal CT scans, and biopsies of the upper digestive tract; one is the injection of the drug infliximab, which is used to treat autoimmune conditions including arthritis and Crohn's disease; and one is physical therapy exercises.

We make some additional sample restrictions to construct our final sample for the demand model. First, we limit the data to only patients insured by Anthem/BCBS, Harvard Pilgrim, Cigna, United, Aetna, MVP, or Tufts. These insurers collectively account for 97 percent of commercial health insurance enrollment in New Hampshire in 2012. The top three, Anthem, Harvard Pilgrim, and Cigna, account for 89 percent. Although our primary focus is on New Hampshire, we observe patients who reside in Massachusetts who cross the border to seek care in New Hampshire. Similarly, we observe patients residing in New Hampshire who seek care in Massachusetts. We include in the sample all enrollees who live in New Hampshire as well as those who live in Massachusetts near the New Hampshire border. Specifically, we include any enrollee living in any Massachusetts zip code within the 75th percentile of distance traveled to a New Hampshire hospital. For every enrollee, we include in the choice set all 26 acute care hospitals in New Hampshire, as well as any Massachusetts hospital within the 75th percentile of distance traveled from any border zip code. The final choice set consists of 40 hospitals, 26 from New Hampshire and 14 from Massachusetts.

For computational feasibility, we take a random sample of approximately 51,000 visits (a 0.5 percent sample) to use in the maximum likelihood estimation. Visits to hospitals or hospital-affiliated providers comprise 31.5 percent of the sample; the majority of the sample consists of visits to standalone physician offices and other providers not affiliated with a hospital. The final sample, including both hospital visits and non-hospital visits, consists of visits from Anthem/BCBS at 55.3 percent, Harvard Pilgrim at 16.9 percent, Tufts at 10.9 percent, Cigna at 8.0 percent, United at 4.7 percent, Aetna at 3.4 percent, and MVP at 0.8 percent. Table E.1 shows the summary statistics for our demand estimation sample, subsetting just to hospital visits. Anthem's share is higher, 61.3 percent, when subsetting to hospital visits. Tufts' sample share is substantially larger than its New Hampshire market share because Tufts is among the top three insurers in Massachusetts, from which we also draw data. Among visits to hospitals and their affiliated providers, 68.0 percent are to New Hampshire hospitals and the remaining 32.0 percent are to Massachusetts hospitals.

Table E.1: Demand Sample Summary Statistics

	Full Sample		NH Narrow Plans	
	Mean	SD	Mean	SD
Patient Characteristics				
Age	49.117	20.049	43.130	17.080
Female	0.596	0.491	0.578	0.494
FAIR Health Weight	63.778	198.095	104.679	306.810
Anthem/BCBS	0.622	0.485	0.000	0.000
Cigna	0.142	0.350	0.000	0.000
Harvard Pilgrim	0.121	0.326	0.000	0.000
Tufts	0.026	0.161	0.635	0.482
United	0.022	0.147	0.000	0.000
Aetna	0.050	0.218	0.000	0.000
MVP	0.015	0.122	0.365	0.482
Distance (Miles)	11.350	15.185	11.487	15.318
Hospital Characteristics				
In Network	0.989	0.104	0.736	0.441
Beds	207.653	140.381	214.719	169.454
Cath Lab	0.847	0.360	0.826	0.379
NICU	0.397	0.489	0.331	0.471
Neuro	0.909	0.288	0.857	0.350
MRI	0.850	0.357	0.827	0.378
Critical Access	0.131	0.338	0.089	0.285
Teaching	0.174	0.379	0.142	0.349

Notes: Summary statistics for demand estimation sample of 2011–2012 outpatient visits. Columns 1–2 summarize the full sample of hospital visits, including Massachusetts residents near the New Hampshire border as well as New Hampshire residents. Columns 3–4 subset to enrollees in New Hampshire-based insurance plans with narrow networks (Tufts and MVP).

In the demand estimation sample, only 1 percent of visits are to out-of-network providers. This is because the majority of visits in the sample are by patients who have access to complete networks; these are all to in-network providers. When subsetting to patients insured through narrow-network insurance plans, 26.4 percent of hospital visits are to out-of-network hospitals. These patients travel somewhat smaller distances to receive care: the mean distance traveled to an out-of-network hospital is 10.5 miles, compared to 16.9 miles to in-network hospitals. Patients appear to trade off disutility from seeking care out of network against the utility of traveling a shorter distance. We leverage this type of variation for estimating the demand model. The out-of-network hospitals comprising these visits are smaller than hospitals comprising in-network visits, with a mean bed count of 160.4 compared to 217.8. The intensity of care obtained at out-of-network hospitals is similar to that in in-network hospitals.

To estimate the bargaining model, we construct the quantities in Section 3.5 from a random sample of 5,000 households. For each household member, we merge in the annual prevalence of health care services in each decile of severity. Severity is measured by weights w_d, described in Appendix C. Severity quantiles and their associated prevalences are calculated from all commercially insured patients in the New Hampshire APCD, separately by sex and five-year age band. We use these prevalences to construct, for each patient, the predicted WTP and hospital volumes under each network configuration. In the bargaining estimation, each patient from the 5,000-household sample is then scaled up by the appropriate sample weight to represent the insurers' complete enrollment panel.

F Hospital Choice

The bargaining model in Section 3 relies on estimates from a model of hospital demand. This section describes the underlying demand estimation, which follows what is now standard in the literature.

Consumers enrolled in health insurance get sick and require health care with some probability. A consumer insured by insurer m and needing procedure d gets the following utility from seeking outpatient care at hospital h (for convenience, we omit time subscript t):

$$u_{imhd} = \lambda_h + \delta\eta_{mh} + \beta x_{ihd} + \varepsilon_{imhd} \tag{11}$$

where λ_h are hospital fixed effects (separate for the hospital's main campus or a secondary location), η_{mh} is an indicator for whether hospital h is in insurer m's network, and x_{ihd} is a vector of observable characteristics of the patient and the hospital. x_{ihd} includes the distance between consumer i's home and hospital h, hospital characteristics, such as its teaching status, patient demographics (in our setting, age, weights of the procedure w_d, and gender), the expected balance bill (calculated as the hospital's charge price less the insurer's out-of-network paid price), and interactions between patient characteristics and service availability at hospital h. Here, d is defined at the level of specific medical procedures (CPT codes). We proxy for its severity with the weights constructed from FAIR Health data for the particular procedure, as described in Appendix C. If consumers prefer to seek care at in-network hospitals, we expect a positive coefficient estimate δ for the in-network indicator. We do not include a finer measure of in-network out-of-pocket price in x_{ihd} because consumers in most plans are not subject to the type of out-of-pocket price structure that results in price-shopping (Prager 2020). The error term ε_{imhd} is assumed to be Type 1 Extreme Value, yielding a discrete choice multinomial logit structure.

