

Equilibrium Provider Networks: Bargaining and Exclusion in Health Care Markets[†]

By KATE HO AND ROBIN S. LEE*

We evaluate the consequences of narrow hospital networks in commercial health care markets. We develop a bargaining solution, “Nash-in-Nash with Threat of Replacement,” that captures insurers’ incentives to exclude, and combine it with California data and estimates from Ho and Lee (2017) to simulate equilibrium outcomes under social, consumer, and insurer-optimal networks. Private incentives to exclude generally exceed social incentives, as the insurer benefits from substantially lower negotiated hospital rates. Regulation prohibiting exclusion increases prices and premiums and lowers consumer welfare without significantly affecting social surplus. However, regulation may prevent harm to consumers living close to excluded hospitals. (JEL C78, D85, G22, H75, I11, I13, I18)

Since the passage of the Affordable Care Act (2010), there has been growing concern among policymakers about “narrow network” health insurance plans that exclude particular medical providers. Selective contracting by insurers, in which only particular providers are accessible, is not a new phenomenon. Dating back to the 1980s, managed care insurers have used exclusion to steer patients toward more cost effective or higher quality hospitals and physicians, and to negotiate lower reimbursement rates. While networks broadened somewhat with the “managed care backlash” of the 1990s (Glied 2003), recent high profile exclusions from state exchange plans have reinvigorated the debate over the desirability of such practices.¹ Amid concerns that restrictive insurer networks may adversely affect consumers by preventing access to high-quality hospitals (Ho 2006), or may be used to “cream

*Ho: Department of Economics, Princeton University, Julis Romo Rabinowitz Building, Princeton, NJ 08544, and NBER (email: kate.ho@princeton.edu); Lee: Department of Economics, Harvard University, Littauer Center, Cambridge, MA 02138, and NBER (email: robinlee@fas.harvard.edu). This paper was accepted to the *AER* under the guidance of Esther Duflo, Coeditor. We thank four anonymous referees, numerous individuals (including David Cutler, Liran Einav, Glenn Ellison, Gautam Gowrisankaran, Paul Grieco, Phil Haile, Barry Nalebuff, Aviv Nevo, Ariel Pakes, Mike Riordan, Bill Rogerson, Mark Shepard, Bob Town, Joel Watson, Mike Whinston, Alexander Wolitzky, and Ali Yurukoglu) and conference and seminar participants for helpful comments and discussion. Lee gratefully acknowledges support from the National Science Foundation (SES-1730063). The authors declare that they have no relevant or material financial interests that relate to the research described in this paper.

[†]Go to <https://doi.org/10.1257/aer.20171288> to visit the article page for additional materials and author disclosure statement(s).

¹An Associated Press survey in March 2014 found that Seattle Cancer Care Alliance was excluded by five out of eight insurers on Washington’s insurance exchange; MD Anderson Cancer Center was included by less than half of the plans in the Houston, Texas area; and Memorial Sloan-Kettering was included by two of nine insurers in New York City and had out-of-network agreements with two more. See “Concerns about Cancer Centers under Health Law,” *US News and World Report*, March 19, 2014.

skim” healthier patients, regulators at the state and federal levels are considering formal network adequacy standards for both commercial plans and plans offered on state insurance exchanges.²

Network adequacy standards and other restrictions on network design are essentially a form of quality regulation (Leland 1979, Shapiro 1983, Ronnen 1991). Generally, the welfare effects of such regulation depend on the extent to which the unregulated equilibrium quality (representing network breadth in our setting) diverges from the social optimum or the regulated level, and on any indirect effects of the regulation. The familiar intuition from Spence (1975), that a profit-maximizing monopolist may choose a socially suboptimal level of quality because it optimizes with respect to the marginal rather than the average consumer, applies here. However, features of the US health care market introduce additional complications. In particular, insurers do not bear the true marginal cost of medical care, but rather reimburse medical providers according to negotiated prices. Thus, the “cost of quality” for the insurer is endogenous, and an insurer may sacrifice social or productive efficiency in its choice of network in order to strengthen its bargaining leverage with respect to providers.

In this paper, we examine the private and social incentives for exclusion of hospitals from insurer networks, and consider the potential effects of network adequacy regulations in the US commercial (employer-sponsored) health insurance market. We begin with a simple framework that isolates the fundamental economic trade-offs when deciding whether to exclude a hospital, and identifies the empirical objects required to measure the costs and benefits from exclusion. We then extend the model of the US commercial health care market developed and estimated in Ho and Lee (2017)—which incorporates insurer-employer bargaining over premiums, insurer-provider bargaining over reimbursement rates, and consumer demand for hospitals and health insurers—to capture exclusionary incentives on the part of insurers. Extensions include incorporating a stage of strategic network formation by an insurer and allowing for endogenous outside options in bargaining. Finally, we use our model to simulate equilibrium market outcomes under hospital networks that would be chosen by an agent maximizing social or consumer welfare, or by a profit-maximizing insurer. By comparing outcomes across networks either maximizing different objectives or required to cover all hospitals in a market, we uncover circumstances when private incentives diverge from social or consumer preferences, and evaluate the effects of certain forms of network regulation.

A central component of our analysis is a model of insurer and hospital demand from Ho and Lee (2017). The model, estimated using detailed admissions, claims, and enrollment data from the California Public Employees’ Retirement System (CalPERS) in 2004, provides key empirical primitives necessary to conduct our simulation exercises. In particular, the model predicts how consumers’ insurance enrollment and hospital utilization decisions (inputs into insurers’ revenues and costs) are affected by counterfactual changes in insurer networks; it also conditions

²Such standards are being actively considered, or implemented, by the Centers for Medicare and Medicaid Services (CMS), several state exchanges, and by state regulators such as the California Department of Managed Health Care. See Ho and Lee (forthcoming) and Giovannelli, Lucia, and Corlette (2016) for additional examples and discussion.

on an individual's age, gender, zip code, and diagnosis to accurately capture insurers' incentives for cream skimming and selection, and to ensure that our findings are relevant for "real-world" environments. During the period of our analysis, CalPERS provided access to three large insurance plans for over one million individuals across multiple geographic markets. The plans offered included: a non-integrated HMO offered by Blue Shield of California; a vertically integrated HMO offered by Kaiser Permanente; and a broad-network PPO plan offered by Blue Cross. The Blue Cross PPO network included essentially every hospital in the markets it covered; in 2004 (and historically), the Blue Shield HMO network included most of these hospitals as well. However, in June 2004, the Blue Shield HMO filed a proposal with the California Department of Managed Health Care (DMHC) to exclude 38 "high-cost" hospitals from its network in the following year "as a cost-savings mechanism."³ After the vetting process, 24 hospitals, including some major systems active across California, were permitted to be dropped from the Blue Shield HMO network the following year. We interpret this process as evidence that the DMHC, which evaluates plans' networks to ensure access and continuity of care for enrollees, imposed binding constraints on the hospital networks that insurers were able to offer.

Our paper conducts simulations that adjust the hospital network and reimbursement rates of the Blue Shield HMO plan across 12 distinct geographic markets in California. We view these simulations as predicting the likely effects of removing network constraints imposed by the DMHC on Blue Shield. We also assess the fit of our model by comparing the hospital systems that we predict would be excluded with those Blue Shield proposed to exclude in 2005. Motivated by our empirical setting, we hold fixed the hospital networks offered by Blue Shield's competitors as they are either "complete" and do not exclude any major hospitals (Blue Cross), or integrated (Kaiser) during the time period of our study; however, we allow for all insurers to adjust their premiums as Blue Shield's network adjusts. While our analysis abstracts away from several institutional realities influencing the 2005 Blue Shield proposal, including political constraints and cross-market linkages induced by state-wide premium setting and multi-market hospital systems, our predictions match the observed number and characteristics of excluded hospitals reasonably well.⁴

The key methodological contribution of this paper is the development of a new bargaining solution concept that extends one used in previous empirical work studying insurer-hospital negotiations (e.g., Gowrisankaran, Nevo, and Town 2015; Ho and Lee 2017) and non-health care settings (e.g., Draganska, Klapper, and Villas-Boas 2010; Crawford and Yurukoglu 2012). Commonly referred to as *Nash-in-Nash* bargaining (cf. Collard-Wexler, Gowrisankaran, and Lee forthcoming), the bargaining solution that we build upon predicts that each hospital is paid

³ See Zaretsky and PMPM Consulting Group Inc. (2005) which documents details of the DMHC's analysis of the Blue Shield HMO narrow network proposal.

⁴ There are additional institutional constraints, common in the health insurance industry, that we condition upon in our analysis. In our sample period, CalPERS constrains premiums to be fixed across demographic groups (e.g., age, gender, or risk category), and only allows them to vary based on household size. Consumer cost-sharing at the point of care is limited: for example, Blue Shield charges a fixed copayment for hospital services. These community-rating and cost-sharing restrictions exacerbate insurers' incentives to exclude high-priced hospitals, since premiums cannot easily be increased solely for the consumers who most value such hospitals and consumers are not exposed to price differences across hospitals.

a fraction of its marginal contribution to an insurer's network given the agreements of all other hospitals; an insurer therefore has an incentive to add hospitals to the network in order to reduce each hospital's marginal contribution and, hence, reimbursement. However, Nash-in-Nash bargaining as typically implemented provides limited guidance as to which network(s) emerge in equilibrium, and does not allow for hospitals outside of an insurer's network to influence negotiated payments. This latter limitation may be problematic if an insurer is able to replace an included hospital with another hospital outside the current network during negotiations. For example, under Nash-in-Nash, if an insurer negotiates with only one of two children's hospitals, its only alternative to contracting with that hospital is having *no* children's hospital in its network; it may be more plausible that the insurer is able to form a contract with the other hospital instead.

