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

Elena Prager
Kellogg School of Management
Northwestern University

Nicholas Tilipman
School of Public Health
University of Illinois at Chicago

This version October 1, 2020. Latest draft here.

Abstract

Recent policy proposals seek to regulate the prices that hospitals can levy for care delivered outside of a patient's insurance network. In this paper, we study the potential effects of such regulations on equilibrium in-network prices and access. We first describe how out-of-network reimbursements affect negotiations over in-network prices and network status. We then conduct a series of counterfactuals to empirically evaluate current policy proposals that would cap out-of-network reimbursements. To do so, we estimate a model of insurer-hospital bargaining that explicitly allows for transactions in the absence of a contract. We operationalize the model using a novel, data-driven measure of out-of-network prices paid by insurers to hospitals. Our results suggest that reducing out-of-network reimbursements would have the intended effect of lowering negotiated prices with in-network hospitals. Aggressive regulation, however, would reduce access to care as a result of narrower networks and some outright service line closures.

JEL codes: C78, I11, L13

*Previously circulated under the title "Disagreement Payoffs and Negotiated Prices: Evidence from Out-of-Network Hospital Payments". We thank Chris Garmon, Kate Ho, Tom Hubbard, Robin Lee, and seminar participants at RAND, DOJ, the Stanford SITE workshop, and the APPAM annual conference for helpful suggestions. Any errors are our own. Contact information: Prager: elena.prager@kellogg.northwestern.edu; Tilipman: tilipman@uic.edu

1 Introduction

Policy-makers have recently sought to regulate the prices that hospitals can levy when they are not in an insurer's provider network. Proponents tout these out-of-network price caps as a mechanism for not only reducing patients' exposure to surprise medical bills outside their insurer's network, but also reducing negotiated prices with in-network providers (Kane 2019; Chernew et al. 2019). If successful, these reductions may also have the unintended—and under-discussed—effect of pushing hospitals to drop out of insurers' networks. In extreme cases, the price reductions may push hospitals to close altogether. This paper evaluates empirically the effects of out-of-network price caps on equilibrium negotiated in-network prices, network breadth, and provider closures.

The paper's first contribution is to propose a practical solution to the empirical challenge of measuring the out-of-network prices actually paid to hospitals by health insurers. The literature recognizes the importance of correctly accounting for disagreement values in the estimation of bargaining models (Ho and Lee 2017a). Nevertheless, existing papers on hospital-insurer bargaining have assumed away the presence of transactions in the case of disagreement, owing in part to the difficulty of measuring off-contract prices. Beyond the fact that health care markets lack posted prices, off-contract prices 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 health insurers' out-of-network payment policies to construct a measure of off-contract prices. Many insurers base their out-of-network reimbursement policies on third-party benchmarks constructed from hospital charge prices in a given geographic market. We replicate the third-party methodology for constructing these benchmarks using the type of data that are readily available to researchers. The resulting measure yields a reasonable approximation of observed out-of-network hospital payments in our data.

Our second contribution is to use our measure of off-contract prices to extend the canonical Nash-in-Nash bargaining framework. The bulk of the existing work defines the disagreement outcome of a negotiation as severing that pair's link outright (Crawford and Yurukoglu 2012; Ho and Lee 2017b; Gowrisankaran et al. 2015; Prager 2016).[1] This setup implies an assumption that no transactions occur between the two non-contracting parties, and the loss in surplus from disagree-

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

ment is equal to the loss of profit associated with the transactions that occur under agreement.[2] In health care, however, the lack of a formal contract between and insurer and a provider does not completely eliminate transactions between them. Instead, that insurer's patients can—and, in our data, often do—still obtain out-of-network care from such providers.

Our model departs from the existing empirical literature by allowing patients to obtain care at out-of-network providers, and allowing insurers to pay those providers strictly positive off-contract prices. We operationalize the model empirically in the context of the hospital market in New Hampshire, a suitable setting for several reasons. First, some health insurers serve the New Hampshire market even though most of their enrollees reside in the neighboring state of Massachusetts. This generates substantial variation in the contract status of New Hampshire hospitals across insurers, partly driven by variation in the distribution of enrollees across New Hampshire and the New Hampshire–Massachusetts border. Second, we document nontrivial volumes at out-of-network New Hampshire providers. In our sample, out-of-network hospitals account for 14.2 percent of transactions for a large New England insurer. Within this insurer, a typical out-of-network hospital has approximately one tenth of the volume of a typical contracted in-network hospital. This volume is quite high relative to the small size of the out-of-network hospitals, which only make up 20.4 percent of all insurers' total hospital volume in the market. Finally, and critically, out-of-network care in this market is nearly always paid for in part or in whole by insurers.

We show, both theoretically and using our measure of off-contract prices, that estimates of marginal costs are often biased upward when hospitals are assumed to have no out-of-network volumes. The upward bias is larger when the true off-contract volume and price are larger. Empirically, we find that ignoring out-of-network payoffs from disagreement results in overstating hospital marginal costs by twenty percent on average. This overestimation of marginal costs in the standard model ultimately results in overly pessimistic evaluations of two policy goals: access to health care providers and prices. Ordinarily, these two policy goals require a trade-off, since a simple method for reducing prices is to exclude high-priced providers from the network. However, compared to estimates from the canonical model with zero disagreement values, estimates from our model predict both broader networks and lower equilibrium prices at low levels of out-of-network prices.

[2] 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 2017a; Ghili 2017), define the surplus from agreement more flexibly. However, those papers maintain the assumption of zero off-contract transactions.

We next consider counterfactual simulations that mimic proposed regulations. Current proposals to cap out-of-network prices come from both major political parties. The Republican-sponsored Lower Health Care Costs Act of 2019 proposes capping insurers' off-contract payments at median in-network rates within each market (Alexander 2019). A high-profile candidate for the 2020 Democratic presidential nomination proposed setting the cap at 200 percent of Medicare (Pete For America 2019). Other third-party proposals have called for rates as low as 120 percent of Medicare (Kane 2019). Notably, out-of-network payments are also a subject of antitrust cases against hospitals. For example, California's high-profile complaint against Sutter Health describes Sutter's out-of-network prices as "punitively high" (Becerra et al. 2018; Ellison 2018).[3]

Our first 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 benchmarks and policies that expand the benchmarks to the point where hospitals are nearly paid their full charge price. The second set of counterfactuals considers capping out-of-network reimbursements at multiples of Medicare rates. We find that in all these counterfactual simulations, our model predicts increasing the off-contract prices gives hospitals bargaining leverage to negotiate above-cost prices. Specifically, doubling the current off-contract price benchmark percent results in a nearly 70 percent increase in average volume-weighted in-network prices. Conversely, reducing off-contract prices to the vicinity of Medicare reimbursements substantially reduces negotiated prices. Pegging off-contract prices to 100 percent of Medicare rates is projected to reduce negotiated prices by nearly one third.

However, while capping out-of-network reimbursements reduces equilibrium prices, it also imposes a trade-off against reduced access to providers. Cutting off-contract prices by half reduces the share of hospitals covered by more than 40 percent. These predictions depart from predictions using the canonical Nash-in-Nash model. Under our counterfactual simulations, the price predictions from the canonical framework are 5 to 10 percent higher than our model with nonzero disagreement values. Moreover, while both sets of predictions produce narrower networks at lower out-of-network price caps, neither set of predictions uniformly dominates the other in terms of network breadth. Finally, our counterfactual simulations suggest that reducing off-contract prices to near the level

[3]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. https://www.nytimes.com/2013/05/17/business/bayonne-medical-center-has-highest-us-billing-rates.html

of Medicare rates would drive substantial hospital exit as in-network prices begin to drop below hospital marginal costs.

Our paper relates to several strands of literature. Several recent papers have proposed approaches to relaxing the Nash assumption that in case of disagreement, all other parties' contracts remain the same (Ho and Lee 2017a; Ghili 2017; Liebman 2017). We view our approach as complementing these important advances by providing a computationally simple alternative for dealing with 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 rates, especially in the context of surprise out-of-network bills (Cooper et al. 2019a; Craig et al. 2019; Cooper et al. 2019c,b). 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 service line closures.

Our main conceptual point carries over to other industries. In television markets, for example, content providers receive revenue directly from advertisers as well as from cable companies. The loss of a contract with a cable company therefore reduces surplus not just by the fees directly associated with that contract, but also by the reduced fees advertisers will be willing to pay as a result of losing access to that cable company's subscribers. In a similar vein, a two-sided platform that loses a brand from among its sellers will likely see an increase in purchases of that brand's products from third-party sellers.[4] Therefore, the importance of defining surplus from agreement more flexibly than the direct value of a contract extends to a variety of industries.

The paper proceeds as follows. Section 2 discusses the details of our algorithm to measure off-contract prices. Section 3 presents our theoretical model and empirical strategy. That section also includes discussion of the direction and magnitude of bias arising from assuming zero out-of-network volumes. Section 4 describes our empirical context, data, and sample. Section 5 presents the parameter estimates, and Section 6 presents counterfactual simulations. Finally, Section 7 concludes.

[4] A high-profile example of this is Nike's withdrawal from its contract with Amazon in fall 2019, following Nike's dissatisfaction with Amazon's handling of counterfeit and third-party merchandise (Hanbury 2019).

2 Measuring Off-Contract Prices

Health insurers do not contract with every health care provider in the United States. Because the U.S. health care system lacks posted prices (Reinhardt 2006), insurers typically put in place explicit policies governing how much they will pay non-contracted providers. While insurers could in principle refuse to pay non-contracted providers at all, in practice they face demand-side incentives 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. Insurers often pay some portion of the bill for out-of-network care, and these payments can be substantial.[5]

Most insurers have policies that rely on "usual and customary" rates to determine payment for out-of-network services. 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, although it is occasionally pegged to fee-for-service Medicare payment rates. Table 1 quotes the relevant language from several insurers' policy documents.

Insurers are not always explicit about how they define the prevailing market rate, 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 privately insured and fee-for-service Medicare claims. Among its flagship products are the FH Charge Benchmarks, which many insurers use as an input to determining out-of-network payment rates. This product reports quantiles summarizing the distribution of charge prices at the level of a geographic area-treatment type pair. It is updated twice a year using a rolling twelve-month window of claims data. Insurers that purchase the Charge Benchmarks can then use a given percentile of the charge price distribution as an input to their determination of out-of-network rates, as indicated by the quotes from Aetna's, Cigna's, and United's policies in Table 1. Although FAIR Health products are not intended to be used as suggested appropriate payment amounts, insurers' payment policies are often informed by them.

We infer insurers' policies with respect to the charge benchmarks by comparing payments for

[5]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 carriers.