We estimate the hospital demand model using maximum likelihood and use it to construct the inputs to the bargaining model. The coefficient on the in-network indicator η_{mh} plays a key role in constructing the expected hospital volumes and willingnesses-to-pay under the observed and unobserved network configurations used to estimate the bargaining model. We therefore instrument for the in-network hospital hospital indicator to address concerns that insurers may differentially include in their networks the hospitals for which their enrollees have an unobservable preference. For example, if insurers differentially include unobservably higher-quality hospitals and the quality heterogeneity is not adequately captured by the hospital fixed effects, then the coefficient on the in-network indicator will be biased upward. We leverage a feature of this market shown in Figure A.2: insurer networks are highly correlated with the geographic distribution of their enrollees. This is likely due to the fact that the fixed cost of entering negotiation in Equation 4 is only recouped when a sufficient number of enrollees highly value access to the hospital. We therefore instrument for network status using the logarithm of the count of enrollees within 20 miles of the hospital. (The 20-mile radius produces the strongest first stage of distances we have tested.)

Because the multinomial logit second stage is nonlinear, we use a control function approach. The control function corrects for the correlation between the endogenous regressor and the error term by approximating the component of the error that is correlated with the endogenous regressor and including it as a separate regressor. In practice, in the first stage, we regress the in-network hospital indicator on the exogenous variables and the enrollee count "instrument," and include the residuals from this first-stage regression in the second-stage multinomial logit. Table F.1 shows the results. Our preferred specification is the control function with the first-degree polynomial. Higher-order polynomial terms are not statistically significant, and the model fit does not improve with their addition.

The demand specification in Equation 11 yields a probability that hospital h is chosen that is given by:

$$\sigma_{imhd} = \frac{\exp\left(\lambda_h + \delta\eta_{mh} + \beta x_{ihd}\right)}{\sum_j \exp\left(\lambda_j + \delta\eta_{mj} + \beta x_{ijd}\right)}$$

where j enumerates the set of all hospitals available to patients (all New Hampshire hospitals and 14 Massachusetts hospitals, as discussed in Section 2.4).

The predicted shares σ_{imhd} from the demand model are used to construct an insurer's volume of patients for each hospital, used in the bargaining model (Equation 3). If hospital h is in insurer m's network, its predicted volume is given by

$$\sigma^1_{mh} = \sum_{i \in I_m} \sum_d w_d f_{id} \sigma_{imhd}$$

Table F.1: Hospital Demand Estimates

	(1) No IV	(2) IV Deg 1	(3) IV Deg 2	(4) IV Deg 3	(5) IV Deg 4	(6) IV Deg 5
In-Network Hospital	0.0858	1.9439***	1.0838	1.0161	1.0146	0.0581
	(0.1190)	(0.2032)	(0.6267)	(0.6213)	(0.6578)	(1.0479)
Other In-Network Provider	2.8057***	2.9066***	2.9088***	2.9084***	2.9078***	2.9078***
	(0.2541)	(0.3565)	(0.3565)	(0.3564)	(0.3563)	(0.3617)
Balance Bill ($)	-0.0005	-0.0003	-0.0005	-0.0005	-0.0005	-0.0006
	(0.0005)	(0.0015)	(0.0016)	(0.0017)	(0.0017)	(0.0017)
Distance (miles)	-0.3023***	-0.3026***	-0.3024***	-0.3024***	-0.3024***	-0.3025***
	(0.0134)	(0.0126)	(0.0126)	(0.0126)	(0.0126)	(0.0125)
Distance^2	0.0006***	0.0006***	0.0006***	0.0006***	0.0006***	0.0006***
	(0.0000)	(0.0000)	(0.0000)	(0.0000)	(0.0000)	(0.0000)
Distance × Age	-0.0001**	-0.0001**	-0.0001**	-0.0001**	-0.0001**	-0.0001**
	(0.0000)	(0.0000)	(0.0000)	(0.0000)	(0.0000)	(0.0000)
Distance × Intensity Weight	-0.0000*	-0.0000**	-0.0000**	-0.0000**	-0.0000**	-0.0000**
	(0.0000)	(0.0000)	(0.0000)	(0.0000)	(0.0000)	(0.0000)
Beds × Intensity Weight	0.0000***	0.0000***	0.0000***	0.0000***	0.0000***	0.0000***
	(0.0000)	(0.0000)	(0.0000)	(0.0000)	(0.0000)	(0.0000)
Critical Access × Intensity Weight	-0.0008	-0.0007	-0.0007	-0.0007	-0.0007	-0.0007
	(0.0004)	(0.0005)	(0.0005)	(0.0005)	(0.0005)	(0.0005)
Teaching × Intensity Weight	-0.0009***	-0.0009***	-0.0009***	-0.0009***	-0.0009***	-0.0009***
	(0.0003)	(0.0002)	(0.0002)	(0.0002)	(0.0002)	(0.0002)
Cath Lab × Intensity Weight	0.0004	0.0004	0.0004	0.0004	0.0004	0.0004
	(0.0003)	(0.0003)	(0.0003)	(0.0003)	(0.0003)	(0.0003)
MRI × Intensity Weight	-0.0000	-0.0000	-0.0000	-0.0000	-0.0000	-0.0000
	(0.0001)	(0.0001)	(0.0001)	(0.0001)	(0.0001)	(0.0001)
NICU × Intensity Weight	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000
	(.)	(0.0000)	(0.0000)	(0.0000)	(0.0000)	(0.0000)
Neuro × Intensity Weight	-0.0002	-0.0002	-0.0002	-0.0002	-0.0002	-0.0002
	(0.0002)	(0.0002)	(0.0002)	(0.0002)	(0.0002)	(0.0002)
1st Stage Residual^1		-0.2963***	-0.2054**	-0.1872**	-0.1864	-0.7355
		(0.0255)	(0.0681)	(0.0714)	(0.1388)	(4.9196)
1st Stage Residual^2			-0.1126	-0.0055	-0.0026	-7.1032
			(0.0843)	(0.1788)	(1.0514)	(55.0185)
1st Stage Residual^3				0.1169	0.1175	-7.3384
				(0.1412)	(0.9159)	(39.5525)
1st Stage Residual^4					-0.0161	-1.9177
					(0.1014)	(2.8778)
1st Stage Residual^5						-0.7387
						(0.5015)
Hospital FEs	Yes	Yes	Yes	Yes	Yes	Yes
Pseudo R^2	0.557	0.558	0.558	0.558	0.558	0.558
Choices	18616	18616	18616	18616	18616	18616

Notes: ***p<0.01, **p<0.05, *p<0.10. Results from multinomial logit provider choice model from years 2011–2012. IV columns estimated using a control function with boostrapped standard errors with 100 replications. *Choices* is the number of choice sets used in estimation. *In-Network Hospital* and *Other In-Network Provider* are indicators for whether the provider is in the patient's insurer's network and is a hospital or non-hospital provider, respectively. *Balance Bill* is the potential (maximum) dollar amount a patient would be charged out-of-pocket in cases where the hospital is out of network.

where f_{id} is the probability that a consumer of type i requires care for procedure d over the course of a plan-year.[37] The term w_d is the resource utilization multiplier used to construct a weighted sum of hospital volume. The terms $\sigma^0_{mh}, \psi^1_{mh}, \psi^0_{mh}$ are defined analogously. These enter into the insurer's bargaining surplus (Equation 2) and the hospital's bargaining surplus (Equation 1) and are used for estimating the bargaining model.