Our new bargaining solution, *Nash-in-Nash with Threat of Replacement* (NNTR), explicitly addresses this restriction by allowing an insurer to threaten to replace an included hospital with an excluded alternative when bargaining. Intuitively, the NNTR solution assigns each hospital in a particular insurer's network the *minimum* of (i) the price the hospital would obtain with the insurer under simultaneous bilateral Nash bargaining (i.e., the Nash-in-Nash price); and (ii) the price that would make the insurer indifferent between keeping the hospital in its network, and replacing that hospital with any excluded alternative hospital at the minimum price the alternative would be willing to accept. The NNTR solution provides the insurer with an additional incentive, beyond those embedded in the Nash-in-Nash framework, to exclude one or more hospitals in order to reduce the rates negotiated with those that remain. We show that while the Nash-in-Nash solution has difficulty rationalizing *any* exclusion by Blue Shield in 2005, our NNTR solution does not.

To motivate our bargaining solution, we develop and analyze a noncooperative network formation and bargaining game among a single insurer and multiple hospitals under which a particular hospital network and set of prices determined by our NNTR solution emerge as equilibrium outcomes. We prove that the NNTR solution always exists in this setting and detail conditions under which the outcome is unique. Our theoretical work builds upon insights from the literature on bargaining with endogenous outside options, and adapts particular results from Binmore, Shaked, and Sutton (1989) and Manea (2018) to our bilateral contracting environment where firms can negotiate with multiple partners and contracting externalities are pervasive. Formally, we show that an insurer's ability to commit to contracting with fewer hospitals than available, and to go back and forth between hospitals during negotiations, are important factors that bestow upon the insurer the credibility to use excluded hospitals as a way of "bidding-down" the prices for included hospitals.

Overview of Empirical Results

For each of our 12 geographic markets in California, we determine the set of *stable* Blue Shield hospital networks (i.e., networks in which no in-network hospital wishes to terminate their contract with Blue Shield at negotiated reimbursement prices) and report outcomes for the networks that maximize social, consumer, or

Blue Shield's surplus.⁵ Overall, we find that the Blue Shield hospital network that maximizes our measure of social surplus is typically quite broad. In half of our markets, this social-optimal network is predicted to include all major hospital systems; when exclusion occurs, it is primarily to improve the utilization of lower-cost hospitals or insurers and involves the exclusion of a single hospital (although realized welfare gains tend to be modest).

In contrast we predict that both a profit-maximizing insurer and consumers often prefer strictly narrower networks than would maximize social surplus. Blue Shield would wish to exclude at least a single hospital system in two-thirds of our markets, and consumers would prefer exclusion in all but one market (with the caveat that we assume no hospital responds to exclusion by reducing fixed expenditures or exiting the market).⁶ We find that incentives to exclude are not driven primarily by steering or cream-skimming incentives, but rather by rate-reduction and premium-setting motives. Under the Blue Shield-optimal network, the insurer negotiates approximately 12 percent lower hospital prices on average across markets than those predicted to be negotiated if Blue Shield had to contract with all hospital systems in all markets (with reductions up to 30 percent in some markets). Under the consumer-optimal network, average rate reductions are even larger (20 percent) because even more hospitals are excluded. We predict that some of these rate reductions are passed along to consumers in the form of lower premiums, which in turn results in average consumer welfare gains compared to the complete network of approximately \$20–\$28 per capita per year.

We thus establish that bargaining motives introduce an economically meaningful incentive to distort network breadth and quality away from the social optimum. In our setting, an insurer committing to negotiate with a narrow hospital network, combined with an ability to play off included against excluded suppliers, enables the firm to obtain substantial reductions in negotiated input prices. Of course, we acknowledge that the precise magnitudes of our predicted effects rely on the distribution of hospital locations and characteristics observed in our data, our estimated model of insurance demand and hospital utilization, and details of our bargaining solution over premiums and hospital rates. Nonetheless, the “pecuniary incentive” to distort a network away from the social optimum in order to negotiate better rates (a feature absent from other bargaining solutions) is likely to be present in other settings, including environments where firms commit *ex ante* to the number or set of agents with which they will negotiate.

Lastly, we use our results to inform the impact of network regulation in health care markets. Both the magnitude of network (or quality) distortions in the absence

⁵Our social surplus measure is defined to be insurer and hospital revenues plus consumer welfare, minus insurer and hospital marginal costs. Our social-optimal network will correspond to the total welfare-maximizing network if any fixed and sunk costs are not influenced by the counterfactual adjustments in Blue Shield's network that we consider. This assumption may be reasonable if an insurer's network adjustment results in minor changes in utilization or demand (e.g., if network changes for an insurer affect only a single employer), but it does not account for the possibility that hospitals may adjust fixed expenditures or exit following any network changes. All results governing changes in social surplus should thus be caveated appropriately.

⁶CalPERS enrollees represent approximately 10 percent of commercial enrollees in California, and a minority of hospital admissions in the state. As noted by the California Health Care Market Report in 2005, “Sutter and the other large hospital systems can absorb the loss of some CalPERS members because these members comprise only a small percentage of these hospitals' total business in most parts of the state” (Baumgarten 2005, p. 6).

of regulation, and the effect of regulation on prices, are empirically substantive. We find that a requirement that Blue Shield engage in negotiations with all major hospital systems, which we refer to as “full-network regulation,” would actively constrain the insurer in all but four markets. Total surplus would be relatively unchanged from such a regulation because gains to consumers from increased access would be offset by reduced steering to low-cost insurers or providers. Meanwhile Blue Shield’s hospital payments and premiums would increase and consumer welfare would fall. There would also be distributional consequences: consumers who lived closer to excluded hospitals would benefit significantly more than those who did not (many of whom are predicted to be worse off as they would no longer experience premium reductions). In the Sacramento health service area, for example, we find that full-network regulation would benefit consumers in certain zip codes by as much as \$70 per capita per year, while rendering others worse off by up to \$40 per capita per year. In our setting, these amounts are equal to approximately 5–9 percent of annual out-of-pocket premiums for single households.

We draw several important lessons from our evaluation of minimum network standards. First, it is critical to accurately account for premium adjustments in response to quality adjustments by insurers. We find that if premiums were instead fixed and not allowed to adjust when networks changed, consumers would always be harmed by *any* exclusion. Second, as is generally the case with complicated interventions, averages mask considerable heterogeneity. Distributional effects of regulation are likely when consumers differ in their preferences over product attributes. This is the case in our setting because consumers have location-based preferences over hospitals. Regulators should thus be attuned to disproportionate harm borne by particularly vulnerable populations. Third, in the presence of multiple insurers, market forces can discipline an insurer from going “too narrow,” potentially reducing the need for network regulation. In our setting, consumers would actually prefer a narrower network than the insurer-optimal choice. Finally, and relatedly, although an insurer’s incentives to exclude may generally be greater than those faced by a social planner, they may be relatively well-aligned with consumer (and employer) preferences. Regulatory intervention might impede these parties from working together to design customized networks, or engaging in other types of sophisticated plan design, in order to control health care spending.

Applicability of Nash-in-Nash with Threat of Replacement (NNTR)

The NNTR bargaining solution is appropriate for use in our application for several reasons. First, its predictions for surplus division coincide with the widely applied Nash-in-Nash bargaining solution when networks are complete, as they essentially are in the markets that we examine. Hence, the empirical primitives that we import from Ho and Lee (2017) are valid under this paper’s alternative bargaining protocol, and are internally consistent with our application. Second, the channel through which exclusion generates bargaining leverage under NNTR matches institutional realities of the medical care market. Managed care plans have for decades used the threat of exclusion to negotiate low prices with providers. Health economics and policy papers note that actual exclusion “on the equilibrium path,” and not a threat of exclusion that may never be realized (as in the Nash-in-Nash model), is needed

to obtain bargaining leverage.⁷ In some cases, insurer-hospital contracts end due to a dispute but are restarted months later, consistent with insurers going back and forth between hospitals as in the NNTR but not the Nash-in-Nash model.⁸ There is also anecdotal evidence that particular insurers have successfully reduced prices through exclusion. For example, in our setting, the California Health Care Market Report for 2005 notes that “CalPERS... has enjoyed some better results in negotiating rates for 2005, apparently vindicating its strategy to reduce the number of health plan options and to cut some hospitals from its Blue Shield network” (Baumgarten 2005, p. 6).

Our NNTR solution may also be useful in other settings where a single firm can commit to a limited number of trading partners in order to obtain favorable prices. Health care contexts include group purchasing organizations negotiating with medical device suppliers (as in Grennan 2013), employers represented by a single pharmacy benefit manager negotiating formularies with pharmaceutical firms, and commercial insurance settings where employers only offer a single insurance plan.⁹ Non-health care examples include a large retailer negotiating wholesale prices with upstream suppliers (as in Villas-Boas 2007); and media markets where a monopoly distributor negotiates with channels (as in Chitty 2001) or a content provider such as a sports team sells distribution rights to potential distributors.¹⁰ In these settings, the usefulness of NNTR will depend on observing the set of potential contracting partners, controlling for the substitutability of goods and services provided by these parties, and accounting for additional exclusionary motives, which may include fixed costs of contracting and non-constant marginal costs due to capacity constraints.