Table 1: 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
Blue Cross Blue Shield 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

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 data by closely following FAIR Health's algorithm. The algorithm is public and is described 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 1 shows the distribution of the ratio of paid amounts to the 60th percentile benchmarks for a large PPO plan run by one of our key insurers. This plan typically pays for out-of-network care at 100 percent of the 60th percentile benchmark, as indicated by the spike in the distribution at 1. 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. Whereas FAIR Health uses the near-universe of privately insured claims and the universe of fee-for-service

Figure 1: Out-of-Network Payments in a Sample Plan

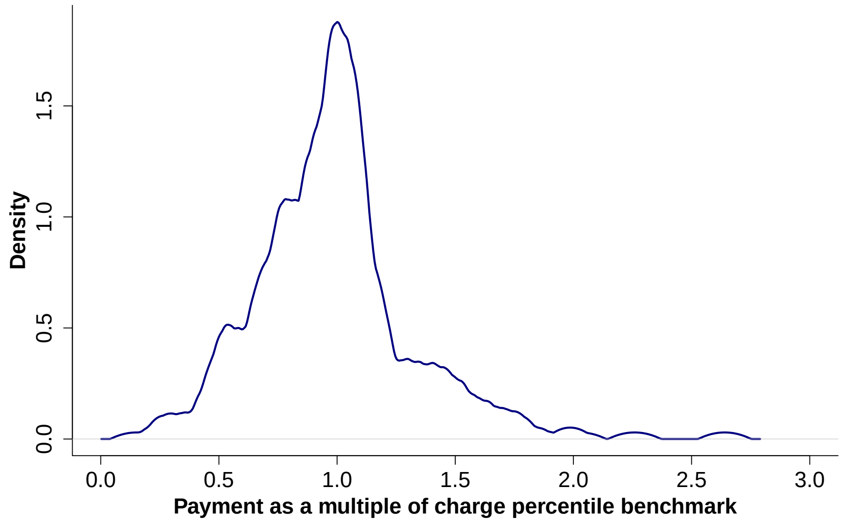

Tufts Health Plan's payment amounts for out-of-network outpatient hospital transactions in a flagship PPO plan, as a multiple of the 60th percentile charge benchmark for the corresponding procedure code. This plan typically pays out-of-network hospitals at 100 percent of the 60th percentile benchmark.

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.[6]

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 our primary sample, this requires examining its claims from other markets in which it has a narrow network. These out-of-network policy inferences are facilitated by comprehensive data on insurers' networks, described in Section 4.3. We then use the inferred policies to construct off-contract prices for pairs of insurers and hospitals that do not necessarily have a contract. These off-contract price measures are a key input to estimating our Nash bargaining model with nonzero disagreement volumes, to which we now turn.

3 Model and Estimation

The goal of the model is to make inferences from the equilibrium network status and equilibrium in-network prices observed in the data. Hospitals do not have posted prices that are systematically

[6] We are in the process of negotiating a purchase of the proprietary FAIR Health data.

paid by purchasers of their services. Instead, health insurers negotiate with hospitals to arrive at a contracted price that the hospitals will be paid for providing services to the insurers' enrollees. We model these negotiations as pairwise Nash bargaining interactions, but depart from the hospital bargaining literature by specifying strictly positive off-contract prices and volumes.

The model proceeds in three stages:

1. Insurer m and hospital h decide whether to enter into negotiations. If so, a contracting cost b is incurred.

2. If they have decided to enter into negotiations, insurer m and hospital h engage in bilateral negotiations that, if successful, determine the in-network price p_{mh}.

3. With some probability f_{id}, patient i enrolled in insurer m's plan gets sick and requires procedure d. The patient chooses a hospital from among the hospitals in the market, 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 3. Second, we estimate the insurer-hospital bargaining model corresponding to Stages 1 and 2 using objects constructed from the hospital choice model estimates. There are four sets of parameters to estimate from the bargaining model: hospitals' marginal costs of treating patients; insurers' weighting of enrollee expected utility relative to hospital expenditures; Nash bargaining weights; and contracting costs. We estimate these parameters using the generalized method of moments.

In the sections that follow, we discuss the model and estimation pertaining to Stages 1 and 2. The discussion of Stage 3 is relegated to Appendix C, as we follow well-established discrete choice models of hospital demand from the literature. Section 3.1 derives the equilibrium prices in case of agreement, conditional on entering into negotiations. Section 3.2 then discusses how an insurer and a hospital decide whether to enter into negotiations in the first place. Section 3.3 describes how we use the equilibrium price conditions from Section 3.1 and the network status conditions from Section 3.2 for estimation. Section 3.4 then shows how the predictions of the bargaining model with positive disagreement volumes depart from the predictions of a model with zero volume in case of disagreement. Finally, Section 3.5 discusses identification.

3.1 Model Setup: Price Negotiation Stage

In Stage 2 of the model, insurer m and hospital h have decided to enter into negotiations. A negotiated contract between them 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.[7] 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, and 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_m^0. The insurer's out-of-network payment rates depend only on the services provided, not the identity or cost structure of the hospital.[8]

Hospital Objectives

We model hospitals as profit maximizers. Conditional on entering into the negotating process, hospital h's surplus from a contract with insurer m at a negotiated price p_{mh} is given by

$$S_h(m, p_{mh}) = (p_{mh} - c_h)\,\sigma^1_{mh} - \left(p^0_m - c_h\right)\sigma^0_{mh} \tag{1}$$

where c_h is the hospital's marginal cost of treating a typical patient, and $\sigma^1_{mh} > \sigma^0_{mh}$ are the hospital's patient volumes from insurer m in the case of agreement and disagreement, respectively. A hospital's volume under a given network configuration is predicted from the hospital demand model discussed in Appendix C. In the empirical application, we weight patient volumes by a measure of resource intensity associated with the services provided, and assume that the price and the hospital's cost both scale linearly by the resource intensity.

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

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

[8] In practice, each insurer's out-of-network payment rates 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.

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 between models assuming zero disagreement volumes and models accounting for positive disagreement volumes that we outline in Section 3.4 obtain for both sets of insurer objectives.

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

$$S_m(h, p_{mh}) = \left(\alpha_m W^1_{mh} - p_{mh}\sigma^1_{mh} - \psi^1_{mh}\right) - \left(\alpha_m W^0_{mh} - p^0_m \sigma^0_{mh} - \psi^0_{mh}\right) \tag{2}$$

where α_m is the insurer's weight on enrollee expected utility, and $W^1_{mh} > W^0_{mh}$ are the expected utilities in the case of agreement and disagreement, respectively. 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:

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

where $\gamma_m \in [0, 1]$ is insurer m's Nash bargaining parameter. Taking the derivative of the logged Nash product with respect to price, the first-order condition describing p^*_{mh} becomes

$$\gamma_m \frac{-\sigma^1_{mh}}{\alpha_m W^1_{mh} - p^*_{mh}\sigma^1_{mh} - \psi^1_{mh} - [\alpha_m W^0_{mh} - p^0_m \sigma^0_{mh} - \psi^0_{mh}]}$$
$$=$$
$$-(1-\gamma_m)\frac{\sigma^1_{mh}}{(p^*_{mh} - c_h)\,\sigma^1_{mh} - (p^0_m - c_h)\,\sigma^0_{mh}}$$

which yields an equilibrium price of

$$p_{mh}^{*} = \frac{1}{\sigma_{mh}^{1}} \left[\begin{array}{l} (1-\gamma_m)\,\alpha_m \left(W_{mh}^{1} - W_{mh}^{0}\right) + p_m^{0}\sigma_{mh}^{0} \\ +\gamma_m c_h \left(\sigma_{mh}^{1} - \sigma_{mh}^{0}\right) - (1-\gamma_m)\left(\psi_{mh}^{1} - \psi_{mh}^{0}\right) \end{array} \right] \tag{3}$$

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

3.2 Model Setup: Network Formation Stage

The price negotiations discussed in Section 3.1 take place only if an insurer and a hospital decide in Stage 1 of the model to enter into negotiations. 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. The parties will then anticipate that no agreement will be reached in Stage 2, and therefore will decide in Stage 1 not to enter into negotiations.

We model negotiations as costly: the insurer and hospital must jointly pay a contracting cost b for each pairwise negotiation. This modeling assumption is 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.[9] We interpret the b parameter as a flavor of Coasian transaction cost. Once the contracting cost is paid, it is sunk. Therefore, the contracting cost does not enter into the price negotiations in Stage 2.

The condition for entering into negotiations 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, less the contracting cost:

$$E_{mh} = \alpha_m W_{mh}^{1} - \psi_{mh}^{1} - \left(\alpha_m W_{mh}^{0} - p_m^{0}\sigma_{mh}^{0} - \psi_{mh}^{0}\right) - c_h\sigma_{mh}^{1} - \left(p_m^{0} - c_h\right)\sigma_{mh}^{0} - b \tag{4}$$

[9] 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. We therefore estimate the network inclusion moments as maximizing a *joint* surplus and interpret the contracting costs as a combination of insurer and hospital negotiating costs.

The Nash-in-Nash structure of Stage 2 guarantees that, if $E_{mh} \geq 0$ and the pair enters negotiations, then an agreement will be reached in Stage 2. We leverage this in the estimation by inferring a weakly positive joint surplus in Stage 1 if and only if an agreement is observed. Note that 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, $p_{mh}\sigma^1_{mh}$.

3.3 Estimation of Model Parameters

We use the model conditions from the price negotiation stage and the network formation stage to form moments for estimation. There are four sets of parameters to estimate from the model: the insurers' weights on expected utility α_m, the insurers' Nash bargaining weights γ_m, the hospital marginal costs c_h, and the joint contracting cost b. All other objects in the model are predicted from the hospital demand model (see Appendix C) and treated as data. This section describes the two sets of moments that enter into our generalized method of moments estimation. The price negotiation stage contributes equality moments, and the network formation stage contributes inequality moments. We defer the discussion of identification to Section 3.5.

Equality Moments from Price Negotiation

Hospital-insurer pairs that have a negotiated contract contribute equality moments from the first-order conditions on negotiated price. In our equilibrium condition for Stage 2, equation 3, prices are observed, whereas hospital marginal costs, c_h, 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_h

$$c_h = \lambda g_h + \nu_h \tag{5}$$

where λ is a parameter vector and ν_h is the unobservable component of hospital costs. The observable characteristics in g_h on which we project costs include hospital fixed effects, which subsume hospital characteristics that remain fixed over the course of our sample period, such as teaching status and system status; and year fixed effects, which allow for flexible statewide trends in cost growth.