Consumers' expected utility from insurer m's network also enters into the bargaining model. This expected utility is a function of the probability of getting sick and needing care, the set of hospitals that are in the network, and the strength of the preference for in-network hospitals. We denote an individual consumer's expected utility for insurer m's network as

$$W_{im} = \sum_d f_{id} \log \left(\sum_j \exp\left(\lambda_j + \delta\eta_{mj} + \beta x_{ijd}\right) \right)$$

The W_{im} terms are summed across an insurer's enrollees to obtain the insurer-wide expected utility of a network that enters into the insurer's bargaining surplus, as defined in Equation 2. When hospital h is in the network, this becomes

$$W^1_{mh} = \sum_{i \in I_m} \sum_d f_{id} \log \left(\exp\left(\lambda_h + \delta \cdot 1 + \beta x_{ihd}\right) + \sum_{j \neq h} \exp\left(\lambda_j + \delta\eta_{mj} + \beta x_{ijd}\right) \right)$$

and W^0_{mh} is defined analogously when the hospital is out of network.

G Detailed Implications of Nonzero Disagreement Values

Empirical work on bargaining typically observes negotiated prices as an equilibrium outcome, and uses them to infer a set of structural parameters pertaining to costs (marginal or fixed) and Nash bargaining weights. Misspecification of the disagreement payments $p^0_m \sigma^0_{mh}$ biases these structural parameters. Any of the parameters may be biased due to misspecification of the disagreement payments. To ease interpretation, it is therefore helpful to consider all but one set of parameters as being fixed at known values, and derive the bias on one set of free parameters. Here, we illustrate the bias arising from assuming that disagreement volume is zero when estimating hospital marginal costs c_{mh} and fixing all other parameters. Biases in marginal cost estimates are of interest to antitrust regulators due to their direct relationship with margins, which are used in merger evaluation.

Consider a simplified empirical setup where the parameters γ_h, α, μ, and $b_m = 0$ are known, leaving only the hospital costs c_{mh} as parameters to estimate. Take an insurer m that has a negotiated contract with hospital h. An expression for the unbiased estimate, $\hat{c}_h$, can be obtained

[37] In specifying f_{id}, we allow for individual consumers to require procedure d more than once in a plan-year.

by rearranging Equation 3:

$$\hat{c}_h = \frac{p^*_{mh}\sigma^1_{mh} - (1-\gamma_h)\,\alpha\left(W^1_{mh} - W^0_{mh}\right) + (1-\gamma_h)\left(\psi^1_{mh} - \psi^0_{mh}\right) - (1-\gamma_h\mu)\,p^0_m\sigma^0_{mh} - \gamma_h\mu p^c_h\sigma^0_{mh}}{\gamma_h\left(\sigma^1_{mh} - \sigma^0_{mh}\right)}$$

If disagreement volume is instead assumed to be zero, then we will obtain a biased estimated of hospital marginal cost, $\tilde{c}_h$:

$$\tilde{c}_h = \frac{p^*_{mh}\tilde{\sigma}^1_{mh} - (1-\gamma_h)\,\alpha\left(\tilde{W}^1_{mh} - \tilde{W}^0_{mh}\right) + (1-\gamma_h)\left(\tilde{\psi}^1_{mh} - \tilde{\psi}^0_{mh}\right)}{\gamma_h\tilde{\sigma}^1_{mh}}$$

where the tilde notation represents quantities calculated from the demand model assuming zero volumes for all out-of-network hospitals. If the insurer has an incomplete network that excludes at least one other hospital $h' \neq h$, then $\tilde{\sigma}^1_{mh} \neq \sigma^1_{mh}$, $\tilde{W}^1_{mh} \neq W^1_{mh}$, and $\tilde{\psi}^1_{mh} \neq \psi^1_{mh}$. That is, the quantities corresponding to hospital h being in the insurer's network under the model assuming zero disagreement volumes depart from the full model.

Regardless of the network configuration, as long as the true disagreement volumes are not equal to zero, the estimated costs will not be equal across the two models. The estimated hospital cost under the assumption of zero disagreement values will be biased upward, i.e. $\tilde{c}_h > \hat{c}_h$, if and only if:

$$\begin{gathered}\alpha\left(W^1_{mh} - W^0_{mh}\right) - \left(\psi^1_{mh} - \psi^0_{mh}\right) + \frac{(1-\gamma_h\mu)\,p^0_m + \gamma_h\mu p^c_h - p^*_{mh}}{1-\gamma_h}\sigma^0_{mh} \\ > \\ \frac{\left(\sigma^1_{mh} - \sigma^0_{mh}\right)}{\tilde{\sigma}^1_{mh}}\left[\alpha\left(\tilde{W}^1_{mh} - \tilde{W}^0_{mh}\right) - \left(\tilde{\psi}^1_{mh} - \tilde{\psi}^0_{mh}\right)\right]\end{gathered} \tag{12}$$

This inequality states that hospital cost estimates will be biased upward if the true payments from the insurer to the hospital in the event of disagreement are "large enough." It is a necessary and sufficient condition for upward bias. For ease of exposition, we present the underlying intuition by discussing two comparative statics rather than the inequality as a whole. The discussion is presented in terms of conditions for the hospital cost estimate being biased upward due to an assumption of zero disagreement volumes, because this is what we believe to be more common empirically; the statements hold in reverse for downward bias.

The first point to note is that, all else equal, the larger is the true out-of-network volume σ^0_{mh}, the larger the upward bias on the hospital cost estimate.[38] To see this, note that the right-hand

[38] This is more easily seen when the inequality is rewritten with respect to the ratio of volume gains:

$$\frac{\left(\sigma^1_{mh} - \sigma^0_{mh}\right)}{\tilde{\sigma}^1_{mh}} < \left[\alpha\left(W^1_{mh} - W^0_{mh}\right) - \left(\psi^1_{mh} - \psi^0_{mh}\right) + \frac{(1-\gamma_h\mu)\,p^0_m + \gamma_h\mu p^c_h - p^*_{mh}}{1-\gamma_h}\sigma^0_{mh}\right]\left[\alpha\left(\tilde{W}^1_{mh} - \tilde{W}^0_{mh}\right) - \left(\tilde{\psi}^1_{mh} - \tilde{\psi}^0_{mh}\right)\right]$$

side of the inequality is scaled by $\left(\sigma^1_{mh} - \sigma^0_{mh}\right)/\tilde{\sigma}^1_{mh} \in [0,1]$. This is the ratio of the hospital's true volume gain from being in-network to its volume gain under the assumption of zero out-of-network volume. The true out-of-network volume σ^0_{mh} also appears on the left-hand side of the inequality (multiplied by a positive constant). The underlying intuition is that, when the true out-of-network volume is large, the hospital's true disagreement value is also relatively large as a result of out-of-network payments. The assumption of zero disagreement volumes therefore overstates the hospital's true surplus from agreement. As a result, the surplus implied by the observed negotiated price p^*_{mh} must instead be rationalized by a high cost estimate $\tilde{c}_h$.