Related Literature

We contribute to a nascent but growing literature examining the observed relationship between narrow health care networks and utilization, costs, and premiums. Consistent with the bargaining mechanisms discussed in our paper, Gruber and McKnight (2016) find that limited-network insurance plans in the Massachusetts Group Insurance Commission led to almost 40 percent reductions in medical care

⁷ Indeed, a major component of the managed care model was to establish selective networks of providers and induce competition for inclusion. For example, Melnick et al. (1992, p. 219), writing about the (then relatively new) managed care phenomenon, say that, “providers, faced with the pressure to reduce prices or risk being locked out of a payor’s network, must... balance trade-offs in negotiating with selective contracting plans.” Cutler, McClellan, and Newhouse (2000, p. 528) explain that, “since patients generally stay within networks of physicians and hospitals, plans bargain hard for low rates to join the network.” Finally, Corlette et al. (2014, p. 3) note, “if an insurer cannot make a bona fide threat to either exclude a provider from its provider network or place it in a disadvantageous cost-sharing tier, it gives up an important source of leverage in payment negotiations.” The NNTR solution captures the idea that providers are played off against one another in this manner (see Section III).

⁸ For example, in November 2014, Grady Health System was removed from the Blue Cross Blue Shield of Georgia network before being added back the following April; Land of Lincoln announced in January 2017 it would drop University of Chicago medical center as of March, but came to an agreement two months later before doing so; and a seven-month dispute in Ohio between Premier Health and UnitedHealthcare ended in January 2018.

⁹ The Kaiser Family Foundation and HRET’s Employer Health Benefits Annual Survey of 2005 found that 80 percent of firms only offer a single plan to their employees (Claxton et al. 2005).

¹⁰ Collard-Wexler, Gowrisankaran, and Lee (forthcoming) discuss applications of Nash-in-Nash in other one-to-many bilateral oligopoly settings. Also related to our analysis is Lee and Fong (2013), which posits a dynamic formation network game with bargaining in bilateral oligopoly among multiple upstream and downstream firms. It also endogenizes networks and outside options (in the form of continuation values) in order to address similar concerns to those raised here regarding static bargaining models, but focuses primarily on the role of dynamic adjustment costs and contracting frictions.

spending, and that the savings reflected a reduction in provider prices for a fixed quantity of care as well as a reduction in quantity of services. Dafny, Hendel, and Wilson (2015) and Dafny et al. (2017) find that state exchange plans with narrow networks have lower premiums, on average, than other similar plans. Shepard (2016) is closer to the structure of our paper: he uses a model of consumer demand and insurer costs to study the incentives facing Massachusetts exchange insurers to exclude providers in order to select enrollees based on health risk and preferences for high-price hospitals.¹¹ Under fixed hospital prices, he finds that insurers have an incentive to exclude a single large, costly hospital system from their networks. Our framework incorporates similar incentives based on adverse selection, but also includes those generated by bargaining motives as we allow hospital prices and networks to adjust.¹²

Our paper also contributes to the literature on bargaining with exclusionary outcomes in health care settings. Gal-Or (1997) uses a stylized bargaining model to show how two differentiated insurers can secure lower reimbursement rates by jointly foreclosing one of two hospitals. Our NNTR bargaining solution shares similar motivations with Ghili (2018) and Liebman (2018), two recent working papers that allow excluded hospitals to affect insurers' negotiated rates with included hospitals. These papers incorporate their amended bargaining frameworks into an estimated model of the insurer-hospital health care market related to that in Ho and Lee (2017), with the primary objective of quantifying the impact of narrow networks on negotiated prices.¹³ Our focus is on the broader issue of network regulation and its implications, where the impact on insurer-hospital negotiations is one input into the welfare and efficiency considerations that we analyze.

Finally, the nature of our empirical exercise is similar in spirit to Handel, Hendel, and Whinston (2015), which studies the trade-off between adverse selection and reclassification risk; it pairs a theoretical model of insurer competition on a competitive health insurance exchange with estimates of the joint distribution of risk preferences and health status, obtained from an alternative empirical setting, in order to simulate equilibria under hypothetical exchange designs.

¹¹ See also Prager (2018), who shows that similar incentives exist when insurers offer tiered hospital networks in which some hospitals are available at lower co-insurance rates than others.

¹² The cross-hospital variation in prices in our setting is smaller, relative to premiums, than in Shepard (2016). At fixed but higher levels of prices, we also find that hospital exclusion would be profitable for Blue Shield based on selection and steering motives alone.

¹³ There are differences in the approaches. Ghili (2018) relies on pairwise stability conditions (cf. Jackson and Wolinsky 1996), which imply that any insurer and excluded hospital in equilibrium must be unable to generate positive gains-from-trade (given equilibrium prices for included hospitals). To rationalize exclusion, Ghili estimates substantial recurring fixed costs that insurers must pay for including hospitals. In contrast, our NNTR solution rationalizes exclusion in our empirical setting without requiring fixed costs of contracting. Liebman (2018), using a bargaining protocol adapted from Collard-Wexler, Gowrisankaran, and Lee (forthcoming), allows for an insurer to commit to the maximum number of hospitals that it will contract with; upon disagreement in bargaining, a replacement is perceived to be randomly chosen from among an exogenously determined set of acceptable hospitals. In contrast, we allow an insurer to strategically choose the hospital with which to threaten replacement; this implies that increasing the number of potential excluded alternatives does not necessarily improve the insurer's bargaining leverage with a particular hospital (whereas it does so mechanically in Liebman 2018). Our approach also has the advantage of using our model and data to select which hospitals are effective substitutes for others, which may be useful in an environment where contracting partners are considerably heterogeneous. See also Arie, Grieco, and Rachmilevitch (2017), which is a theoretical exercise that incorporates repeated interaction and limits on the number of simultaneous negotiations by the same insurer.

I. Network Design in US Health Care Markets

While the concept of selective contracting has been present in health care markets since at least the emergence of HMO plans in the 1980s, recent publications and press articles suggest that provider network breadth has lessened over time. In a sample of 43 major US markets in 2003, Ho (2006) found that 85 percent of potential hospital-HMO pairs in the commercial market agreed on contracts, suggesting that realized networks were not very selective a decade ago. In contrast, Dafny, Hendel, and Wilson (2015) document that only 57 percent of potential links were formed by HMO plans on the 2014 Texas exchange. The 2017 Employer Health Benefits Survey, released by the Kaiser Family Foundation, suggests narrowing is occurring also in the employer-sponsored health insurance market: for 15 percent of employers offering health benefits, the largest health plan offered had high performance or tiered networks that provided financial or other incentives for enrollees to use selected providers. Nine percent offered a plan considered to be a narrow network plan, and 3 percent of employers reported that their plan eliminated hospitals or a health system to reduce costs.¹⁴

A. The Benefits and Costs of Narrow Networks

Why do insurers choose to exclude medical providers? In this section, we present the fundamental economic trade-offs behind such a decision. We focus on a stylized setting in order to highlight the key reasons why the network choices of an insurer may diverge from those of the social planner.

Key Institutional Details.—Several institutional features of the US private commercial health care market are important to have in mind before continuing. First, different medical providers often negotiate different reimbursement rates with a particular insurer: rate variation in the commercial insurance market is substantial, and may reflect cost and quality differences (Cooper et al. 2018). Second, consumer cost-sharing at the point of care is typically quite limited. After paying a premium, enrollees pay relatively small fees to access providers, and the amount they pay exhibits limited variation across providers. In our setting, hospital co-insurance rates (the percentage of hospital charges that a consumer pays) are zero for both HMO providers. Thus, insurers bear the majority of the incremental price differences when consumers visit a high-priced versus low-priced hospital, and narrow networks may represent an important instrument for insurers to steer patients toward lower-priced (and potentially lower-cost) providers. Third, community rating rules and other premium-setting constraints prevent plans from basing premiums on particular enrollee characteristics (e.g., age, gender, or risk category). These types of requirements are intended to reduce enrollee risk exposure. However, they may also exacerbate insurers' incentives to exclude high-priced hospitals by making it difficult to increase premiums for the groups of enrollees who most value those providers.

¹⁴This survey was released jointly by the Kaiser Family Foundation and the Health Research & Educational Trust. Survey results available at <https://www.kff.org/health-costs/report/2017-employer-health-benefits-survey/>.




FIGURE 1. REMOVING A HOSPITAL FROM AN INSURER'S NETWORK

Notes: Panel A provides demand $D(\cdot)$ and costs $C(\cdot)$ for a hypothetical monopolist insurer offering a product with a given hospital network at fixed premium ϕ . Panel B illustrates new demand $D'(\cdot)$ and costs $C'(\cdot)$ upon the removal of a hospital from the network: areas A and B represent reduction in premium revenues and savings in costs (if the insurer reimburses hospitals at cost); area E represents the reduction in consumer surplus. Panel C depicts potential adjustment in reimbursement prices $P(\cdot)$ to $P'(\cdot)$ upon removal of a hospital: areas A' and B' represent reduction in insurer premium revenues and savings in payments to hospitals.