The econometric error for the GMM estimator is then defined as the difference between the projected cost from Equation 5 and the cost implied by the first-order conditions on equilibrium prices from Equation 3. That is, we define the econometric error for a hospital-insurer pair as

$$\xi_{mh} = \lambda g_h - \frac{1}{\gamma_m \left(\sigma^1_{mh} - \sigma^0_{mh}\right)} \left[\begin{array}{l} p^*_{mh}\sigma^1_{mh} - (1-\gamma_m)\,\alpha_m \left(W^1_{mh} - W^0_{mh}\right) \\ -p^0_m \sigma^0_{mh} + (1-\gamma_m)\left(\psi^1_{mh} - \psi^0_{mh}\right) \end{array} \right] \tag{6}$$

We then search for parameters λ to set the vector of ξ_{mh} across pairs orthogonal to a set of assumed exogenous variables z_{mh}. Following Gowrisankaran et al. (2015), we include in z_{mh}: a hospital's predicted contribution to enrollees' expected utility, $W^1_{mh} - W^0_{mh}$; its expected per-enrollee contribution to expected utility; and predicted hospital quantity. The equality moment that enters into the GMM estimation is then

$$\mathbb{E}\left[\xi_{mh} | z_{mh}\right] = 0 \tag{7}$$

This gives us one moment per hospital-insurer pair in each year that the pair has a negotiated contract. Out-of-network hospitals do not contribute to this set of moments, as the Nash bargaining first-order condition on which the moments are based is not defined in the absence of a negotiated price contract.

Inequality Moments from Network Formation

In addition to the equality moments contributed by hospital-insurer-years in which we observe an agreement (Equation 7), 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 disagreement volumes and out-of-network reimbursement benchmarks. However, varying the level of out-of-network payments may result in carriers or hospitals deciding it is more profitable to enter into a formal contract (and negotiate an in-network rate) rather than remain out-of-network under counterfactual policies. As such, our model needs to incorporate carrier and hospital decisions surrounding network formation with currently out-of-network hospitals, as several recent papers have done (Ho and Lee 2017a; Ghili 2017; Liebman 2017). 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.

We therefore incorporate into the estimation additional moments from the network status determination decisions discussed in Section 3.2. To construct these moments, we follow closely 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:

$$\begin{aligned} E_{mh}(\boldsymbol{\theta}) &= \alpha_m W^1_{mh} - \psi^1_{mh} - \left(\alpha_m W^0_{mh} - p^0_m \sigma^0_{mh} - \psi^0_{mh}\right) - c_h \sigma^1_{mh} - \left(p^0_m - c_h\right)\sigma^0_{mh} - b \\ &= \alpha_m \left(W^1_{mh} - W^0_{mh}\right) - \psi^1_{mh} + \psi^0_{mh} + \left(-\sigma^1_{mh} + \sigma^0_{mh}\right) c_h - b \end{aligned}$$

where c_h is projected from Equation 5, and $\sigma^1_{mh}, \sigma^0_{mh}, W^1_{mh}, W^0_{mh}, \psi^1_{mh}, \psi^0_{mh}$ are predicted from the demand model. 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 for any contract and that they predict these gains with error.[10] 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.[11] Then:

$$\begin{aligned} \mathbb{E}[E_{mh}(\boldsymbol{\theta})|\mathcal{J}] &= E_{mh}(\boldsymbol{\theta}) - \omega_{mh} \\ &= \left[\alpha_m \left(W^1_{mh} - W^0_{mh}\right) - \psi^1_{mh} + \psi^0_{mh} + \left(-\sigma^1_{mh} + \sigma^0_{mh}\right) c_h - b\right] - \omega_{mh} \end{aligned}$$

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:

$$0 \leq \mathbb{E}[E_{mh}(\boldsymbol{\theta})|\mathcal{J}] = E_{mh}(\boldsymbol{\theta}) - \omega_{mh} \tag{8}$$

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

[10]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 rates and estimates of gains from trade.

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

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.[12] The resulting inequality for estimation is given by:

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

Each hospital-insurer pair that is observed not to have a negotiated contract therefore contributes a single inequality, defined by Equation 9, that imposes upper bounds on the surpluses from agreement. We refer to these inequalities as network exclusion moments. In the estimation, if insurer m and hospital h are observed not to have a network but Equation 9 is violated—that is, if the implied total surplus at the current parameter values is positive—we penalize the objective function by the magnitude of the violation.

Collectively, the network inclusion and exclusion conditions are what Ghili (2017) calls network 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:

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

We search for a full set of parameters, $\boldsymbol{\theta}$,that satisfies this full system of inequalities. If no set of parameters satisfies all of inequalities, we construct a moment equation that minimizes the absolute deviations for any inequality violated. 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.

[12] 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.

3.4 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 volume σ^0_{mh} and the disagreement payments $p^0_m \sigma^0_{mh}$ biases these structural parameters. Here, we illustrate the bias arising from assuming that disagreement volume is zero when estimating hospital marginal costs c_h.

Consider a simplified empirical setup where the parameters γ_m and α_m are known, leaving only the hospital costs c_h 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 by rearranging Equation 3:

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

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

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

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:

$$\alpha_m \left(W^1_{mh} - W^0_{mh} \right) - \left(\psi^1_{mh} - \psi^0_{mh} \right) + \frac{\sigma^0_{mh} \left(p^0_m - p^*_{mh} \right)}{1 - \gamma_m}$$

$$>$$

$$\frac{\left(\sigma^1_{mh} - \sigma^0_{mh} \right)}{\tilde{\sigma}^1_{mh}} \left[\alpha_m \left(\tilde{W}^1_{mh} - \tilde{W}^0_{mh} \right) - \left(\tilde{\psi}^1_{mh} - \tilde{\psi}^0_{mh} \right) \right] \tag{10}$$

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.[13] To see this, note that the right-hand 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. 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

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

$$\frac{\left(\sigma^1_{mh} - \sigma^0_{mh} \right)}{\tilde{\sigma}^1_{mh}} < \left[\alpha_m \left(W^1_{mh} - W^0_{mh} \right) - \left(\psi^1_{mh} - \psi^0_{mh} \right) + \frac{\sigma^0_{mh} \left(p^0_m - p^*_{mh} \right)}{1 - \gamma_m} \right] \Big/ \left[\alpha_m \left(\tilde{W}^1_{mh} - \tilde{W}^0_{mh} \right) - \left(\tilde{\psi}^1_{mh} - \tilde{\psi}^0_{mh} \right) \right]$$

negotiated in-network price p^*_{mh}, scaled by the inverse of the hospital's Nash bargaining weight.[14] 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 $\alpha_m \left(W^1_{mh} - W^0_{mh}\right) > 0$ is the contribution to surplus of enrollees' willingness-to-pay to include hospital h in the network. This term is typically smaller than its right-hand side analog $\alpha_m \left(\tilde{W}^1_{mh} - \tilde{W}^0_{mh}\right)$, because the WTP gain from an included hospital is smaller when consumers can 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.[15] 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.

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 5 shows how this expectation plays out in the data: the bias is of

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

[15] 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.

the same sign as $p^0_m - p^*_{mh}$ for most hospitals, and is increasing in $p^0_m - p^*_{mh}$.

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.[16] In Section 6, we show how the biased cost estimates affect the predicted effects of counterfacual policies.

3.5 Identification of Bargaining Parameters and Contracting Costs

Identification of hospital marginal costs, c_h, and bargaining weights, γ_m, is similar to Gowrisankaran et al. (2015). The equality moments from Stage 2 of the model (Equation 6) help pin down these parameters. Estimation of these moments relies on exogenous instruments, z_{mh}. We use all the fixed effects included in the cost equation (Equation 5) as well as the the instruments described above. Hospital marginal costs c_h are identified primarily through variation in observed prices within insurer across hospitals. Intuitively, for given guesses of γ_m, α_m, and b, hospitals that have higher observed negotiated prices, p_{mh}, will be predicted to have higher marginal costs. Figure 2 displays the variation in negotiated prices observed in the data that helps to identify c_h. Hospitals that have high negotiated prices with one insurer tend to have high negotiated prices with other insurers. Conversely, γ_m are identified primarily from variation in observed prices within hospitals, across insurers. Suppose, for instance, that two insurers have enrollee distributions with similar WTP for a particular hospital, but those insurers negotiated very different prices with that hospital. This variation would map into different values of γ_m for each insurer, reflecting their differential ability to extract surplus from negotiations. Figure 3 shows that for most hospitals, Harvard Pilgrim negotiates lower prices than Tufts. All else equal, these price differences will map into a higher estimated γ_m for Harvard Pilgrim.

Identification of the insurer's weight on enrollee surplus, α_m, and the contracting cost, b, relies

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

Figure 2: Insurer Distributions of In-Network Prices

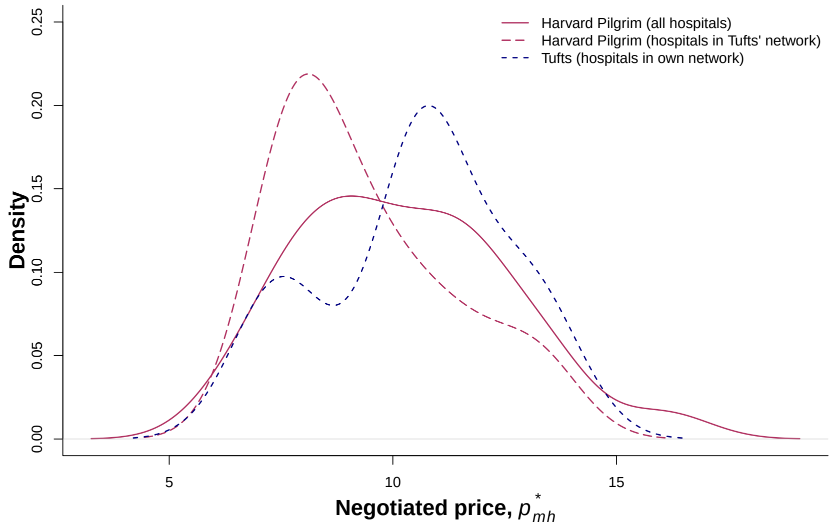

This figure plots the price indices for Harvard Pilgrim and Tufts Health Plan across in-network hospitals in New Hampshire in 2010. The dashed curves are subsetted to hospitals that are in-network for both insurers (because Harvard Pilgrim has a complete network, this is equivalent to subsetting to hospitals that are in Tufts' network).

largely on the inequality moments from Stage 1 of the model in Equations 8 and 9. Since we estimate both γ_m and α_m at the insurer level, it is empirically difficult to separately identify them using the same variation from the equality moments. For example, if one insurer negotiates a higher price with a particular hospital relative to another insurer, this may be because that insurer's bargaining ability is greater or because that insurer places higher weight on enrollee surplus. The network inclusion moments help separate these. If, conditional on a guess of γ_m, an insurer is observed to cover a hospital, even if doing increases costs "more" than it increases enrollees' WTP, then the implication is that this insurer values its enrollees' surplus quite highly. Figure 4 suggests that insurers are indeed responsive to their enrollees' surplus. Harvard Pilgrim, which has a substantially larger 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.