The second comparative static is that, all else equal, the higher is the out-of-network price, the larger the upward bias on the hospital cost estimate. The last term on the left-hand side of the inequality is the product of the true out-of-network volume σ^0_{mh} and the difference between the out-of-network price p^0_m and the negotiated in-network price p^*_{mh}, scaled by the inverse of the hospital's Nash bargaining weight.[39] If the out-of-network price is higher than what the negotiated price would be under agreement, as is typically the case in practice, then this term is positive. The underlying intuition is analogous to the previous paragraph. The higher is the true out-of-network price, the more severely the assumption of zero disagreement volumes will overstate the hospital's true surplus. The surplus implied by the observed price must instead be rationalized by a high cost estimate.

The remaining terms in the inequality measure the insurer's surplus from including hospital h in the network, modulo the change in spending on that hospital itself. On the left-hand side, the term $\alpha\left(W^1_{mh} - W^0_{mh}\right) > 0$ is enrollees' willingness-to-pay to include hospital h in the network, scaled by its contribution to the insurer's surplus. This term is typically smaller than its right-hand side analog $\alpha\left(\tilde{W}^1_{mh} - \tilde{W}^0_{mh}\right)$, because the WTP gain from an included hospital is smaller when consumers can still seek care at that hospital even if it is out of network. The term $\left(\psi^1_{mh} - \psi^0_{mh}\right)$ is the change in the insurer's payments to other hospitals $h' \neq h$ as a result of including h and having consumers re-sort across hospitals. This difference is negative, because hospital h loses volume to other hospitals when it is out of network. The terms in brackets on the right-hand side of the inequality are the analogs calculated under the assumption of zero volume for all out-of-network hospitals (scaled by the the volume gain ratio discussed above). The calculated savings in payments to other hospitals will be larger on the right-hand side, $\tilde{\psi}^0_{mh} - \tilde{\psi}^1_{mh} > \psi^0_{mh} - \psi^1_{mh}$, because the assumption of zero disagreement volume will mean there is more of hospital h's volume to be reallocated to other hospitals in the event of disagreement.[40] Therefore, upward bias in the hospital cost estimates cannot result from the WTP and other-hospital savings alone, as the sum of these terms is greater on the right-hand side.

[39] Recall that γ_h is the insurer's bargaining weight, and the two parties' weights sum to one.

[40] Both components of the $\left(\tilde{\psi}^1_{mh} - \tilde{\psi}^0_{mh}\right)$ term may depart from their right-hand side analogs, because if any hospitals besides h are out-of-network, then the right-hand-side will assume they have zero volumes.

Instead, the assumption of zero disagreement volumes will only bias the hospital cost estimates upward if the true disagreement payments are large enough to reverse the inequality. As discussed above, this can obtain from a combination of large out-of-network volumes σ^0_{mh} and high out-of-network prices p^0_m. In our setting, it is usually the case that $p^0_m > p^*_{mh}$ and $\sigma^0_{mh} > 0$, so we expect the majority of the cost estimates to be biased upward under the model that assumes zero disagreement volumes. Figure A.3 shows how this expectation plays out in the data.

Bias in hospital cost estimates has important implications for counterfactual exercises. When cost estimates are biased upward, counterfactual simulations of policies whose goal is to reduce negotiated prices will understate the true magnitude of price reductions. This arises from an understatement of true hospital markups due to the upward-biased cost estimates. The downward-biased estimate of hospital markups gives the impression that there is little room to reduce prices without inducing hospital exit. Moreover, if policy-makers rely on economists' estimates of markups, they may craft policies that erroneously assume hospitals are capturing little producer surplus.[41] In Section 5.2, we show how the biased cost estimates affect the predicted effects of counterfacual policies.

H Bargaining Model Identification Details

This section discusses identification of the bargaining model parameters that were not covered in Section 3.6.

Marginal Costs: As discussed in the main text, identification of hospital marginal cost parameters, λ, and bargaining weights, γ_h, relies primarily on the two sets of equality moments presented in Section 3.5, leveraging variation in observed negotiated prices and hospital average costs. Each set of moments helps discipline parameters implied by the other set, and this interaction between model-implied costs from the first-order conditions (Equation 3.5) and average cost moments (Section 3.5) is essential for separating estimates of marginal costs from bargaining power. In particular, without the average cost moments, high observed prices could be rationalized either by high marginal costs or high ability of a hospital to extract surplus.

Marginal cost parameters, λ, are identified through two channels. First, the marginal cost projections in Equation 8 must align with observed cost data, allowing hospital characteristics to explain patterns in average costs. Second, conditional on bargaining weights, these cost projections must match implied costs from the first-order conditions. While the relationship between average costs and hospital characteristics helps pin down the majority of the cost parameters, this second channel serves two roles: it provides additional variation in the relationship between prices and hospital characteristics across hospitals, and it helps identify insurer fixed effects in λ through

[41] See Berry et al. (2019) for a forceful argument in favor of careful estimation of markups.

systematic variation in implied costs within-hospital across insurers.

Bargaining Weights: Bargaining weights, γ_h, are also separately identified through two key sources of variation. First, since the average cost moments in Section 3.5 pin down costs independently of bargaining weights, the pass-through of these costs to observed prices helps identify bargaining parameters. For example, if a teaching hospital consistently negotiates prices 30% above the level predicted by its cost characteristics, this suggests high bargaining power (low γ_h) for that hospital. (Figure H.2 shows that some hospitals have systematically high negotiated prices across insurers.) Second, the relationship between willingness-to-pay (WTP) differences and price differences helps separate bargaining weights from costs. When a hospital with similar observed costs negotiates higher prices with an insurer, this could in principle reflect either higher (unobserved) marginal costs or greater bargaining power. Higher WTP for that hospital among the insurer's enrollees would suggest the latter. Figure H.1 shows the substantial variation in negotiated prices both within and across hospitals that aids in this identification.