Baseline Analysis.—To build intuition, consider the incentives facing a monopolist insurer choosing the set of hospitals to include in its network.¹⁵ For now, assume that premiums for this insurer remain fixed, and that the insurer is able to reimburse hospitals at their marginal costs. Hospitals may be differentiated with different qualities and utilities that they generate for patients; they may also have heterogeneous marginal costs. Consistent with limited cost-sharing or lack of price transparency, assume that consumers do not internalize the cost differences between hospitals when choosing providers.

Panel A of Figure 1 depicts a hypothetical demand curve $D(\cdot)$ facing this insurer. At a fixed premium ϕ that it charges for its plan (which is higher than its average costs per-enrollee, but potentially less than the monopoly price if there are premium-setting constraints such as those imposed by regulators or employers), there are q_1 enrollees. Let $C(\cdot)$ represent the (social) marginal costs of insuring each enrollee, including enrollees' drug, hospital, and physician utilization.¹⁶

Consider what might occur if the insurer drops a hospital from its network, depicted by moving from panel A to panel B. There may be several changes. First, the insurer's demand curve shifts inwards from $D(\cdot)$ to $D'(\cdot)$ for at least two reasons. On the intensive margin, enrollees' valuation for the insurer's network decreases, implying a lower willingness-to-pay for the plan. On the extensive margin, some enrollees are likely to leave the plan for the outside option, thereby also changing the identity of the marginal consumer. At the same time as an inward shift in demand, the marginal cost curve might also shift down, particularly if the excluded hospital has a higher

¹⁵ Such an insurer may be thought of as optimizing relative to a nonstrategic outside option: either the choice of no insurance, or an alternative plan or plans whose networks do not respond to this insurer's choices. We thus use the phrase "monopolist insurer" despite the fact that there may exist other insurance plans that consumers view as potential choices.

¹⁶ Note that the marginal cost curve may be downward sloping in the presence of adverse selection.

cost of serving patients than others in the insurer's network. This is due to both the improved steering of enrollees to lower-cost hospitals and the selection of possibly healthier, lower-cost enrollees into the plan. This second effect, commonly referred to as "cream-skimming," will occur if excluding the hospital disproportionately induces higher cost enrollees to switch to the outside option.

If the costs that an insurer faces are given by $C(\cdot)$, which will be the case if it can reimburse medical providers at their respective marginal costs, then a profit-maximizing insurer will choose to exclude the hospital if the size of area A is less than the size of area B in Figure 1, panel B. A represents the loss in premium revenues due to loss of enrollees, and B is the reduction in costs due to both reallocation of patients across hospitals and cream skimming.

However, a social planner would also consider the change in inframarginal consumer surplus for current enrollees if the hospital were removed (represented by E in panel B of Figure 1), a consideration ignored by the profit-maximizing monopolist optimizing over quality (Spence 1975), as well as the loss in social surplus from consumers switching out of the insurance plan and into the outside option. This last object will be significant if the insurer, by dropping a hospital, shifts enrollees to higher-cost plans or to being uninsured (thus potentially resulting in adverse health consequences or spillovers to other parts of the economy). Thus, instead of examining whether $A < B$ (as a monopolist insurer would) to determine whether a hospital should be excluded, a social planner would consider whether $A + E + F < B$, where F is the impact on the outside option (not depicted in the figure).

This analysis highlights the key distortions relative to socially optimal networks if $E + F$ is nonzero, with socially excessive (insufficient) exclusion if their sum is positive (negative). The direction of the distortion is theoretically ambiguous. For example, excessive exclusion can occur if the insurer is more efficient than the outside option and E is large. On the other hand, there may be insufficient exclusion if the outside option is more efficient than the insurer (so that inducing consumers to enroll elsewhere is desirable) and if the insurer's remaining enrollees have a low valuation for the excluded hospital (E is relatively small).

Extending the Analysis.—

Hospital Rate Negotiations: The previous discussion did not distinguish between an insurer's marginal costs and the underlying social cost of providing medical services. It would be reasonable to abstract away from the difference in a setting where insurers reimbursed providers based on marginal costs (perhaps together with a fixed fee transfer). However, in reality, hospitals treating commercial patients are usually paid a price per patient treated (or sometimes per inpatient day), and insurer-hospital pairs engage in pairwise negotiations to determine linear prices: i.e., markups over costs. This feature of the market has important implications for insurer incentives and network choices.

Panel C of Figure 1 illustrates the trade-off facing an insurer if it reimburses providers according to negotiated prices. If excluding a hospital allows the insurer to reduce its marginal reimbursement prices from $P(\cdot)$ to $P'(\cdot)$, then the insurer will exclude if the loss in its premium revenues, given by A' , is less than its savings on reimbursement rates, given by B' . The insurer does not consider the difference

between A and A' , which represents hospital profits.¹⁷ Nor does it consider social cost savings (B), because provider reimbursement rate adjustments do not typically reflect marginal cost adjustments from network changes. As drawn in panel B, $A > B$ so that if the insurer reimbursed providers at cost, it would not wish to exclude the hospital. This coincides with the social planner's preference if $F = 0$, since $A + E > B$. However, in panel C, $A' < B'$, indicating that if the insurer anticipated that excluding a hospital would substantially lower its reimbursement rates, it would choose to do so. Thus, in this example, accounting for the divergence between reimbursement rates and marginal costs leads the insurer to exclude when the social planner would not.

In our subsequent analysis, we show that if an insurer can commit to including or excluding particular hospitals prior to negotiations, this may strengthen its bargaining leverage with those that remain. To the extent that hospital rates are affected by exclusion, there will also be an incentive for the insurer to distort the network away from the industry surplus-maximizing choice: i.e., to “shrink the pie” in order to capture a larger share of it. We consider different bargaining models in our application, and note that since they have different implications for the effect of exclusion on rate negotiations, they also differ in their predictions over the networks that will be chosen by a profit-maximizing insurer.

Adjusting Premiums: Now consider the impact of permitting premiums to vary with the insurer's network. The sign and magnitude of any premium adjustments for the insurer depend on the extent of cost changes for all inframarginal consumers, and on how the elasticity of demand changes for the marginal consumer. If the plans making up the outside option also adjust their premiums in response, this complicates the model further. For these reasons, the breadth of the equilibrium network, and the difference between the monopoly and socially optimal equilibrium outcome, may either increase or decrease once premiums are allowed to adjust. A detailed empirical model of both demand and costs is needed to evaluate these effects.

B. Takeaways

The previous discussion highlights three reasons why a profit-maximizing insurer might choose to exclude a high-cost hospital (e.g., a center of excellence). The first relates to selection or cream-skimming: sick consumers who have an ongoing relationship with the hospital may select out of a plan that excludes it, reducing that plan's costs (Shepard 2016). The second is steering: relatively healthy consumers might prefer to visit the higher-cost provider for standard or routine care if it remains in-network. Excluding the hospital is an effective way for the insurer to steer patients to lower-priced providers. Finally, price negotiations with providers may be affected by network breadth: by excluding some hospitals, the insurer may be able to negotiate lower prices with those that remain.

¹⁷There are also potential issues related to double marginalization, since the premium set by the insurer introduces a second markup in the vertical chain. Since hospital markups differ, the inefficiency of double marginalization may be reduced if high-markup hospitals are excluded from the network. Double marginalization may also imply an additional social gain from prompting consumers to switch to lower-margin outside options.

The discussion also suggests that the network chosen by a profit-maximizing insurer may differ from that preferred by the social planner. A private firm choosing its network breadth will optimize with respect to the marginal rather than the average consumer. Steering patients to low-priced providers may be welfare improving if those providers also have low underlying costs, but this may not always be the case. The welfare effects of cream skimming by one insurer depend critically on the costs and characteristics of other options available to enrollees. In addition, hospital prices are negotiated and may be influenced by the network that is chosen: depending on the particular model of insurer-hospital rate negotiations, this can lead to a “network distortion” either toward or away from the social optimum.

The incentives to exclude, and hence the welfare effects of network regulation, will depend on the characteristics of the particular market (including consumer locations, demographics and preferences, hospital characteristics, and the attributes of the outside option). Accurate empirical estimates of both consumer demand (for insurance plans and hospitals) and health care costs are needed to understand these issues. The demand model must be sufficiently flexible to predict selection of consumers, by health risk and preferences, across providers and insurers when networks change.

Relation to Network Adequacy Regulation and Minimum Quality Standards.—Minimum quality standards may intensify price competition because they require low-quality sellers to raise their qualities, hence reducing product differentiation (e.g., Ronnen 1991). All consumers may be better off as a result of increased quality and reduced hedonic prices compared to the unregulated equilibrium. In our setting, under the interpretation that network breadth may be interpreted as a dimension of insurers’ quality, insurer-hospital rate negotiations imply that the cost of quality provision is endogenous and thereby generates a different intuition. Since insurers may use exclusion to negotiate reduced rates, imposing minimum network requirements may in fact lead to rate increases and corresponding increases in premiums.

We abstract away from possible consumer gains due to minimum quality standards in the presence of incomplete information about product quality (Leland 1979, Shapiro 1983), and assume that consumers are informed about the hospital networks offered by insurers in their choice set. If provider networks are not adequately publicized by insurers, or if consumers are not aware of network composition when making enrollment decisions, there may be benefits from regulation that are outside of the scope of our analysis.