To identify contracting costs, b, we require additional assumptions. The first assumption is that the contracting costs are identical across all insurer-hospital pairs in our data. This assumption aids in identification in two ways. First, given the limited number of insurers in our sample, it allows us

Figure 3: Hospitals' In-Network Prices Across Insurers

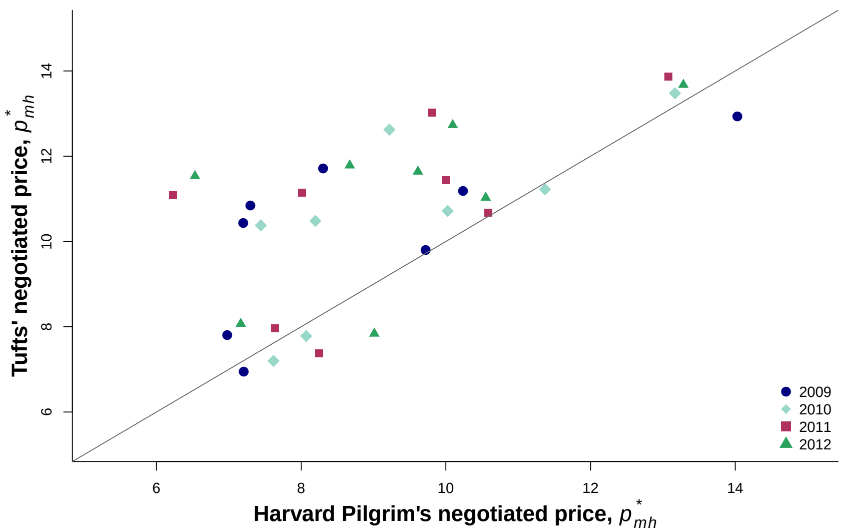

This figure plots the relationship between Harvard Pilgrim's and Tufts Health Plan's price indices for New Hampshire hospitals that are in-network for both insurers. In the majority of hospital-year pairs, Harvard Pilgrim negotiates a lower price than Tufts.

to use information estimated from moment conditions from one insurer in order to inform guesses for other insurers, thereby maximizing efficiency. Second, it implicitly assumes away any structural errors that might bias the estimate of b. This problem is discussed at length in Eizenberg (2014) and Pakes (2010). Specifically, by assuming that the contracting costs are the same across hospitals, we rule out the possibility that coverage decisions are made, in part, because of private information indicating heterogeneity in contracting costs across hospitals. Similar assumptions have been made in Ho (2009) and Nosko (2014). Once the cost of entering into the negotiation is paid, that cost becomes sunk regardless of whether agreement is reached.[17]

A second assumption we make is that b reflects annual *fixed* costs of negotiation that are incurred irrespective of whether an insurer had a contract with a hospital in prior years. 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

[17]This implication does not require any further assumptions. In the model, an insurer-hospital pair will not enter into negotiations in the first place if the expected surplus from agreement (relative to disagreement) is not large enough to offset the bargaining cost.

burden of dealing with contract negotiations adds considerable expense and complexity on both the insurer and provider sides (Wikler et al. 2012). Sometimes, contracting disputes arise between parties that have a history of successful negotiations. Such disputes can require prolonged and costly negotiating before the parties ultimately agree.

While these are not innocuous assumptions, our primary interest in this paper is demonstrating the impact of disagreement values on prices and equilibrium networks. 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. Appendix A reports a robustness check in which we set $b = 0$ and re-estimate all other parameters of the bargaining model. The estimates of the remaining key parameters c_h and γ_m remain quite similar, suggesting that our results are robust to the value of b.

4 Data

In this section, we provide context for our empirical application: the private health insurance market in New Hampshire. We then describe the data used in estimation and the details of sample construction.

4.1 Empirical Setting

Our empirical setting is large New England insurers' negotiations with hospitals in New Hampshire. The insurance market is highly concentrated, with the largest three insurers accounting for at least 85 percent of commercial enrollment throughout our sample period. Two of the top three insurers are large national insurers. As in many states, the top insurer is the local Blue carrier, which is Anthem. Depending on the year, Cigna, another large national carrier, is in second or third place. The third of the top three is Harvard Pilgrim, a smaller, regional carrier that draws the bulk of its enrollment from New England (Prager and Tilipman 2019). The remainder of the insurance market is divided between a number of other regional insurers and small local affiliates of national insurers, such as Aetna and United.

New Hampshire has 32 hospitals, including a Veterans Affairs hospital and five rehabilitation or psychiatric hospitals. We focus on the remaining 26 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 by 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 volumes of Massachusetts residents or locals who are insured by Massachusetts insurers. For example, Harvard Pilgrim was originally based in Massachusetts.

Most insurers with substantial operations in New Hampshire have complete hospital networks within the state. That is, they have negotiated contracts with each of the state's 26 acute care hospitals. Unsurprisingly, among the insurers with complete networks are the three top insurers in the state. This pattern is not peculiar to New Hampshire; it is common for insurers to have locally complete hospital networks for their broadest-network plans.

Outside of New Hampshire's top three insurers, however, some hospital networks are incomplete. Notably, Massachusetts-based Tufts Health Plan, which is among the smaller insurers in the state throughout our sample period, has negotiated contracts with only eight of the state's 26 hospitals. The Tufts network includes four of the five highest-volume hospitals in the state, among them the Dartmouth-Hitchcock flagship hospital. The other four hospitals within Tufts' network are all within a 35-minute drive of the state's southern border with Massachusetts, where the bulk of Tufts' enrollees are located. None of the hospitals in the northern half of New Hampshire is in Tufts' network. The fact that Tufts' network only covers a small share of the New Hampshire market, despite having enrollees residing in the state, plays an important role in identifying parameters in our demand and bargaining models.

4.2 Health Care Claims Data

Data for estimating the hospital choice model and constructing other inputs to the bargaining model are drawn from the 2009–2012 Massachusetts All-Payer Claims Database (APCD). Private health insurers contribute data for the APCD to the state agency that manages the data and uses it for policy-relevant analysis, the Center for Health Information and Analysis (CHIA) (CHIA 2014). The data include privately managed Medicare Part C and Medicaid Managed Care plans, but not traditional Medicare or Medicaid.

The APCD contains approximately 150 million health care claims per year. These include claims originating both within and outside of Massachusetts, as long as they are attributable to enrollees of Massachusetts insurers that contribute data. 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 APCD. Charge prices measure what the provider bills the insurer or the patient. Allowed amounts and insurer paid amounts measure the insurer's contracted price with the provider, in case of an in-network provider with a negotiated price contract; or the amount the insurer pays the provider off-contract, in case of an out-of-network provider. We use the allowed and paid amounts to construct measures of equilibrium negotiated prices for use with the first-order conditions in Equation 3. We use the ratio of paid amount to charge price to infer insurer's out-of-network payment policies. Also reported in the data are amounts for which patients are directly responsible under their insurance plan: deductibles, copays, and coinsurance.

We supplement the APCD 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, 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.

4.3 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 (2018).[18]

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

[18] Many claims databases, including the one used in this paper, 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.

Figure 4: Hospital Networks and Enrollees by Health Plan in New Hampshire

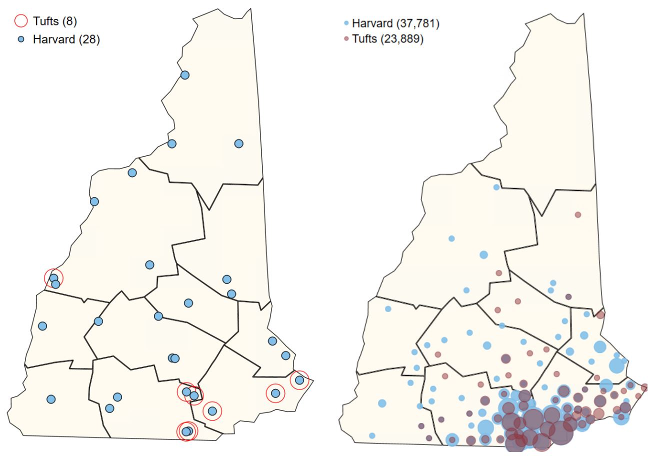

(a) Hospital Networks (b) Enrollees

Notes: Panel (a) plots the hospital networks for Harvard Pilgrim and Tufts Health Plan in New Hampshire. Small blue circles represent Harvard Pilgrim's hospital network and large red circles represent Tufts' hospital network. Panel (b) plots the distribution of enrollees for each health plan from a random sample of members in the state, by 5-digit zip code. Blue circles represent Harvard Pilgrim's enrollees and red circles represent Tufts enrollees.

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. If a hospital is not classified as in-network even in the insurer's broadest-network plans, then it is defined as lacking a contract with the insurer. As described in Section 4.1, the largest insurer with an incomplete hospital network in New Hampshire is Tufts Health Plan.

Figure 4 shows the hospital networks and distribution of enrollees for two carriers in New Hampshire: Harvard Pilgrim and Tufts Health Plan. Figure 4a shows that Harvard Pilgrim has full coverage in the state, whereas Tufts' largest PPO network only covers 8 hospitals. Those hospitals tend to be clustered in the southeastern part of the state, while several counties in the mid-to-northern part of the state have zero coverage. Figure 4b shows the geographic distribution of enrollees for each of those plans, pulled from a random sample of 5,000 members. Harvard Pilgrim's enrollees are widely distributed across the state, whereas Tufts members are concentrated in the southeast, matching the geographic distribution of hospitals covered. However, Tufts also does have some enrollees residing in counties in the northern and western part of New Hampshire, where network coverage is much sparser.

4.4 Outpatient Hospital Sample

In the empirical implementation, we restrict our attention to health care services that are performed in an outpatient, rather than inpatient, setting. We do this for two primary reasons. First, in our sample, out-of-network inpatient hospitalizations are rarer than out-of-network claims for outpatient services. Second, the data we use to construct off-contract prices is based on the FAIR Health outpatient benchmark data (see 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 APCD. Reconstructing DRG classifications from the data without proprietary software would introduce additional noise into our off-contract price measures.

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.

For each procedure in the sample, we assign a measure of resource intensity by merging in Medicare Relative Value Units (RVUs) from the Center for Medicare and Medicaid Services (CMS).[19] RVUs are updated annually by CMS and are used to determine Medicare payment rates for professional services in Part B.[20] RVUs also vary geographically to reflect local variation in resource utilization for particular procedures. As such, a patient living in Boston may have a different RVU weight for a colonoscopy than a patient living in New Hampshire. In our setting, we use RVU as a continuous measure of severity in our demand model.