Because negotiated prices are endogenous, there is still concern of potential bias in our parameter estimates, particularly for bargaining weights γ_h. For example, observing a hospital with high average costs but low negotiated prices with a particular insurer could indicate either low hospital bargaining power (high γ_h) or unobserved low costs specific to that insurer-hospital pair. To address this endogeneity, we instrument negotiated prices using the log of insurer m's enrollees within 20 miles from each hospital (see Appendix F). These instruments affect negotiating leverage through market share but should be uncorrelated with idiosyncratic cost shocks, following similar logic to our demand model identification.

Fraction of Balance Bill Recouped: The fraction of balance bills that hospitals can recoup, μ, is identified through variation in potential balance bills (the gap between charge prices p_h^c and out-of-network reimbursement prices p_m^0). If hospitals with higher potential balance bills negotiate higher in-network prices, conditional on costs and bargaining weights, this suggests a higher μ. The network formation inequalities provide additional identifying variation. Hospitals with high potential balance bills being out-of-network more frequently suggests a high μ, as these hospitals require higher negotiated prices to join networks. Cross-insurer variation is particularly informative: insurers with systematically lower p_m^0 create larger potential balance bills, and if these insurers pay systematically higher in-network prices, this helps identify μ.

Insurer Weight on Enrollee WTP: Identification of the weight insurer objective functions place on enrollee surplus, α, relies on both price first-order conditions and network formation inequalities. Higher α implies insurers are willing to pay more to include high-WTP hospitals in their networks. The inequality moments are particularly informative here: if insurers consistently include high-WTP hospitals despite high prices, this suggests a high α. Figure A.2 suggests that insurers are indeed responsive to their enrollees' surplus. Harvard Pilgrim, which has a relatively large number of enrollees in the northern part of New Hampshire, includes many more northern hospitals in its network than does Tufts, whose enrollees are clustered near the southern border. MVP, which has a lumpier distribution of enrollees throughout New Hampshire and additional enrollees to the west in Vermont (not pictured), includes hospitals near its enrollee clusters and some hospitals near the Vermont border. All three insurers include the premier academic medical center in the state, Dartmouth-Hitchcock Medical Center, which patients value highly.

Contracting Costs: Contracting costs, b_m, are identified primarily through network formation inequalities, with three key assumptions. First, costs are identical across all hospitals for a given insurer. Second, these reflect annual fixed costs of negotiation that are incurred irrespective of whether an insurer had a contract with a hospital in prior years.[42] Third, b_m is identified only for insurers with incomplete networks, as they provide both upper and lower bounds; we assume zero costs for insurers with complete networks. While these assumptions are strong, our counterfactual predictions about out-of-network reimbursement regulation primarily operate through costs and bargaining parameters, making them relatively invariant to b_m estimates.[43]

[42] That is, we assume that the negotiating process is costly, even for renegotiations of existing contracts. This is motivated by two facts. First, insurers and hospitals employ dedicated staff for contract negotiations with the other party. Second, existing evidence has shown that the administrative burden of dealing with contract negotiations adds considerable expense and complexity on both the insurer and provider sides (Wikler et al. 2012). Contracting disputes sometimes arise between parties that have a history of successful negotiations. Such disputes can require prolonged and costly negotiating before the parties ultimately agree.

[43] As we argue in Section 3.4, the primary mechanism for this effect is through model estimates of hospital marginal costs and bargaining parameters. Therefore, while precise estimates of contracting costs do help to rationalize the observed networks in the data at baseline, our counterfactual predictions of the effect of regulating out-of-network reimbursements are largely invariant to our estimates of b.

Figure H.1: Distributions of In-Network Prices by Insurer

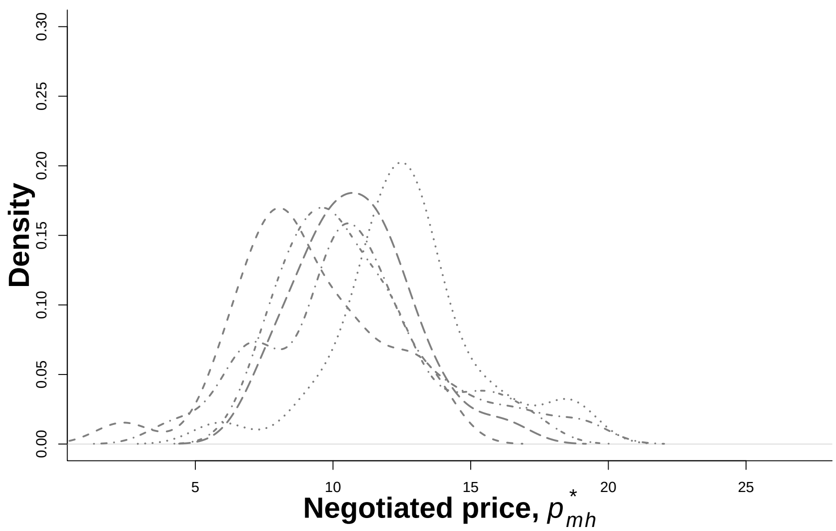

This figure plots the price indices for the five insurers with complete networks across hospitals in New Hampshire in 2012. Each curve represents the distribution of one insurer's negotiated hospital prices. The variation in negotiated prices within an insurer across hospitals contributes to the identification of the hospital marginal cost estimates.

Figure H.2: Hospitals' In-Network Prices Across Large Insurers

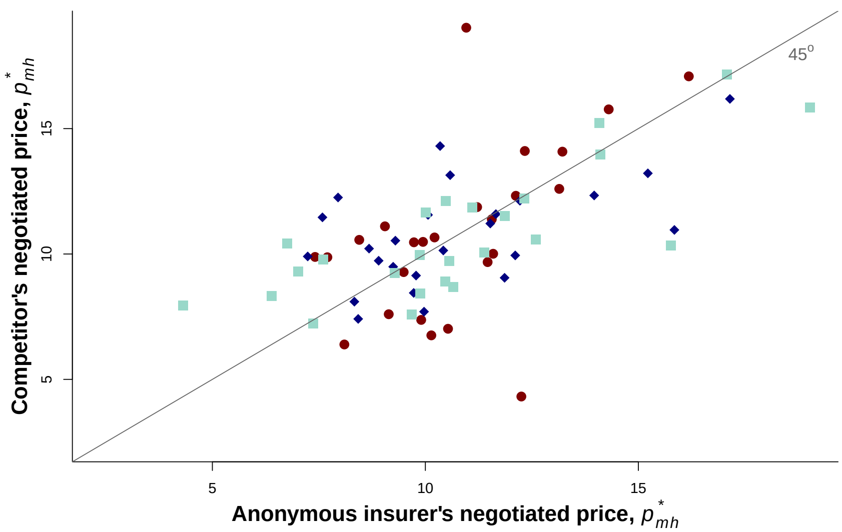

This figure plots the relationship between the three largest insurers' (by market share) negotiated prices. Each dot represents a hospital. Each dot type represents a single (anonymous) pair of insurers. Hospitals that have a high negotiated price with one insurer typically also have high prices with other insurers. This within-hospital correlation contributes to the identification of the hospital marginal cost estimates