II. Empirical Setting and Overview of Model

The remainder of our paper examines a particular setting in which we quantify the incentives explored in the previous section. Following Ho and Lee (2017), we focus on the set of insurance plans offered by California Public Employees’ Retirement System (CalPERS), an agency that manages pension and health benefits for California state and public employees, retirees, and their families. It is the second-largest employer-sponsored health benefits purchaser in the United States after the federal government; its enrollees represent approximately 10 percent of the total commercially insured population of the state. Our data (described in detail in the

online Appendix) contain the set of insurance plans offered to CalPERS enrollees, their enrollment choices, and medical claims and admissions information in 2004.

A. Empirical Setting

For over a decade starting in 2004, CalPERS employees were primarily able to access plans from three large carriers: a PPO plan from Anthem Blue Cross (BC), an HMO from California Blue Shield, and an HMO plan offered by Kaiser Permanente. During the period of our study, BC was a broad network plan that offered access to essentially every hospital in its covered markets; Blue Shield's network was somewhat narrower, containing approximately 85 percent of the hospitals in BC's network;¹⁸ and Kaiser Permanente, as a vertically integrated entity that owned its own hospitals, had the narrowest hospital network and did not generally allow its enrollees to access non-Kaiser hospitals.

The Blue Shield HMO, along with other managed care plans in California, is regulated by the Knox-Keene Health Care Service Plan Act (1975) and overseen by the California Department of Managed Health Care (DMHC) to ensure, among other things, that the plan meets minimal standards of access and continuity of care. Until 2004, the standards were interpreted as requiring network breadth: insurers were expected to contract with all or most hospitals in their markets.¹⁹ However, in June 2004, Blue Shield filed a proposal with the DMHC to exclude 38 providers, including 13 hospitals from the Sutter hospital system, from their 2005 network. The proposal was vetted by the DMHC for compliance with the accessibility standards set out in the Knox-Keene Act, and described as offering "a vastly different approach to cost savings" compared to other employers' use of copayments, deductibles, or cost sharing (Zaretsky and PMPM Consulting Group Inc. 2005, p. 1). The idea was to "exclude high cost hospitals" from the provider network and hence, presumably by steering patients to lower-cost providers, to provide "alternative mechanisms for the control of rising health care premiums that do not involve greater cost sharing" on the part of consumers.

The DMHC's approval process was intended to verify that the 2005 provider network would provide access to hospital services in each bed service category for CalPERS enrollees.²⁰ Some of the hospitals that Blue Shield proposed to drop were required to be reinstated; these were predominately small community hospitals in relatively isolated communities. In the end, 24 hospitals were excluded from the network in 2005: these are listed in online Appendix Table C.1, along with those included in the original proposal but later withdrawn or denied. Consistent with the motivation of steering patients away from high-cost hospitals, the excluded

¹⁸The number of in-network hospitals only count those with least ten admissions in our data for a particular insurer. We obtain BC hospital network information directly from the insurer; for Blue Shield, we infer the hospital network by including all hospitals that had claims data indicating that the hospital was a "network provider."

¹⁹Insurers were not required to include every hospital in their networks; such a requirement would have substantially reduced their leverage in price negotiations. Instead, regulatory guidelines included specific physician-enrollee ratios and distance standards for the location of providers to enrollee residences and workplaces. The realized outcome was that by far the majority of hospitals were included in both Blue Cross and Blue Shield networks.

²⁰The stated guideline was that enrollees should have access to medical services within 30 minutes or 15 miles of an enrollee's residence or workplace. This rule was applied in terms of distance/travel time between the discontinuing hospitals and other hospitals in the network.

providers included major Sutter hospitals across northern California, and medical centers such as USC University Hospital in Los Angeles. However, other components of the proposal may point to different motivations. The broad geographic spread of the excluded hospitals, in 12 out of 14 health service areas in California, is consistent with the rate-setting motivation: lower rates can be negotiated with an included hospital when a reasonable substitute (a nearby hospital) is excluded and used as a replacement threat. Attempting to exclude academic centers of excellence such as City of Hope National Medical Center (which was eventually withdrawn from the proposal) is also consistent with cream-skimming relatively healthy enrollees. We return to examining the characteristics of excluded hospitals when discussing the results of our empirical simulations.

The setting provided by CalPERS is suitable for studying network design for several reasons. First, we have sufficiently detailed data (from 2004, the year before the Blue Shield network change) to estimate the detailed demand model and cost primitives needed to understand the trade-offs faced by the insurer. Second, we know that Blue Shield offered a relatively broad hospital network in 2004, that it chose to exclude hospitals in the following year, and that it was permitted to exclude fewer hospitals than requested.²¹ Our simulations therefore capture interesting empirical variation because the hospitals offered by Blue Shield in 2004 are likely to differ in terms of costs, and consumer valuations, in ways that generate exclusionary incentives for some but not others. We use the hospitals in the Blue Shield proposal to test the predictive ability of our model. The question of whether potential interventions by the DMHC (or a social planner, more generally) to ensure access are welfare-improving is also clearly empirically relevant. Finally, institutional factors support our modeling assumption that only one insurance plan, the Blue Shield HMO, is able to strategically adjust its hospital network. Motivated by the institutional differences between PPOs and HMOs, we assume that the BC PPO network is exogenously determined and complete.²² Kaiser Permanente is vertically integrated, and cannot flexibly adjust its set of hospitals.

B. Model Overview

To move beyond the abstract discussion of costs and benefits provided in the previous section, and to examine the welfare impact of hospital exclusion and selective contracting, we develop a model of how insurers, hospitals, employers, and consumers interact in the US commercial health care market. We rely on an estimated

²¹The fact that California's regulatory standards with respect to access generated Blue Shield networks that were essentially complete in 2004 is useful for our exercise: when networks are complete, the Nash-in-Nash bargaining model used for estimation in Ho and Lee (2017) has the same outcome as the NNTR model developed in this paper. See Section IVC for further discussion.

²²Preferred Provider Organization (PPO) plans are a specific type of managed care plan that offers greater flexibility and provides greater coverage than a typical HMO. Consistent with this, the coverage offered by HMOs in California, but not most PPOs, is regulated by the Knox-Keene Act (1975). Only a small number of studies investigate the networks of different plan types, but that small literature provides evidence, consistent with our data, that PPO networks are much more complete than those offered by HMOs. For example, a recent McKinsey report on hospital network breadth for plans on health care exchanges found that, in the silver tiered exchange plans they analyzed, 76 percent of ultra-narrow plans were HMOs while only 18 percent were PPOs (McKinsey and Company 2013). Of broad network plans, 67 percent were PPOs while only 19 percent were HMOs. Polsky and Weiner (2015) contains similar evidence relating to plans' physician networks.

version of this model to simulate equilibrium market outcomes if Blue Shield were to adjust its hospital network. For any set of networks chosen by the insurers in a market, the model predicts equilibrium: (i) negotiated hospital prices; (ii) premiums that insurers charge to enrollees; (iii) consumer enrollment in insurance plans; and (iv) consumer utilization of (or “demand” for) hospital services. These objects enable counterfactual profit and welfare evaluation.

We build on the model of the commercial health care market developed in Ho and Lee (2017) that was used to examine the welfare effects of insurer competition. As in this prior work, we condition on the set of insurers, referred to as managed care organizations (MCOs), and hospitals that are available in a market, and assume a one-shot game with the following timing of actions:

- (1a) Network Formation and Rate Determination: MCOs bargain with hospitals over inclusion in their networks and the reimbursement rates that are paid.
- (1b) Premium Setting: Simultaneously with the determination of hospital networks and negotiated rates, the employer and all MCOs bargain over per-household premiums.²³
- (2) Insurance Demand: Given hospital networks and premiums, households choose to enroll in an MCO, determining household demand for each MCO.
- (3) Hospital Demand: After enrolling in a plan, each individual becomes sick with some probability. Individuals who are sick visit a hospital in their network, determining hospital utilization, payments, and costs.

These assumptions approximate the timing of decisions in the commercial health insurance market, in which insurers negotiate networks and choose premiums in advance of each year’s open enrollment period. During that period, households observe insurance plan characteristics and choose a plan in which to enroll for the following year. Individual enrollees’ sickness episodes then arise stochastically throughout the year.

Our point of departure from Ho and Lee (2017) concerns the actions taken by firms in Stage 1a. Whereas Ho and Lee (2017) conditions on the hospital network observed in the data when examining the determination of rates and holds hospital networks fixed in its simulations, we allow an insurer’s hospital network to be endogenously determined. In addition, unlike prior work, we assume that an insurer is able to leverage hospitals that are excluded from its network when bargaining with hospitals within its network.

In the next section, we present our new bargaining solution and network formation protocol that are assumed to take place in Stage 1a. This part of the analysis relies on anticipated actions taken in Stages 1b, 2, and 3 of the model; details of these other stages follow Ho and Lee (2017), and are summarized in Section IV.

²³ See Ho and Lee (2017) and Crawford et al. (2018) for further discussion regarding the assumption that reimbursement rate and premiums are determined simultaneously.

III. Equilibrium Hospital Networks and Reimbursement Rates

In this section, we extend the framework of the US commercial health care market developed in Ho and Lee (2017) by (i) incorporating a bargaining solution that allows an insurer to play hospitals included in its network off against those excluded in order to negotiate more advantageous rates, and (ii) providing a way to predict an insurer's hospital network.