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 regional New England insurers with a nonnegligible presence in New Hampshire: Harvard Pilgrim and Tufts Health Plan. We have longitudinal data on the hospital networks of each of those carriers. 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. As such, we include 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. We hereafter refer to these as "border zip codes." 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. The outside option is defined as obtaining the same care from a different provider type—typically a physician office—in New Hampshire or in border zip codes in Massachusetts.

Table 2 shows the summary statistics for our final outpatient sample. The first two columns reflect average characteristics for the full sample. Patients in our sample are, on average, 52 years old and seek care for an RVU weight of 7.22. Approximately 47 percent of our sample are insured by Blue Cross Blue Shield, with the remainder evenly split between Harvard Pilgrim and Tufts. On average, patients travel about 11 miles for one of our selected procedures. Column 1 also displays the average hospital characteristics where patients sought care. On average, hospitals have about 200 beds, about 40 percent have a cardiac catheterization lab (often a signal of expensive service

[19] The RVU data can be downloaded from https://www.cms.gov/Medicare/Medicare-Fee-for-Service-Payment/PhysicianFeeSched/PFS-Relative-Value-Files.html.

[20] In this way, they are analogous to DRGs, but for physician services.

Table 2: Outpatient Sample Summary Statistics

	Full Sample		Tufts NH Sample	
	Mean	Std Dev	Mean	Std Dev
	Patient Characteristics			
Age	52.46	16.85	50.04	11.28
Female	0.50	0.50	0.48	0.50
RVU Weight	7.22	2.89	6.75	2.27
BCBS	0.47	0.50	–	–
Tufts	0.26	0.44	–	–
Harvard	0.27	0.44	–	–
Distance in Miles	11.12	10.82	8.27	9.99
In Network Hospital	0.99	0.09	0.93	0.24
	Hospital Characteristics			
Beds	207.83	101.61	178.02	65.18
CathLab	0.39	0.49	0.07	0.25
NICU	0.54	0.50	0.50	0.50
Neuro	0.98	0.14	0.99	0.11
MRI	0.87	0.34	0.99	0.11
Critcal Access	0.04	0.19	0.01	0.12
Teaching	0.38	0.49	0.29	0.45

Notes: Outpatient sample summary statistics 2009-2013. First two columns reflect the full sample, including Massachusetts residents on the border of New Hampshire and New Hampshire residents. Second two columns reflect only New Hampshire residents who are insured by Tufts Health Plan.

lines), about 54 percent have a NICU, and about 40 percent are teaching hospitals.

In the full sample, 99 percent of patients seek care from an in-network hospital for our selected services. This pattern changes, however, when limiting the sample to only Tufts enrollees and only those residing in New Hampshire (the second two columns). Here, patients travel somewhat smaller distances to receive care (about 8 miles), seek care for somewhat lower-intensity procedures, and from hospitals that are notably smaller with fewer expensive service lines. For example, the share of patients going to hospitals with a cardiac catheterization lab in this Tufts sample is only 7 percent. Most importantly, however, the share of procedures performed in-network hospitals drops from 99 percent to 93 percent. This variation is critical for identifying patient disutility from out-of-network hospitals in our demand model.

4.5 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 (2017b), 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.[21] 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_h, 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 are 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: the FAIR Health charge benchmark

[21] Other papers making analogous assumptions include Shepard (2016), Ghili (2017) and our own work in Prager (2018) and Tilipman (2018).

percentiles.[22] We normalize the weights such that $w_d = 1$ for venipuncture (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.

Figure 2, first described in Section 3.5, plots the price indices computed for our two focal insurers across in-network hospitals in New Hampshire. The price indices reflect negotiated prices for procedures with the resource intensity of a routine venipuncture. Both insurers' negotiated prices fall primarily in the $6 to $13 range. However, in most cases, Harvard Pilgrim's negotiated prices are lower than Tufts' for the same hospitals (dashed curves in Figure 2; see also Figure 3). Because Tufts has a narrow network in New Hampshire whereas Harvard Pilgrim has a complete network, Harvard Pilgrim has in its network many hospitals that are out-of-network for Tufts. The hospitals excluded from Tufts' network have disproportionately high negotiated prices. This variation in in-network prices and network status helps to identify the insurers' respective bargaining weights (see Section 3.5).

5 Results

5.1 Hospital Demand Estimates

Table 3 shows the results of the hospital demand model for outpatient care. Consistent with the literature on hospital and physician demand, distance enters negatively and significantly into the utility function. Older patients are less willing to travel for colonoscopies, endoscopies, and arthroscopies, but patients in need of procedures with higher RVU weights (particularly the arthroscopies) are more willing to travel farther distances.

Most of the interactions between patients and hospital characteristics follow the expected signs. Patients are more willing to travel for hospitals with a cardiac catheterization lab, larger hospitals, and teaching hospitals. More puzzling is that patients are more willing to travel to critical access hospitals. This is partially, but not entirely, explained by multicollinearity between critical access

[22] The benchmark construction algorithm is described in detail in Appendix B.

status and bed size, as critical access hospitals are small: the majority of the ones in our sample have 25 beds. 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. Female patients receive more utility from hospitals with neonatal intensive care units.

The key 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 a four additional miles to receive care from an in-network facility as opposed to an out-of-network facility, or about 36 percent farther than the average distance traveled in our sample (11 miles). This preference for hospitals to be in the insurer's network generates positive consumer willingness-to-pay, which then enters into the insurer objective function (Equation 2).

5.2 Hospital Costs and Bargaining Parameters

The first column of Table A.1 shows the results of the bargaining estimation. The estimated hospital costs for routine venipunctures (the baseline procedure with weight $w_d = 1$) in 2010 are all positive, ranging from a low of \$2.27 to a high of \$18.61, with most cost estimates in the \$5–13 range. These are sensible magnitudes. Given that most hospitals in New Hampshire are reimbursed between \$15 and \$20 for this procedure, this suggests that hospitals are, on average, making an markup of 150–300 percent relative to estimated costs, with substantial heterogeneity. For example, Dartmouth Hitchcock Medical Center, a prestigious academic hospital, is estimated to make a markup of about 300 percent in our model, on the higher end of the spectrum.

Harvard Pilgrim Health Care's estimated Nash bargaining weight is 0.99, while Tufts Health Plan's is 0.79, suggesting that on average Harvard Pilgrim is able to extract more surplus from New Hampshire hospitals relative to Tufts. This aligns closely with the fact that for the same procedures, Harvard Pilgrim is observed to pay lower prices than Tufts to the same hospitals. Moreover, Harvard Pilgrim maintains a larger presence in the New Hampshire market than Tufts, both in terms of number of hospitals in-network and enrollment. The estimated MCO weight on consumer surplus relative to spending, α, is approximately 53,121 for Harvard and 105 for Tufts. Though the magnitude of the estimate is difficult to interpret, as our WTP is in utils rather than dollars, its direction is informative. Both Harvard and Tufts are estimated to place strictly positive

Table 3: Results of Hospital Demand

Variable	Utility Parameter	Standard Error
Distance	-0.1525***	0.0066
Distance^2	0.0008***	0.0000
DistxAge	-0.0011***	0.0000
DistxRVU	0.0011*	0.0002
In Network	1.1641***	0.1237
BedsxAge	-0.0001***	0.0000
BedsxRVU	0.0004***	0.0000
BedsxDist	0.0001***	0.0000
CathLabxAge	0.0247***	0.0018
CathLabxRVU	0.0351***	0.0078
CathLabxDist	0.0097***	0.0022
NICUxDist	-0.0094***	0.0016
NICUxFemale	0.1369***	0.0256
NeuroxAge	-0.0091	0.0057
NeuroxRVU	-0.0379*	0.0222
NeuroxDist	-0.0531***	0.0058
MRIxAge	-0.0086***	0.0019
MRIxRVU	-0.0411***	0.0080
MRIxDist	-0.0264***	0.0022
CritAccessxAge	-0.0079*	0.0047
CritAccessxRVU	0.0767***	0.0206
CritAccessxDist	0.0318***	0.0057
TeachingxAge	0.0012	0.0012
TeachingxRVU	-0.0541***	0.0051
TeachingxDist	0.0393***	0.0021
Hospital FE	Yes	
Obs.	1,157,062	
Pseudo R2	0.52	

Notes: ***p<0.01, **p<0.05, *p<0.10. Results from hospital demand model from years 2009-2013. Each observation reflects a visit x hospital pair. “CathLab” refers to whether the hospital has a cardiac catheterization lab. “Neuro” refers to whether the hospital has a neurology unit. “CritAccess” refers to whether the hospital is a critical access hospital.

weights on enrollee surplus relative to costs arising from hospital expenditures. The larger estimate of α for Harvard Pilgrim is driven by its substantially broader hospital network in New Hampshire, particularly in northern regions where enrollment is observed to be minimal. The implication is that Harvard Pilgrim could reduce hospital expenditures by reducing its network breath in northern New Hampshire. The fact that Harvard Pilgrim's network nevertheless includes those hospitals produces a high estimate of α.[23]

Finally, the estimated contracting cost, b, is approximately \$4,593, consistent with estimates in the literature (Ghili 2017). This suggests that the fixed cost of forming and maintaining contracts is non-negligible. However, the contracting cost is not a key driver of our other parameter estimates. In Appendix A, we report these estimates assuming that the bargaining costs are 0. The marginal cost estimates remain very similar, suggesting that the primary channel through which disagreement payoffs affect our model is through their effect on marginal cost estimates, rather than the contracting costs or α.

5.3 Estimates Under Zero Disagreement Volumes

We now turn to the impact that nonzero disagreement volumes have on the estimated cost parameters of the model. To do so, we hold fixed the estimated bargaining weights ($\gamma_{Harvard}$ and γ_{Tufts}), the MCO weight on enrollee surplus ($\alpha_{Harvard}$ and α_{Tufts}), and contracting costs, b, and re-estimate the hospital marginal costs (c_h) under the standard Nash-in-Nash framework.We first remove all out-of-network hospitals from each individual's choice set, and then use the demand model from Table 3 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 from the supply side of the model. For this exercise, we run the bargaining model on a single year (2010, midway through our sample), as this is sufficient to empirically illustrate the implications of nonzero disagreement values (see Section 3.4).

The bargaining model estimates for 2010 using the standard Nash-in-Nash model are reported in Table A.1 column 2. In most cases, incorporating nonzero disagreement volumes into the es-

[23]This is almost certainly an overestimate for Harvard Pilgrim, driven by under-counting of Harvard Pilgrim enrollees in the northern regions of the state. Because we are relying on enrollment data from Massachusetts, our data are skewed towards households in southern New Hampshire and northern Massachusetts. We are currently in the process of obtaining claims data from New Hampshire, which will improve the accuracy of our enrollment counts.