Figure H.3: Hospitals' In-Network Prices Across Complete- Vs. Narrow-Network Insurers

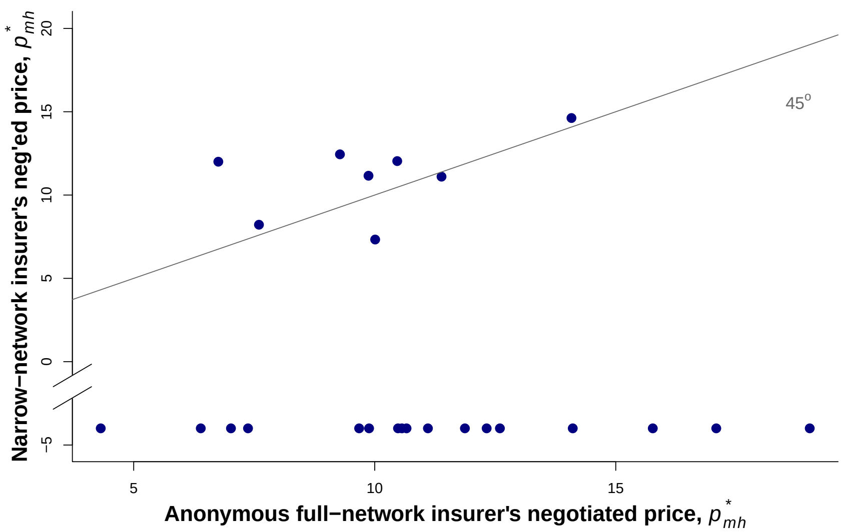

This figure plots the relationship between the 2012 negotiated prices of a high-market share insurer with a complete hospital network (horizontal axis) and the negotiated prices of one of the two narrow-network insurers (vertical axis). Each dot represents a hospital. The insurer on the horizontal axis typically negotiates a lower price than the insurer on the vertical axis with the same hospital. The hospitals excluded from the narrow-network insurer's network have disproportionately high negotiated prices with complete-network insurers. This type of variation across insurers contributes to the identification of the Nash bargaining weight estimates.

I Details of Counterfactuals

I.1 Counterfactual Simulation Algorithm

Algorithm without Hospital Closures: We begin by describing the version of the counterfactual simulation algorithm that allows hospitals to turn away patients whose price is below marginal cost, and therefore does not permit a notion of hospital exit. This algorithm proceeds as follows at each iteration t:

1. Use the bargaining first-order conditions in Equation 3 to simulate in-network negotiated prices p_{mh}^t given the set of estimated parameters $\hat{\boldsymbol{\theta}}$ when we set p_m^0 to the counterfactual reimbursement.

2. Given the new in-network prices in Step 1, use the network inclusion and exclusion conditions (Equation 9 and Equation 10) to check whether any new network links form or whether any existing network links sever. Denote each network link by I_{mh}^t.

3. If a new link forms, assign the predicted in-network price p_{mh}^t from Step 1. If a link severs, assign the counterfactual out-of-network reimbursement p_m^0 to the severed link.

4. If $\max_{m,h} \left|p_{mh}^t - p_{mh}^{t-1}\right| < \epsilon$ and $\max_{m,h} \left|I_{mh}^t - I_{mh}^{t-1}\right| = 0$, stop. Otherwise, return to Step 1 using the updated p_{mh}^t, I_{mh}^t.

The convergence criterion requires that network links do not change between iterations $t-1$ and t, and that prices change by no more than \$0.01 ($\epsilon = 0.01$). Because network links are allowed to change, finding an equilibrium is not guaranteed. We search for equilibria starting from the observed equilibrium in the data, and incrementing the policy parameter sequentially away from the observed parameter. This procedure reduces the risk of switching equilibria across different counterfactuals.

Algorithm with Hospital Closures: The amended counterfactual simulation algorithm that permits closures but has naive negotiators is as follows. At each iteration t:

1. Use the bargaining first-order conditions in Equation 3 to simulate in-network negotiated prices p_{mh}^t given the set of estimated $\hat{\boldsymbol{\theta}}$, when we set p_m^0 to the counterfactual reimbursement.

2. Given the new in-network prices in Step 1, use the network inclusion and exclusion conditions (Equation 9 and Equation 10) to check whether any new network links form or whether any existing network links sever. Denote each network link by I_{mh}^t.

3. If a new link forms, assign the predicted in-network price p_{mh}^t from Step 1. If a link severs, assign the counterfactual out-of-network reimbursement p_m^0 to the severed link.

4. Calculate each hospital's total variable profit across insurers. If total variable profit is negative, assign hospital h to exit the market. Denote each closure by C_{mh}^t.

5. Given the price assignments from Step 3 and the exits from Step 4, check whether any exited hospital can profitably re-enter the market. If so, add it back to the set of hospitals negotiating in the next iteration.

6. If $\max_{m,h} \left|p_{mh}^t - p_{mh}^{t-1}\right| < \epsilon$, $\max_{m,h} \left|I_{mh}^t - I_{mh}^{t-1}\right| = 0$, and $\max_{m,h} \left|C_{mh}^t - C_{mh}^{t-1}\right| = 0$, stop. Otherwise, return to Step 1 using the updated p_{mh}^t from Step 1, I_{mh}^t from Step 2, and updated C_{mh}^t from Step 4.

With sophisticated negotiators, the counterfactual simulation algorithm that permits closures is, at each iteration t:

1. Set p_m^0 to the counterfactual out-of-network price. Given the set of estimated $\hat{\boldsymbol{\theta}}$ and current network status of hospital h with other insurers, calculate h's variable profit if it were to be out of network for insurer m.[44]

2. If the variable profit from Step 1 is weakly positive, then use the bargaining first-order conditions in Equation 3 to solve for in-network negotiated price p_{mh}^t and set hospital h to remain open. Denote each network link by I_{mh}^t and closure status by C_{mh}^t.

3. If the variable profit from Step 1 is negative, then redefine the disagreement payoff to be the hospital's closure. Check whether there exists an in-network price insurer m is willing to pay that would return hospital h to positive variable profits. If yes, solve for that new in-network price p_{mh}^t using first-order conditions defined based on closure as the disagreement payoff and set hospital h to be in insurer m's network.[45] If not, set hospital h to close (for all insurers). Denote each network link by I_{mh}^t and closure status by C_{mh}^t.

4. Given the price assignments from Steps 2 and 3 and the closures from Step 3, check whether any exited hospital can profitably re-enter the market. If so, add it back to the set of hospitals negotiating in the next iteration.