A. Intuition and Overview of Theory

With regards to the determination of hospital rates, our starting point is a commonly used surplus division rule in applied work on bilateral oligopoly. Generally referred to as the *Nash-in-Nash* bargaining solution (cf. Collard-Wexler, Gowrisankaran, and Lee forthcoming), this solution, following its use in Horn and Wolinsky (1988), has been leveraged in several recent applied papers to model bargaining between firms with market power in both non-health care (e.g., Draganska, Klapper, and Villas-Boas 2010; Crawford and Yurukoglu 2012; Crawford et al. 2018) and health care (e.g., Grennan 2013; Gowrisankaran, Nevo, and Town 2015; Ho and Lee 2017) settings.²⁴ Defined for a particular network, the Nash-in-Nash solution in our context specifies that the reimbursement price negotiated between each hospital and MCO solves that pair's Nash bargaining problem given that all the prices for all other bilateral pairs in the network are determined in the same fashion.

However, there are several limitations of the Nash-in-Nash bargaining solution as commonly implemented that restrict its direct application here (cf. Lee and Fong 2013). First, as the Nash-in-Nash solution is often characterized as a surplus division rule for a given network, it typically provides limited guidance as to which network(s) might form in equilibrium.²⁵ Second, the Nash-in-Nash solution does not allow parties to possess outside options or disagreement points that involve adjusting or forming new agreements. In our setting, this implies that when an MCO negotiates with a particular hospital, it can only threaten to drop that hospital while holding its contracting decisions with all other hospitals fixed.

This issue is not innocuous. Failing to accurately account for parties' true outside options and disagreement points when bargaining may lead to erroneous predictions with substantively important economic implications. To understand why, first note that in our setting, the Nash-in-Nash solution implies that any hospital that is excluded from an insurer's network has no direct effect on negotiated rates with in-network providers. Hospitals are reimbursed based solely on their *marginal contribution* to an insurer's given network. That is, holding fixed the other hospitals that are in the MCO's network, a hospital captures a proportion of the incremental

²⁴This solution's name comes from the possibility of interpreting it as a "Nash equilibrium in Nash bargains": i.e., the Nash equilibrium of a game among pairs of firm, with each pair maximizing its Nash product given the actions of other pairs.

²⁵Often, Nash-in-Nash only imposes the restriction that all agreements in a network generate positive "gains-from-trade" for parties involved. Alternative interpretations may also include conditions that restrict all agreements not formed to generate negative gains-from-trade at any contract, holding fixed all other agreements. See footnote 32 for further discussion.

value it generates when contracting with an insurer.²⁶ An MCO thus has an incentive to reduce the marginal contribution of a hospital in order to negotiate lower rates. One effective way of doing so is by including additional hospitals in its network: if hospitals are substitutable, a broader network implies a smaller marginal contribution of every hospital that is added. This tendency toward broader and more inclusive hospital networks partly explains why, as we will show, the Nash-in-Nash solution has difficulty rationalizing observed levels of exclusion in our empirical application.²⁷

We thus depart from a direct application of Nash-in-Nash in order to capture MCOs' exclusionary incentives, and develop a new bargaining solution concept that we refer to as the *Nash-in-Nash with Threat of Replacement* (NNTR) solution. Intuitively, the NNTR solution is interpretable as one in which each bilateral hospital-insurer pair engages in simultaneous Nash bargaining over their combined gains-from-trade (as in Nash-in-Nash); however, crucially, an insurer can threaten not only to drop its bargaining partner, but also to replace it with an alternative hospital that is not on the insurer's network. In a sense, this feature imposes an endogenous cap on Nash-in-Nash prices, where the effectiveness of the cap depends on whether there exists a credible alternative negotiating partner: if there are no alternative hospitals that generate sufficient surplus with the insurer, then the insurer does not obtain any additional leverage in its negotiations with a given hospital. By allowing hospitals that are excluded from an insurer's network to affect the negotiated prices for hospitals that are included in the network, our NNTR solution removes the requirement that an insurer include additional hospitals in order to benefit from their presence when negotiating reimbursement prices. Furthermore, as we will show, this solution can rationalize observed levels of exclusion in our application.

An overview of the rest of this section follows. In Section IIIB, we define the NNTR bargaining solution for any particular hospital network that can be formed. It conditions on the set of premiums that insurers charge and (anticipated) insurance enrollment and hospital utilization decisions of consumers. For a given insurer's network, the NNTR solution is a set of reimbursement prices such that the price for each hospital included in the network is the minimum of (i) the price that hospital would obtain under Nash bargaining with the insurer; and (ii) the price that would make the insurer indifferent between keeping the hospital in the network, and replacing the hospital with some excluded alternative hospital at the alternative's *reservation price*, defined to be the minimum price the alternative would be willing to accept to be included in the insurer's network. Such a solution will only be admissible for networks that we refer to as *stable*, a condition that implies no party has a unilateral incentive to terminate a relationship based on negotiated prices. We prove that stability is equivalent to each hospital included in a network generating higher joint surplus with the insurer than any excluded hospital. Thus, to affect equilibrium bargaining outcomes, excluded hospitals need to generate enough surplus to provide

²⁶In our setting, this value comprises the higher premium revenues that the insurer obtains as a result of having the hospital in its network, net of any cost adjustments borne by the insurer or hospital.

²⁷Nonlinear costs or fixed costs of contracting, abstracted away from here, may increase insurers' incentives to exclude under Nash-in-Nash bargaining.

the insurer with additional leverage in bargaining, but not so much that their exclusion makes the network unstable.

Next, in Section IIIC, we discuss the relation between our NNTR solution and prior work on bargaining with outside options. In particular, we relate our solution to the work of Binmore, Shaked, and Sutton (1989), who provide both theoretical and experimental motivation for why bargaining outcomes adjust in the presence of outside options only when those options are credible to exercise; i.e., they deliver payoffs greater than would be achievable through a bargaining game without them. We also discuss how we extend such intuition, developed in a two-player setting, to our environment characterized with multiple contracting partners, externalities, and endogenous outside options.

Last, in Section IIID, we provide a noncooperative extensive form that, under certain conditions, admits a unique equilibrium network and set of negotiated reimbursement rates. As firms become patient, the network that arises coincides with what we refer to as the *insurer-optimal stable network*, and the negotiated rates converge to the NNTR bargaining solution. This exercise highlights why aspects of the NNTR solution emerge when an insurer is able to commit to negotiate a particular network and “go back and forth” between counterparties until agreements are reached. Being able to go back and forth allows the insurer to use excluded hospitals reimbursed at their reservation prices as credible threats. The exercise also provides support for the reasonableness of the NNTR bargaining solution, how it might emerge in practice, and why, in our empirical application, it is plausible that Blue Shield can commit to and form the network that maximizes its equilibrium profits under NNTR bargaining.

We close this overview with two additional points. First, though the Nash-in-Nash solution may understate the extent to which an insurer can form selective networks and play hospitals off against one another, the NNTR and the Nash-in-Nash bargaining solutions *coincide* when an insurer’s hospital network is complete (i.e., all hospitals are included) because an insurer then would have no alternative out-of-network hospitals to employ as bargaining leverage. Only when networks are incomplete, as is the case with narrow networks, may predictions between the two solutions differ.²⁸

Second, our analysis considers only adjustments to the hospital networks and reimbursement rates for a single MCO, Blue Shield. We condition on the networks and reimbursement rates for the other MCOs. This is motivated by our empirical setting where, as discussed in Section II, the hospital networks offered by Blue Shield’s competitors are either complete for institutional reasons (Blue Cross PPO) or integrated (Kaiser Permanente) and are assumed to be fixed. Our setting also motivates our choice to explicitly adjust the insurer’s outside option, but not those of hospitals, when developing our NNTR bargaining solution: for any hospital negotiating with Blue Shield, it already contracts with Blue Cross and cannot contract with Kaiser. Thus, hospitals do not have alternative insurers with which to threaten to replace Blue Shield.

²⁸ Even when an insurer’s hospital network is not complete, the two solutions may still coincide. This may occur if none of the excluded hospitals generate sufficient levels of surplus if brought in network to replace an included hospital (e.g., if the excluded hospitals are sufficiently high cost or low quality).

B. Nash-in-Nash with Threat of Replacement

Setup and Notation.—In a given market, consider a set of MCOs \mathcal{M} that are offered by an employer, and hospitals \mathcal{H} ; these sets are assumed to be exogenous and fixed. As noted above, we are primarily concerned with the determination of equilibrium hospital networks and reimbursement rates for a single MCO (represented by index j). We condition on, and do not adjust, the networks and reimbursement rates for other MCOs, denoted by $-j$. For exposition in this section, we also initially hold fixed the set of premiums for all MCOs; we thus omit premiums from notation, and re-incorporate them when presenting our full model in the next section.