Table 4: Hospital Cost Estimates With and Without Non-Zero Disagreement Volumes

Variable	Full Model	NiN
Hospital Costs (c_h)		
Alice Peck Day Memorial Hospital	6.03	9.15
Androscoggin Valley Hospital	14.56	13.60
Catholic Medical Center	7.64	9.12
Cheshire Medical Center	2.27	6.20
Concord Hospital	9.57	10.13
Cottage Hospital	2.83	6.82
Dartmouth Hitchcock Medical Center	5.93	10.74
Elliot Hospital	9.52	11.15
Exeter Hospital	12.49	13.46
Franklin Regional Hospital	8.42	9.44
Frisbie Memorial Hospital	7.67	10.38
Huggins Hospital	10.60	11.71
Lakes Region General Hospital	7.91	9.13
Littleton Regional Hospital	11.84	11.52
Memorial Hospital	18.61	16.05
Monadnock Community Hospital	10.75	11.13
New London Hospital	5.46	8.18
Parkland Medical Center	3.04	6.73
Portsmouth Regional Hospital	6.32	9.88
Southern New Hampshire Medical Center	4.52	7.35
Speare Memorial Hospital	12.46	12.07
St Joseph Hospital	6.37	8.84
Upper Connecticut Valley Hospital	12.71	12.14
Valley Regional Hospital	12.59	13.49
Weeks Medical Center	10.68	10.93
Wentworth Douglas Hospital	5.14	8.52
Bargaining Weights		
$\gamma_{Harvard}$	0.99	0.99
γ_{Tufts}	0.79	0.79
Bargaining Fixed Costs		
b	\$4,593	\$4,593
MCO weight on WTP		
$\alpha_{Harvard}$	53,121.49	53,121.49
α_{Tufts}	104.67	104.67
Obs.	34	34

Results from bargaining estimation 2010. First column reflects estimates from the full model, allowing for non-zero disagreement volumes and payoffs constructed from Fair Health benchmarks. Second column reflects estimates with disagreement volumes set to zero, as in canonical Nash-in-Nash estimation. All models fix b, γ, and α at their estimated values from the full model. Each observation reflects an insurer-hospital pair. Sample is limited to Harvard Pilgrim, Tufts Health Plan, and only New Hampshire hospitals. Hospital marginal costs reflect a “standardized” cost measure for performing a routine venipuncture.

Figure 5: Bias in Hospital Cost Estimates Under Zero Disagreement Volumes

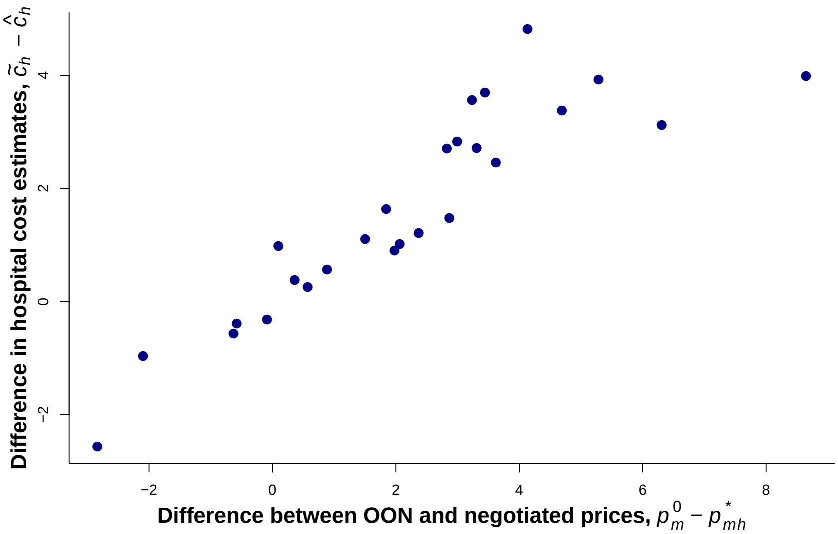

This figure illustrates the direction of the bias arising from assuming zero disagreement volumes using our estimates for 2010. The horizontal axis is the difference between the hospital's out-of-network price and its negotiated price with Harvard Pilgrim, which has a complete network. The vertical axis is the difference between the hospital cost estimates from the standard Nash-in-Nash framework (assuming zero disagreement volumes) and the hospital cost estimates from the full model with nonzero disagreement values. The bias of the Nash-in-Nash estimates increases with the difference between the out-of-network and in-network prices.

timation yields substantially lower hospital marginal cost estimates than assuming that volumes are zero to out-of-network hospitals. The magnitude of the bias is large: on average, standardized 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. Figure 5 plots the empirical analog of the comparative static on prices. It shows that, as out-of-network prices p_m^0 rise above negotiated prices p_{mh}^*, the standard Nash-in-Nash model tends to overestimate hospital marginal costs to a greater degree. Because out-of-network prices are greater than negotiated prices in most markets, the majority of standard Nash-in-Nash estimates of hospital costs are biased upward. Consequently, prior models may have been systematically overestimating hospital marginal costs, limiting the predicted scope of potential policy interventions to reduce hospital reimbursement prices without resulting in exit. In the next section, we detail how such overestimates may affect changes in negotiated rates through a series of counterfactual policy experiments.

6 Policies to Reduce Negotiated Prices

We conduct a series of policy counterfactual simulations using our bargaining model estimates by imposing various restrictions on the out-of-network reimbursement policies and then simulating equilibrium *in-network* negotiated rates between insurers and providers in our sample.

One set of counterfactuals mirrors current federal legislation surrounding surprise out-of-network billing, but applies them more broadly to all out-of-network payments. The Lower Health Care Costs Act of 2019 proposes to regulate surprise out-of-network billing by capping insurers' off-contract payments at median in-network rates in a given market, while also establishing strong balance-billing protections for patients (Alexander 2019). Other policy proposals include fixing out-of-network reimbursements to multiples of Medicare payment rates. 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 rates as low as 120 percent of Medicare (Kane 2019).

Medicare rates are substantially lower than the current standard based on FAIR Health benchmarks. 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 amount (Luthi 2019). As such, we also simulate policies that vary the multiples of the FAIR Health benchmark themselves.

In order to predict the impacts of these policies, we focus specifically on Tufts Health Plan (which has an incomplete network in New Hampshire) and on the year 2010, using our estimates from Column 2 of Table A.1. Under standard Nash-in-Nash, the procedure would involve using our estimated parameters and computing in-network rates, 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 at each iteration t:

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

2. Given the new in-network prices in Step 1, use the network inclusion and exclusion conditions (equations 8 and 9) to check whether any new networks form or whether any existing networks sever. Denote each network link by I^t_{mh}.

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

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

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.

Based on the first-order condition for equilibrium prices (Equation 3), equilibrium in-network prices are linear in counterfactual out-of-network reimbursements. This is because, conditional on which hospitals are in the insurer's network, transaction volumes to each hospital are fixed. Our counterfactual simulations shift p^0_m for all hospitals simultaneously, which will shift a given hospital h's equilibrium price by $\sigma^0_{mh}/\sigma^1_{mh} + (1-\gamma)\psi^0_{mh}$. This linearity is a consequence of hospital demand being independent of price, conditional on network structure. As discussed in Appendix C, this is a sensible approximation for the majority of consumers in our sample. However, if consumers were responsive to price, then p_{mh} would be nonlinear in p^0_m even conditional on the network. Without consumer price responsiveness, nonlinearities in the relationship between p^0_m and p_{mh} can only occur due to changes in the networks themselves.

6.1 Alternate Multiples of Charge Price Benchmarks

We first consider rescaling the disagreement values to be alternate multiples of the current benchmarks (the current benchmark for Tufts Health Plan is the 60th percentile of charges, as described in Appendix B). 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.

The solid blue dots in Figure 6 plot the results of this simulation. In-network negotiated rise with increases in the off-contract prices that insurers pay to out-of-network hospitals. By increasing off-contract prices, hospitals disagreement value is improved, while the insurer's disagreement value

Figure 6: Predicted Negotiated Prices Against Multiples of Current Off-Contract Prices

This figure plots results of counterfactual simulations varying out-of-network prices to be various multiples of current out-of-network payments. The vertical axis plots the counterfactual negotiated price for hospitals predicted to be in-network (volume-weighted average prices). Plot is for Tufts Health Plan in 2010. Gaps represent counterfactuals for which no equilibrium was found.

worsens. Hospitals and therefore they gain considerable bargaining leverage to raise prices. The slope is quite dramatic.[24] At current off-contract prices (multiple of 1.0 on the horizontal axis), the average predicted in-network price is \$10.00 (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 70 percent to an average of about \$16.79. On the other hand, reducing the benchmark to half of the current benchmark would drive predicted in-network rates to below \$6.54, substantially below the median hospital's marginal cost.[25]

While equilibrium price reductions are desirable to policy-makers, access to health care is also an important policy goal. As shown in Appendix Figure A.1, which adds equilibrium networks to the plot, these goals are in direct competition. As negotiated prices fall, so too does the fraction of hospitals that are in the equilibrium network.

Figure 6 also illustrates how conclusions about the counterfactual policies would differ under

[24]Note that in the vicinities of equilibrium network transitions, an equilibrium cannot always be found; this is the source of the gaps in Figure 6.

[25]Such agreements are still possible in equilibrium because the hospital's outside option is to remain out-of-network but still treat some of the insurer's patients at an even lower off-contract price.

estimates from the standard Nash-in-Nash model that assumes zero disagreement volumes. The hollow red triangles plot the results of the same simulation, but using our bargaining model estimates from the last column of Table A.1. Due to the higher estimated hospital marginal costs, the counterfactual in-network prices are always higher than those using our baseline model. Moreover, despite the higher prices, the equilibrium networks are often narrower, as shown in Appendix Figure A.1. This is a good illustration of the importance of accurately estimating hospital costs when conducting policy simulations whose goal is to reduce equilibrium prices. The standard Nash-in-Nash model both overstates equilibrium prices and misstates network breadth. In evaluating a policy proposal, this would cause overly pessimistic predictions about spending and, under some parameter values, about access to care.

6.2 Medicare-Based Out-of-Network Payment Caps

Next, we consider policy proposals that peg insurer reimbursements to out-of-network hospitals at multiples of Medicare reimbursement rates. Medicare reimbursements for the outpatient procedures we study are approximately one quarter of the in-network prices we observe in New Hampshire (see Figure 2), and for many hospitals, less than half of the marginal costs estimated in Table A.1. It is therefore not surprising that most proposals use multiples of Medicare reimbursements greater than one. We simulate the counterfactual equilibrium in-network prices and networks for a range of multipliers strictly above one.