5. If $\max_{m,h} \left|p_{mh}^t - p_{mh}^{t-1}\right| < \epsilon$, $\max_{m,h} \left|I_{mh}^t - I_{mh}^{t-1}\right| = 0$, and $\max_{m,h} \left|C_{mh}^t - C_{mh}^{t-1}\right| = 0$, stop. Otherwise, return to Step 1 using the updated p_{mh}^t from Steps 2 and 3, I_{mh}^t from Steps 2 and 3, and updated C_{mh}^t from Step 3.

[44] The variable profit is given by $\sum_{m' \neq m} \left(p_{m'h} - c_{m'h}\right) \sigma_{m'h} + \left(p_{mh}^0 - c_{mh}\right) \sigma_{mh}^0$, where $p_{m'h}$ and $\sigma_{m'h}$ are its price and volume from other insurers m' if it is out of network with insurer m and has its current network status with all other insurers.

[45] With some abuse of notation, this first-order condition is given by $p_{mh}^t = \frac{1}{\sigma_{mh}^1} \left[(1-\gamma_h) \left[\alpha \left(W_{mh}^1 - W_m^{hc}\right) - \left(\psi_{mh}^1 - \psi_m^{hc}\right) - b_m\right] - \gamma_h \pi_{m'h} + \gamma_h \sigma_{mh}^1 c_{mh}\right]$, where $\pi_{m'h}$ is h's variable profit from all insurers besides m and superscript hc indicates hospital h is closed.

The convergence criterion for both amended algorithms requires that market exit status and network links do not change between iterations $t-1$ and t, and that prices change by no more than \$0.01 ($\epsilon = 0.01$). Because exit, entry, and network links are allowed to change, finding an equilibrium is not guaranteed. We deal with multiple equilibria by attempting to trace out the equilibrium path that corresponds to the observed equilibrium played in the data. We do so by adjusting out-of-network prices in small increments, and using the previous computed equilibrium as the starting values for the next increment of out-of-network prices.

I.2 Alternate Multiples of Charge Price Benchmarks

In this section, we consider rescaling the out-of-network prices to alternate multiples of the current benchmarks. This is meant to approximate the impact on in-network hospital prices of proposals to set out-of-network reimbursements closer to hospitals' current charge prices.

Figure I.1 plots the results of this simulation. Regulating out-of-network prices to levels below their current value yields similar effects on prices and networks as pegging reimbursements to low multiples of Medicare prices. Conversely, in-network negotiated prices rise with increases in the off-contract prices, as this improves the hospital's disagreement value and worsens the insurer's disagreement value. Hospitals therefore gain considerable bargaining leverage to raise prices.[46] At current off-contract prices (multiple of 1.0 on the horizontal axis), the average predicted in-network price is \$11.09 (for a routine venipuncture). However, if off-contract prices were to increase to twice the current benchmark, then average negotiated prices are predicted to increase by approximately 7 percent to an average of \$11.86. On the other hand, reducing the benchmark to half of the current benchmark would drive average predicted in-network prices down by approximately 47 percent to \$5.92.

Two forces moderate the magnitude of the effect of regulating out-of-network prices. First, we continue to allow hospitals to turn away unprofitable out-of-network patients. Second, hospitals are paid at most their charge price, even when the benchmark exceeds it. For most insurers, the price effects flatten out when the regulated price reached about 1.2 times its current value. Interestingly, in the case of Aetna and Cigna, this results in slightly narrower networks. The logic is intuitive: as hospitals receive the maximum allowable price for out-of-network care, some may prefer to strategically opt out of insurance networks if they believe they can retain a sufficient volume.

[46] Note that in the vicinities of equilibrium network transitions, an equilibrium cannot always be found.

Figure I.1: Predicted Negotiated Prices Vs. Multiples of Current Off-Contract Prices

(a) Aetna

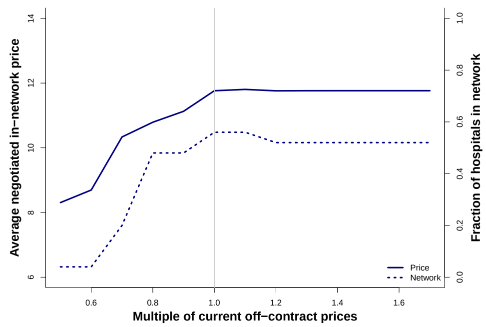

(b) Anthem

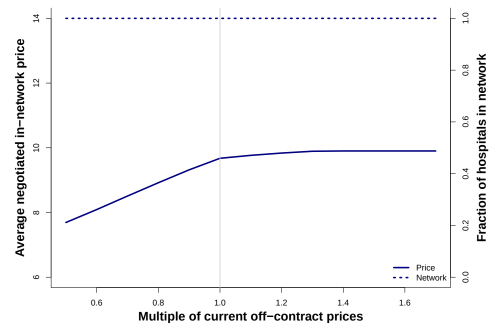

(c) Cigna

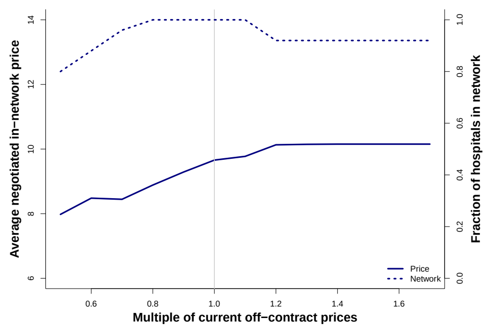

(d) Harvard Pilgrim

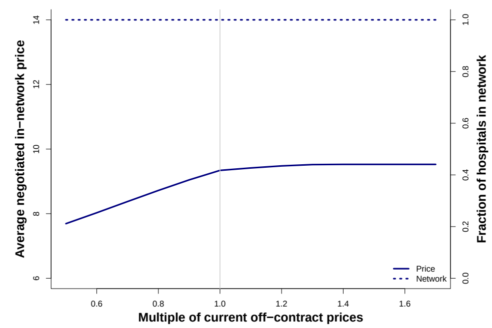

(e) MVP

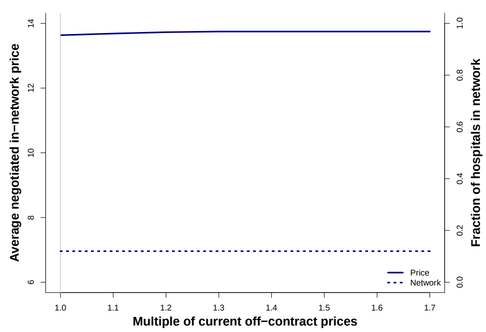

(f) Tufts

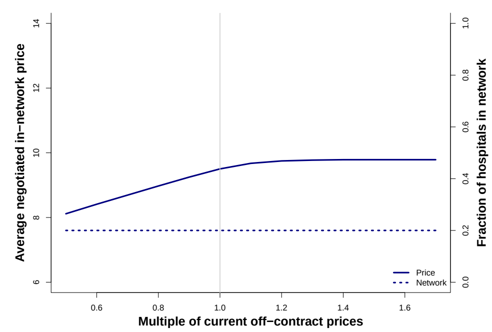

(g) United

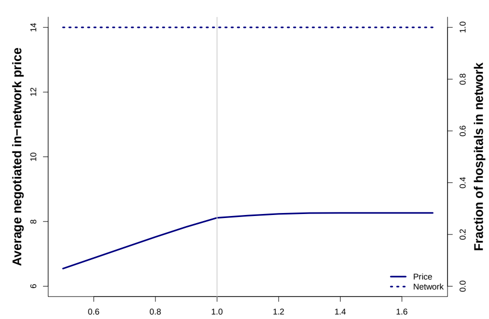

Notes: This figure plots results of counterfactual simulations varying out-of-network prices to be various multiples of current out-of-network payments. The left vertical axis plots the counterfactual negotiated price for hospitals predicted to be in-network (volume-weighted average prices). The right vertical axis plots the fraction of hospitals predicted to be in-network. Gray vertical line is value of observed out-of-network payments. Gaps represent counterfactuals for which no equilibrium was found.