Let \mathcal{G}_j denote the set of all potential hospital networks that MCO j can form: for a given $G \in \mathcal{G}_j$, we say that hospital i is included on MCO j 's network if $i \in G$, and excluded if $i \notin G$. Denote by $\pi_j^M(G, \mathbf{p}) \equiv \tilde{\pi}_j^M(G) - \sum_{i \in G} (D_{ij}^H(G) \times p_{ij})$ and $\pi_i^H(G, \mathbf{p}) \equiv \tilde{\pi}_i^H(G) + \sum_{n \neq j} (D_{in}^H(G) \times p_{in})$ to be MCO j 's and hospital i 's profits for any network G and vector of reimbursement prices $\mathbf{p} \equiv \{p_{ij}\}_{i \in \mathcal{H}, j \in \mathcal{M}}$, where each hospital-MCO specific price p_{ij} represents a linear payment per admission made by MCO j to hospital i , and $D_{ij}^H(G)$ represent admissions of MCO j 's enrollees into hospital i given MCO j 's network G . These profit functions depend on realized demand and utilization patterns following the determination of hospital networks, prices, and premiums (i.e., Stages 2 and 3 of our industry model), and are taken as primitives for this section's analysis. The key assumptions that we rely upon are that (i) negotiated payments enter linearly into profits, with nonpayment related components of profits (represented by $\tilde{\pi}_j^M$ and $\tilde{\pi}_i^H$) dependent only on the realized hospital network and other objects that are given or held fixed; and (ii) demand for hospital services, $D_{ij}^H(\cdot)$, are not a function of negotiated prices. This last assumption is consistent with limited cost-sharing faced by patients. We take profit and demand functions as primitives for now, and provide explicit parameterizations for them in the next section.

Let $[\Delta_{ij}\pi_j^M(G, \mathbf{p})] \equiv \pi_j^M(G, \mathbf{p}) - \pi_j^M(G \setminus i, \mathbf{p}_{-ij})$ and $[\Delta_{ij}\pi_i^H(G, \mathbf{p})] \equiv \pi_i^H(G, \mathbf{p}) - \pi_i^H(G \setminus i, \mathbf{p}_{-ij})$ denote the *gains-from-trade* to MCO j and hospital i from forming a contract with one another over their respective disagreement points (denoted $\pi_j^M(G \setminus i, \cdot)$ and $\pi_i^H(G \setminus i, \cdot)$), which are profits when i is removed from network G ; these gains-from-trade are computed given that other agreements in $(G \setminus i)$ are formed at prices \mathbf{p}_{-ij} (representing the vector of all prices other than p_{ij}). Additionally, let $[\Delta_{ij}\Pi_{ij}(G, \mathbf{p})] \equiv [\Delta_{ij}\pi_j^M(G, \mathbf{p})] + [\Delta_{ij}\pi_i^H(G, \mathbf{p})]$ denote the total bilateral gains-from-trade (or surplus) created by MCO j and hospital i . One important feature to emphasize is that bilateral surplus between i and j , $[\Delta_{ij}\Pi_{ij}(G, \mathbf{p})]$, does not depend on the level of p_{ij} given our assumptions on profit functions (as any terms affected by or interacted with p_{ij} cancel out).

Definition.—Let $G \in \mathcal{G}_j$ represent the set of hospitals with which MCO j contracts. We define the *Nash-in-Nash with Threat of Replacement* (NNTR) prices for MCO j associated with network G to be a vector of prices $\mathbf{p}^* \equiv \{p_{ij}^*(G, \mathbf{p}_{-ij}^*)\}_{i \in G}$ such that for all $i \in G$:

$$(1) \quad p_{ij}^*(G, \mathbf{p}_{-ij}^*) = \min \{ p_{ij}^{Nash}(G, \mathbf{p}_{-ij}^*), p_{ij}^{OO}(G, \mathbf{p}_{-ij}^*) \},$$

where

$$(2) \quad p_{ij}^{Nash}(G, \mathbf{p}_{-ij}^*) = \arg \max_p [\Delta_{ij} \pi_j^M(G, \{p, \mathbf{p}_{-ij}^*\})]^\tau \times [\Delta_{ij} \pi_i^H(G, \{p, \mathbf{p}_{-ij}^*\})]^{(1-\tau)},$$

is the solution to the bilateral Nash bargaining problem between MCO j and hospital i with Nash bargaining parameter $\tau \in [0, 1]$; and $p_{ij}^{OO}(G, \mathbf{p}_{-ij}^*)$, referred to as the *outside option price*, solves

$$(3) \quad \pi_j^M(G, \{p_{ij}^{OO}(\cdot), \mathbf{p}_{-ij}^*\}) = \max_{k \notin G} [\pi_j^M((G \setminus i) \cup k, \{p_{kj}^{res}(G \setminus i, \mathbf{p}_{-ij}^*), \mathbf{p}_{-ij}^*\})],$$

where $p_{kj}^{res}(G \setminus i, \cdot)$ represents hospital k 's *reservation price* of being added to MCO j 's network $G \setminus i$, and is defined to be the solution to

$$(4) \quad \pi_k^H((G \setminus i) \cup k, \{p_{kj}^{res}(\cdot), \mathbf{p}_{-ij}^*\}) = \pi_k^H(G \setminus i, \mathbf{p}_{-ij}^*).$$

In this definition for $p_{ij}^*(\cdot)$, the price $p_{ij}^{Nash}(\cdot)$ represents the solution to hospital i and MCO j 's bilateral Nash bargain, given that disagreement results in ij 's removal from G with the negotiated payments for all other bargains fixed at \mathbf{p}_{-ij}^* ; and the price $p_{ij}^{OO}(\cdot)$ represents the lowest reimbursement rate that MCO j could pay hospital i so that MCO j would be indifferent between having i in its network, and replacing i with some other hospital k that is not included in j 's network at hospital k 's reservation price.²⁹ Hospital k 's reservation price, in turn, is defined to be the reimbursement rate that k would accept so that it would be indifferent between replacing hospital i on MCO j 's network at this price, and having neither hospital i nor k on MCO j 's network.³⁰

In Lemma A.1 contained in the online Appendix, we prove that (2)–(4) imply that (1) is equivalent to

$$(5) \quad p_{ij}^*(\cdot) = \begin{cases} p_{ij}^{OO}(G, \mathbf{p}_{-ij}^*) & \text{if } \tau[\Delta_{ij} \Pi_{ij}(G, \mathbf{p}_{(ij=0)}^*)] \\ & < \max_{k \notin G} [\Delta_{kj} \Pi_{kj}((G \setminus i) \cup k, \mathbf{p}_{-ij}^*)], \\ p_{ij}^{Nash}(G, \mathbf{p}_{-ij}^*) & \text{otherwise} \end{cases}$$

where $\mathbf{p}_{(ij=0)}^* \equiv \{0, \mathbf{p}_{-ij}^*\}$. In other words, the NNTR solution for a given hospital $i \in G$ is equivalent to the outside option price $p_{ij}^{OO}(\cdot)$ only when there is some other hospital $k \notin G$ that generates bilateral gains-from-trade with MCO j upon replacing hospital i that exceed τ -share of the bilateral gains-from-trade generated between hospital i and MCO j ; otherwise, it is given by the Nash bargaining solution $p_{ij}^{Nash}(\cdot)$.

²⁹The solution can straightforwardly be extended to allow for an insurer to threaten to swap a hospital i with some subset of hospitals (as opposed to a single hospital). In our empirical application, an insurer is allowed to swap an included hospital system with any excluded hospital system.

³⁰In other words, the reservation price holds hospital k to its Nash-in-Nash disagreement point when bargaining with the MCO given agreements in $G \setminus i$ have formed.

In Section IIIC, we discuss why the NNTR price $p_{ij}^*(\cdot)$ for any hospital i can be interpreted as the outcome of a Nash bargain between two parties (MCO j and hospital i), where the disagreement point to the bilateral bargain is hospital i being dropped from MCO j 's network (holding fixed the outcomes of all other agreements), and MCO j has an outside option of replacing i with some excluded hospital at its reservation price; and why this is a credible outside option.

Example.—Assume that there is a single MCO j and two hospitals, i and k ; and the MCO negotiates with hospitals over a price per-admission for inclusion in its network. Assume for simplicity that hospital profits excluding payments received ($\tilde{\pi}_i(\cdot)$ and $\tilde{\pi}_k(\cdot)$ using our notation), which usually include hospital costs, are zero for any network; and that there is only a single consumer who enrolls in the MCO and requires admission to a hospital with certainty. Further assume that $\tilde{\pi}_j^M(\{i\}) - \tilde{\pi}_j^M(\emptyset) = 10$ and $\tilde{\pi}_j^M(\{k\}) - \tilde{\pi}_j^M(\emptyset) = x$, where $x < 10$: i.e., the gains-from-trade the MCO obtains by contracting with hospital i (or hospital k) at a payment of zero, given the alternative of having no hospitals in its network, is 10 (or x). These profits can be interpreted as the premium revenues the MCO is able to obtain from the consumer given its hospital network. Assume the MCO Nash bargaining parameter is $\tau = 1/2$.

Let us compute the objects that are required to construct NNTR prices for the network involving only hospital i ($G = \{i\}$). First, under Nash bargaining with hospital i , the MCO splits its gains-from-trade equally; thus, $p_{ij}^{Nash}(G) = 5$. Next, given our assumptions, hospital k will accept any nonnegative payment for inclusion in the MCO's network; thus hospital k 's reservation price $p_{kj}^{res} = 0$. Lastly, note that $p_{ij}^{OO}(G) = 10 - x$, as this ensures the MCO is indifferent between (i) having only k on its network at $p_{kj}^{res} = 0$, and (ii) having only i on its network at price $p_{ij}^{OO}(G)$ (both outcomes leave the insurer with a surplus of x). Thus, our NNTR price when $G = \{i\}$ will be $p_{ij}^*(G) = \min\{p_{ij}^{Nash}(G), p_{ij}^{OO}(G)\} = \min\{5, 10 - x\}$.