Figure 7 plots the results of this simulation. As before, the solid blue dots represent simulations using the hospital cost estimates that take nonzero disagreement values into account, while the hollow red triangles represent simulations using estimates from the standard Nash-in-Nash model. It is clear from comparing these counterfactuals to Figure 6 that Medicare reimbursements are substantially lower than current off-contract reimbursements: current equilibrium prices are achieved when out-of-network prices are pegged to approximately 400 percent of Medicare.

Negotiated prices are lower using the smaller hospital cost estimates from the model with nonzero disagreement values. As shown in Appendix Figure A.2, equilibrium network breadth is dramatically reduced as out-of-network reimbursements approach Medicare rates. At 200 percent of Medicare, the equilibrium network using our model includes just under half of New Hampshire's 26 hospitals. At 120 of Medicare, the equilibrium network includes only eight hospitals, further

Figure 7: Predicted Negotiated Prices Against Multiples of Medicare Reimbursements

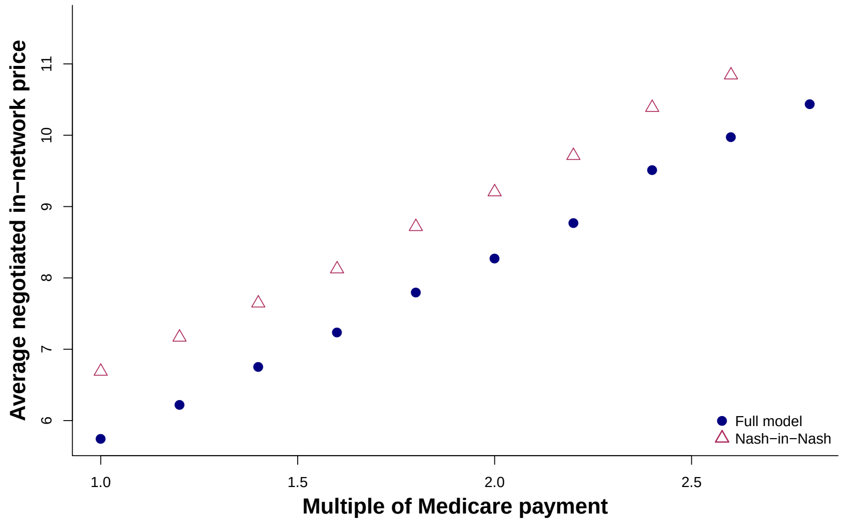

This figure plots results of counterfactual simulations varying out-of-network prices to be various multiples of Medicare reimbursements. The vertical axis plots the counterfactual negotiated price for hospitals predicted to be in-network (volume-weighted average prices). Plot is for Tufts Health Plan in 2010. Gaps represent counterfactuals for which no equilibrium was found.

dropping to seven when we use the cost estimates from the Nash-in-Nash model assuming zero disagreement values.

These results suggest that proposals to peg out-of-network reimbursements to as low as 125 percent of Medicare would likely cause substantial disruptions to provider networks and prices. Equilibrium prices may fall below hospitals' marginal costs, inducing exits or reducing hospitals' capital investment, service availability, and quality of care.

6.3 Forecasting Hospital Closures

The counterfactual simulations discussed in Sections 6.1 and 6.2 allow hospitals to leave insurers' networks in equilibrium. The narrowing of networks that we document is a central concern raised by opponents of regulation to cap out-of-network reimbursements. By contrast, the possibility of outright closures of hospital service lines—or, in extreme cases, of entire hospitals—has received little attention. If reductions in out-of-network reimbursements prompt sufficient reductions in in-network prices, some hospitals may be forced to exit service lines for which prices fall below their marginal costs. In existing models of hospital-insurer bargaining, the assumption of zero out-of-network volumes precludes the possibility of care being reimbursed at below marginal cost. On one hand, a hospital will only enter into a contract if the in-network price exceeds its marginal cost; on the other hand, remaining out-of-network means no marginal costs are incurred. This section evaluates the impact of regulating out-of-network prices on hospital exit.

We proceed by amending the counterfactual simulation algorithm to allow hospitals to exit when price falls below marginal cost. The amended algorithm iterates through the following steps at each iteration t:

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

2. Given the new in-network prices in Step 1, use the network inclusion and exclusion conditions (equations 8 and 9) to check whether any new networks form or whether any existing networks sever. Denote each network link by I^t_{mh}.

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

4. If a link forms and $p^t_{mh} < c_h$; or if a link severs and $p^0_m < c_h$, assign hospital h to exit the market. Denote each closure by C^t_{mh}.

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 $\max_{m,h} \left|p^t_{mh} - p^{t-1}_{mh}\right| < \epsilon$, $\max_{m,h} \left|I^t_{mh} - I^{t-1}_{mh}\right| = 0$, and $\max_{m,h} \left|C^t_{mh} - C^{t-1}_{mh}\right| = 0$, stop. Otherwise, return to Step 1 using the updated p^t_{mh} from Step 1, I^t_{mh} from Step 2, and updated C^t_{mh} from Step 4.

The convergence criterion 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.

Modeling hospital exit in the counterfactuals requires several assumptions. First, we assume that the hospital service lines used in our empirical analyses are separable from hospitals' other service lines (see Section 4.4 for a detailed description of the sample). If that is the case, then price dropping below marginal cost for these service lines is a sufficient condition for a hospital to close the affected service lines. We therefore interpret our hospital closure results as pertaining only to the service lines included in our sample.

Second, we assume that if the focal insurer m's price drops below the hospital's marginal cost, that induces the hospital to exit. This assumption substantially reduces the computational burden of the counterfactuals by avoiding the need to search for multi-insurer equilibria, but it is a simplification in two important ways. 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 rates are generally lower than private insurers' prices (see Section 6.2) 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. However, it remains true that hospitals may cross-subsidize losses from one private insurer's patients using higher prices from a different private insurer. Our counterfactuals do not account for this possibility.

Figure 8 plots the results of the counterfactuals from Sections 6.1 and 6.2, now accounting for hospital exit. Consistent with the earlier results, the fraction of hospitals that remain in the insurer's network (dark green in the figure) drops as out-of-network reimbursements p_m^0 drop and in-network negotiated prices p_{mh}^* follow. Beyond the narrowing networks, however, Figure 8 also makes clear that severe price reductions will also induce some hospitals to close service lines. Capping out-of-network reimbursements at Medicare rates is predicted to induce half 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 lend credence to concerns about providers exiting in response to various payment-reducing policy proposals.

Figure 8: Predicted Networks and Hospital Service Line Closures

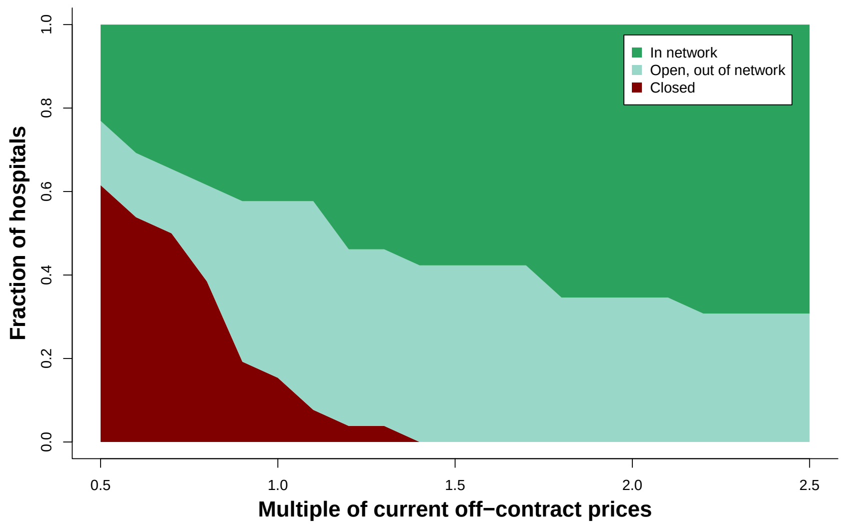

(a) Under Multiples of Current Off-Contract Prices

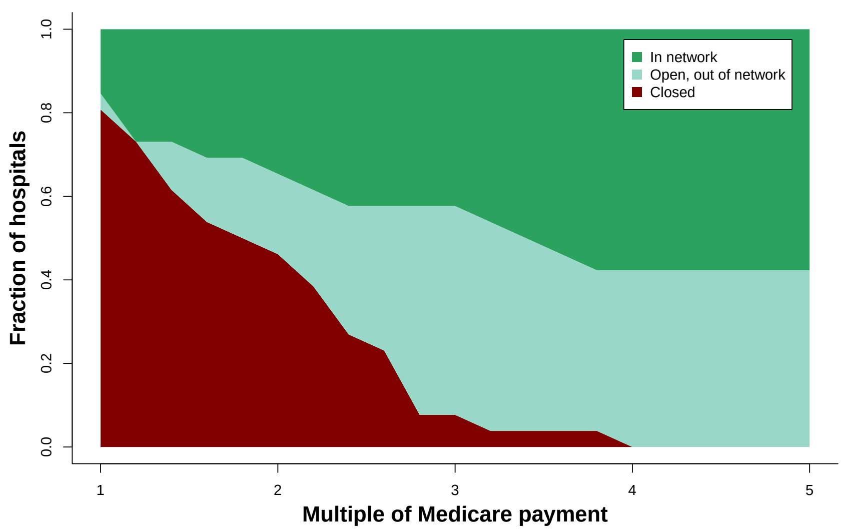

(b) Under Multiples of Medicare Reimbursements

This figure plots results of counterfactual simulations varying out-of-network prices to be various multiples of current out-of-network payments (Figure 8a) or Medicare reimbursements (Figure 8b). 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). Plot is for Tufts Health Plan in 2010, using the estimates from the full bargaining model.

7 Conclusion

Nash-in-Nash bargaining models are a workhorse tool of empirical work studying markets with negotiated prices. While the importance of correctly specifying disagreement values in these models is well understood, there is a practical barrier to measuring prices and transaction volumes in the absence of an agreed-upon contract. This paper proposes a tractable measure of off-contract prices in the context of hospital-insurer negotiations, and uses the measure to evaluate policy proposals surrounding out-of-network hospital reimbursements. Those policy evaluations require a new modeling feature relative to the existing literature: without a way for out-of-network reimbursement rates to enter into the bargaining model, it is not possible to simulate the effects of changing those rates on equilibrium prices and networks.

Incorporating out-of-network transactions into the empirical model results in substantially lower estimates of hospital costs for the majority of hospitals in our data. Because our proposed measure of out-of-network prices is simple to implement in the types of datasets used in the insurer-hospital bargaining literature, it should be straightforward for researchers to correct for this bias in future empirical work without an additional computational burden. This difference in costs also has important implications for the predicted effects of proposed policies. Under a range of counterfactual policies, cost estimates from our model predict lower equilibrium prices and broader equilibrium networks than do cost estimates from the standard model. The counterfactual simulations suggest that policies that cap out-of-network payments at prices close to Medicare rates would severely reduce network breadth, and may even cause hospitals to exit in equilibrium due to in-network prices dropping below marginal costs. Policies that set *all* prices in the health care market to Medicare rates, such as some versions of Medicare For All proposals, may generate even more dramatic market adjustments.