I.3 Dynamic Payment Caps

Next, we consider the Alexander (2019) proposal to peg out-of-network prices to median negotiated in-network prices in the corresponding geographic market.[47] In the majority of markets, median negotiated prices are substantially lower than current out-of-network prices. Forcing hospitals to accept the median in-network price when they are out of network would therefore amount to worsening their bargaining leverage vis-à-vis insurers.[48] This is precisely the mechanism that the policy proposal is designed to leverage in order to reduce negotiated in-network prices.

In principle, this policy is dynamic and self-reinforcing. The policy would peg the current year's out-of-network prices to the preceding year's median negotiated prices. In the first year of implementation, this would reduce in-network prices due to the worsening of hospitals' bargaining leverage. The following year, the median will be calculated from a negotiated price distribution that has shifted to the left, further reducing negotiated in-network prices.

Figure I.2 simulates these effects over the first ten years following policy implementation. Both negotiated and out-of-network prices drop substantially, but so too does network breadth for several insurers. Nearly the entire reduction in in-network prices obtains within the first two years, despite out-of-network prices continuing to fall through the fourth year.

[47] In this set of counterfactuals, we replace the geographic delineations used by insurers in the data for calculating out-of-network prices with the geographic markets defined in the proposed legislation.

[48] We ignore the possible countervailing effect of hospitals internalizing the effect of today's negotiated prices on tomorrow's disagreement payoffs, if hospitals are dynamically strategic.

Figure I.2: Predicted Negotiated Prices with Dynamic Out-of-Network Payment Caps

(a) Aetna

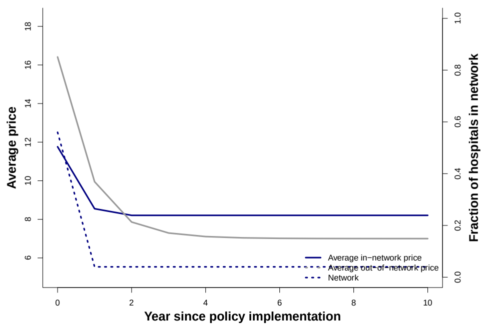

(b) Anthem

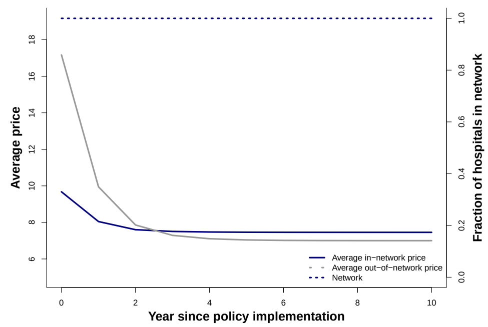

(c) Cigna

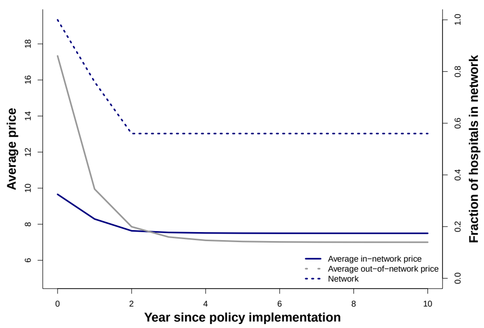

(d) Harvard Pilgrim

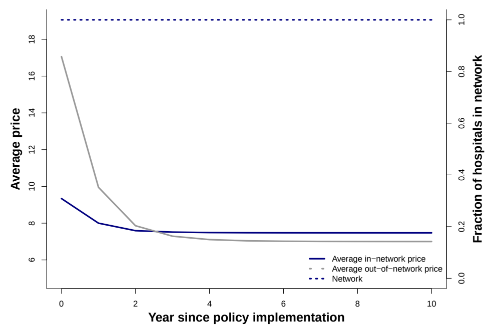

(e) MVP

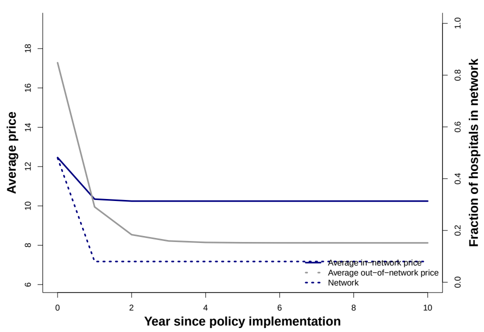

(f) Tufts

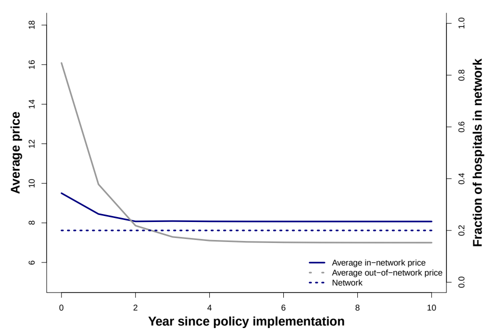

(g) United

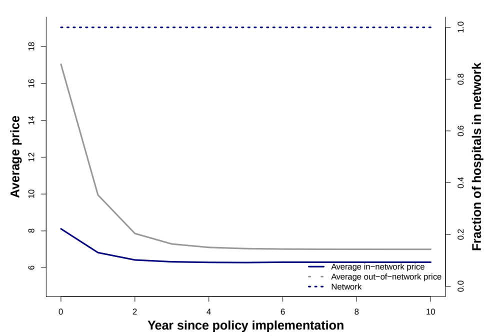

Notes: This figure plots results of counterfactual simulations pegging out-of-network prices to the median of the prior year's negotiated in-network prices. The left vertical axis plots the counterfactual negotiated price for hospitals predicted to be in-network (volume-weighted average prices). The right vertical axis plots the fraction of hospitals predicted to be in-network. Gaps represent counterfactuals for which no equilibrium was found.