Regulation of health insurers' out-of-network payments is currently limited to a small handful of jurisdictions. As a result, insurers are free to change their policies determining out-of-network prices. If, for instance, hospitals in a market strategically inflate their charge prices in order to raise the benchmark charge prices on which insurers often base out-of-network payments, then insurers can amend their policy to pay a smaller fraction of the benchmark. Policy-makers should therefore consider pairing any regulation of out-of-network payments with regulations that take determination of the benchmark price out of the hands of providers. Pegging to a (large) multiple of Medicare would achieve this goal, whereas pegging to any form of charge prices would not.

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.”

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.

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

CHIA, (Massachusetts Center for Health Information and Analysis) (2014) *All-Payer Claims Database*: Massachusetts Center for Health Information and Analysis.

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, Zach, Fiona Scott Morton, and Nathan Shekita (2019b) “Surprise! Out-of-Network Billing for Emergency Care in the United States,” *Working Paper*.

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

Craig, Stuart, Keith Ericson, and Amanda Starc (2019) “How Important is Price Variation Between Health Insurers?” *Working Paper*.

Crawford, Gregory S. and Ali Yurukoglu (2012) “The Welfare Effects of Bundling in Multichannel Television Markets,” *American Economic Review*, Vol. 102, No. 2.

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

Eizenberg, Alon (2014) "Upstream Innovation and Product Variety in the U.S. Home PC Market," *Review of Economic Studies*, Vol. 81.

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

Ghili, Soheil (2017) "Network Formation and Bargaining in Vertical markets: The Case of Narrow Networks in Health Insurance," *Working Paper*.

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.

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

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 (2017a) "Equilibrium Provider Networks: Bargaining and Exclusion in Health Care Markets," *Working Paper*.

——— (2017b) "Insurer Competition in Health Care Markets," *Econometrica*, Vol. 85, No. 2, pp. 379–417.

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

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

Liebman, Eli (2017) "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*.

Nosko, Chris (2014) "Competition and Quality Choice in the CPU Market," *Working Paper*, Vol. Forthcoming.

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*.

——— (2018) "Consumer Reponsiveness to Simple Health Care Prices: Evidence from Tiered Hospital Networks," *Working Paper*.

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.

Shepard, Mark (2016) "Hospital Network Competition and Adverse Selection," *NBER Working Paper Series, 22600*.

Tilipman, Nicholas (2018) "Narrow Physician Networks, Switching Costs, and Product Variety in Employer Markets," *Working Paper*.

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

Appendices

A Additional Tables and Figures

Figure A.1: Predicted Negotiated Prices Against Multiples of Current Off-Contract Prices

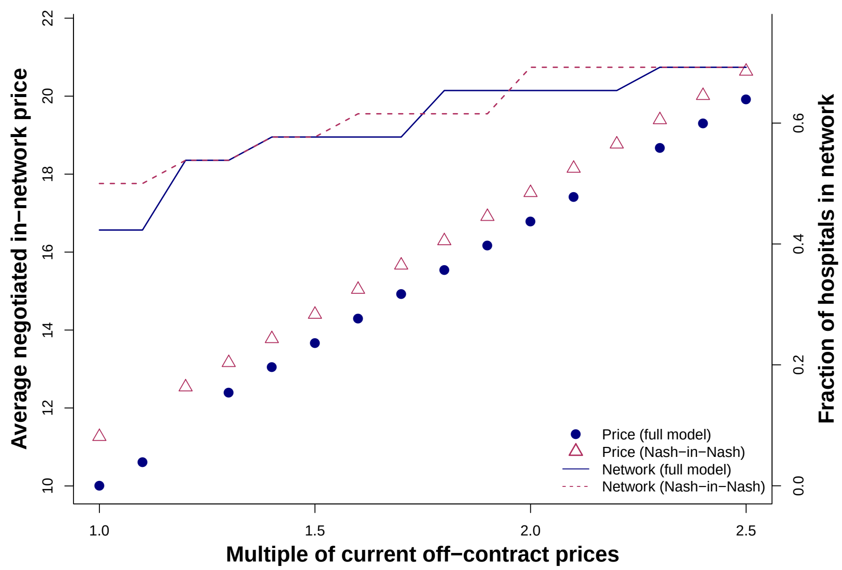

This figure plots results of counterfactual simulations varying out-of-network prices to be various multiples of current out-of-network payments. The vertical axis plots the counterfactual negotiated price for hospitals predicted to be in-network (volume-weighted average prices). Plot is for Tufts Health Plan in 2010. Gaps represent counterfactuals for which no equilibrium was found.

Table A.1: Hospital Cost Estimates With and Without Non-Zero Disagreement Volumes

Variable	Full Model	NiN
Hospital Costs (c_h)		
Alice Peck Day Memorial Hospital	6.06	8.73
Androscoggin Valley Hospital	14.37	13.37
Catholic Medical Center	6.65	8.60
Cheshire Medical Center	1.64	5.47
Concord Hospital	9.15	9.75
Cottage Hospital	2.84	6.25
Dartmouth Hitchcock Medical Center	5.01	9.39
Elliot Hospital	7.88	10.75
Exeter Hospital	11.16	12.86
Franklin Regional Hospital	7.99	9.05
Frisbie Memorial Hospital	7.18	9.92
Huggins Hospital	9.96	11.16
Lakes Region General Hospital	7.32	8.60
Littleton Regional Hospital	9.34	10.56
Memorial Hospital	18.32	15.79
Monadnock Community Hospital	10.21	10.77
New London Hospital	5.50	7.61
Parkland Medical Center	2.02	5.97
Portsmouth Regional Hospital	5.36	9.16
Southern New Hampshire Medical Center	3.61	6.66
Speare Memorial Hospital	12.04	11.65
St. Joseph Hospital	4.61	8.26
Upper Connecticut Valley Hospital	12.49	11.91
Valley Regional Hospital	12.36	13.33
Weeks Medical Center	10.29	10.56
Wentworth Douglas Hospital	4.52	7.93
Bargaining Weights		
$\gamma_{Harvard}$	0.96	0.96
γ_{Tufts}	0.75	0.75
Bargaining Fixed Costs		
b_m	\$0 (fixed)	\$0 (fixed)
MCO weight on WTP		
$\alpha_{Harvard}$	906.05	906.05
α_{Tufts}	249.67	249.67
Obs.	34	34

Results from bargaining estimation 2010 assuming no bargaining costs. First column reflects estimates from the full model, allowing for non-zero disagreement volumes and payoffs constructed from Fair Health benchmarks. Second column reflects estimates with disagreement volumes set to zero, as in canonical Nash-in-Nash estimation. All models fix γ, and α at their estimated values from the full model. Each observation reflects an insurer-hospital pair. Sample is limited to Harvard Pilgrim, Tufts Health Plan, and only New Hampshire hospitals. Hospital marginal costs reflect a “standardized” cost measure for performing a routine venipuncture.

Figure A.2: Predicted Negotiated Prices Against Multiples of Medicare Reimbursements

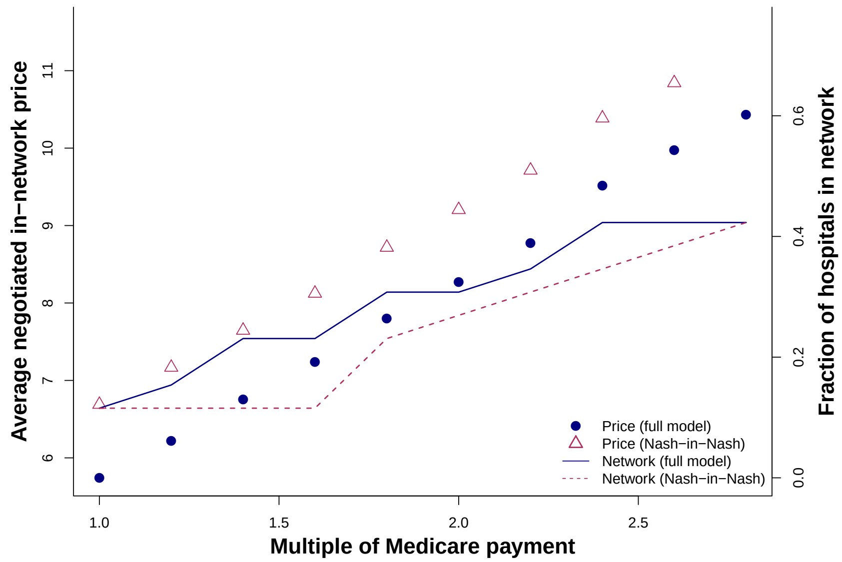

This figure plots results of counterfactual simulations varying out-of-network prices to be various multiples of Medicare reimbursements. The vertical axis plots the counterfactual negotiated price for hospitals predicted to be in-network (volume-weighted average prices). Plot is for Tufts Health Plan in 2010. 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.

B.1 The FAIR Health Algorithm

FAIR Health is the source of charge price benchmarks for many insurers (see Table 1). 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 approximate the outpatient price benchmarks using the near-universe of private insurance claims in New Hampshire from the state's All-Payer Claims Database. As the FAIR Health benchmarks additionally use the universe of fee-for-service Medicare claims, our measure of the benchmark percentiles is somewhat noisy.

However, we follow the FAIR Health benchmark algorithm as faithfully as possible within the available data. 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). Finally, we match the rolling one-year windows and their release dates in May and September.

We are in the process of negotiating a purchase of the proprietary FAIR Health data. If that purchase succeeds, we will update the paper to use the benchmarks from FAIR Health instead of our approximations.

C 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):

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

where λ_h are hospital fixed effects, 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, RVU weights of the procedure, and gender), and interactions between patient characteristics and service availability at hospital h. Here, d is defined at the level of specific medical procedures (CPT codes), and we proxy for it with the RVU weight for the particular procedure, as described in Section 4.5. If consumers prefer to seek care at in-network hospitals, we expect a positive coefficient estimate δ for the in-network indicator. The coefficient δ includes the demand effect of higher expected out-of-pocket payment for out-of-network hospitals.[26] We do not include a finer measure of 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 2018). 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.

This specification 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

[26] The implicit assumption in this specification is that consumers know that they are likely to incur *some* cost for receiving out-of-network care, though they do not necessarily observe what those specific costs are.

14 Massachusetts hospitals, as discussed in Section 4.3).

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

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

where f_{id} is the probability that a consumer of type i requires care for procedure d over the course of a plan-year.[27] 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.

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