This solution has an intuitive interpretation. If $x < 5$, then the MCO obtains less surplus with hospital k if it pays k its reservation price of 0 (yielding x) than it would with hospital i under standard Nash bargaining with a disagreement point of 0 (yielding 5). The option of contracting with k is thus not a credible threat to exercise (in a sense made more formal below); as a result, the NNTR price with hospital i coincides with the Nash bargaining solution ($p_{ij}^* = 5$). On the other hand, if $x \geq 5$, then the MCO can obtain more from contracting with k at its reservation price than it would get by contracting with i at $p_{ij}^{Nash} = 5$; in this case, the NNTR price with hospital i is $p_{ij}^* = 10 - x$, which is less than the Nash bargaining solution and guarantees the MCO a surplus of x .

Existence.—To guarantee that NNTR prices exist for any network G , we restrict each NNTR price to lie on a compact interval of the real line.³¹ We then establish the following result.

³¹ Formally, we redefine (1):

$$(6) \quad p_{ij}^*(\cdot) = \max\{-\bar{p}, \min\{p_{ij}^{Nash}(G, \mathbf{p}_{-ij}^*), p_{ij}^{OO}(G, \mathbf{p}_{-ij}^*), \bar{p}\}\} \quad \forall i \in G,$$

for some $0 < \bar{p} < \infty$ sufficiently large. This restriction does not affect our analysis: as nonpayment related profits $\tilde{\pi}_j^M(G)$ and $\tilde{\pi}_i^H(G)$ are bounded, defining \bar{p} to be appropriately high (e.g., the maximum value of any firm's

PROPOSITION 1: *For any G and negotiated prices for other MCOs $\mathbf{p}_{-j} \equiv \{p_{ik}\}_{i \in \mathcal{H}, k \neq j}$, there exists a vector of NNTR prices \mathbf{p}^* for MCO j .*

(All proofs are contained in the online Appendix). We do not provide a general proof of uniqueness; however, in the next subsection, we prove that if NNTR prices involve fixed lump-sum payments, they will be unique. In our empirical application, multiple sets of (linear) NNTR prices \mathbf{p}^* have not been found for any network.

Stability.—Although there will exist NNTR prices for any network G , it may be that at such prices, a hospital would prefer to unilaterally reject a contract with MCO j . We interpret NNTR prices associated with those networks to be unreasonable outcomes. We thus focus only on the set of *stable* networks, formalized as follows. We define an agreement $i \in G$ to be *stable* at prices \mathbf{p} if $[\Delta_{ij}\pi_j^M(G, \mathbf{p})] \geq 0$ and $[\Delta_{ij}\pi_i^H(G, \mathbf{p})] \geq 0$, and *unstable* otherwise; we define a network G at prices \mathbf{p} to be *stable* if all agreements $i \in G$ are stable. Thus, if a network G is stable at prices \mathbf{p} , any agreement $i \in G$ does not wish to be unilaterally terminated by either party involved.³²

The following proposition states that examining bilateral surplus is sufficient for determining whether a network G is stable under NNTR prices.

PROPOSITION 2: *Network G is stable at NNTR prices \mathbf{p}^* if and only if, for all $i \in G$,*

$$[\Delta_{ij}\Pi_{ij}(G, \mathbf{p}^*)] \geq \max \left\{ 0, \max_{k \in (\mathcal{H} \setminus G)} [\Delta_{kj}\Pi_{kj}((G \setminus i) \cup k, \mathbf{p}_{-ij}^*)] \right\}.$$

The proposition establishes the equivalence between hospital $i \in G$ preferring to reject an agreement with the MCO at the NNTR price $p_{ij}^*(\cdot)$ (given all other agreements in $G \setminus i$ are established at prices \mathbf{p}_{-ij}^*), and the MCO generating lower bilateral surplus with hospital i than with some excluded hospital k or there being negative total gains-from-trade. To understand this result, note first that if a particular agreement $i \in G$ is unstable given G and prices \mathbf{p}^* , then one of the following must hold: (i) $p_{ij}^*(\cdot) = p_{ij}^{Nash}(\cdot)$, and the Nash bargaining problem represented by (2) has no solution that both parties would agree to accept; or (ii) $p_{ij}^*(\cdot) = p_{ij}^{OO}(\cdot)$, and hospital i would rather reject than accept its outside option price. Case (i) requires that bilateral surplus is negative ($[\Delta_{ij}\Pi_{ij}(G, \cdot)] < 0$). Case (ii), a new source of instability in our setting, is proven to occur whenever $[\Delta_{ij}\Pi_{ij}(G, \cdot)] < [\Delta_{kj}\Pi_{kj}((G \setminus i) \cup k, \cdot)]$ for some $i \in G, k \in (\mathcal{H} \setminus G)$. Thus, for a network to be stable with respect to NNTR prices, the network cannot exclude any hospital that generates greater bilateral surplus than any included hospital.

nonpayment-related profits) implies that if prices are outside this support, then there will be some firm that would prefer not to contract at such prices.

³²Stability does not rule out the possibility that agreements not contained in G would be profitable to form if all agreements in G remained fixed at the same set of prices. We do not require that agreements not included in G must be unprofitable to form, holding agreements in G as fixed: in our setting, we do not believe that it is reasonable to assume that new agreements can be formed without affecting others (as forming a new agreement requires a joint deviation by two agents, whereas terminating only requires a unilateral action).

Note that Proposition 2 and (5) imply that for a given network G , if some hospital k is the best replacement for i (i.e., $k = \arg \max_{h \in (\mathcal{H} \setminus G)} [\Delta_{hj} \Pi_{hj}((G \setminus i) \cup h, \mathbf{p}^*)]$), then k will constrain i 's price under the NNTR solution (i.e., $p_{ij}^*(\cdot) = p_{ij}^{OO}(\cdot)$) without inducing G to be unstable only if

$$(7) \quad \tau \times [\Delta_{ij} \Pi_{ij}(G, \cdot)] \leq [\Delta_{kj} \Pi_{kj}((G \setminus i) \cup k, \cdot)] \leq [\Delta_{ij} \Pi_{ij}(G, \cdot)].$$

Let $\mathcal{G}_j^S \subseteq \mathcal{G}_j$ denote the set of all networks that are stable at NNTR prices, and let $G^{*,ins}$, referred to as the *insurer-optimal stable network*, be the stable network that maximizes insurer's profits at NNTR prices \mathbf{p}^* : i.e., $G^{*,ins} = \arg \max_{G \in \mathcal{G}_j^S} \pi_j^M(G, \mathbf{p}^*)$.³³

C. Relation of NNTR to Literature on Outside Option Bargaining

Note that if MCO j and hospital i were the only parties bargaining, and MCO j 's outside option involved credibly paying some other hospital k (not in j 's network) its reservation price, then $p_{ij}^*(\cdot)$, given by (5), would emerge as the outcome of certain noncooperative implementations of the Nash bargaining solution. For example, Binmore, Shaked, and Sutton (1989) formally examines an extension of the Rubinstein (1982) alternating offers bargaining game to include the possibility that either party can terminate negotiations and exercise an outside option (assumed to generate *less* surplus than the main bargain) after rejecting an offer; they show that the unique subgame perfect equilibrium outcome converges to the Nash bargaining solution as the discount factor approaches one *unless* the outside option is binding, in which case the side for which the option is binding obtains exactly that amount. This insight, which Binmore, Shaked, and Sutton refer to as the “deal-me-out” form of the “outside-option-principle,” emphasizes the different roles that outside options and disagreement points play. In many circumstances, outside options only affect bargaining outcomes if they deliver payoffs greater than would be achievable in the bargaining game without such outside options, and hence become credible for agents to exercise: i.e., by asking to be “dealt out” of the bargain.³⁴

There are two important complications when applying these insights from a two-party bargaining environment to our setting. The first involves allowing for firms to contract with multiple parties and exert externalities on one another (i.e., so that profits for firms may depend on the entire network of agreements that are realized). The second involves appropriately determining what an MCO's outside

³³ If there are multiple sets of NNTR prices for a given network G , \mathcal{G}_j^S represents the set of networks that are stable at some vector of NNTR prices, and the insurer-optimal stable network is the one that maximizes insurer's profits among all NNTR prices for which the network is stable. We assume that the insurer-optimal stable network is unique for the statement of our subsequent results.

³⁴ An alternative treatment of outside options is one in which such options directly affect disagreement points regardless of whether they deliver higher payoffs than would be obtainable in a bargaining game without options. This treatment is referred to as “split-the-difference” by Binmore, Shaked, and Sutton (1989) (as outcomes lie between agreement points and outside options), and can emerge from an extension of the Rubinstein (1982) alternating offers bargaining game without discounting and with an exogenous probability of breakdown between rounds (see also Binmore, Rubinstein, and Wolinsky 1986). Binmore, Shaked, and Sutton provide experimental evidence that deal-me-out performs better than split-the-difference in predicting outcomes.

option is; in particular, when bargaining with hospital i , why an MCO can credibly pay some hospital k (that is not included in network G) its reservation price $p_{kj}^{res}(\cdot)$.

We deal with the first complication by building on the Nash-in-Nash bargaining solution, which admits multiple contracting partners and externalities across agreements. For a given network G , the Nash-in-Nash solution in our health care context is the solution to each hospital i and MCO j 's bilateral Nash bargaining problem, assuming that all other bilateral pairs in $G \setminus i$ come to agreement: i.e., Nash-in-Nash prices satisfy (2) if $p_{ij}^*(\cdot) = p_{ij}^{Nash}(\cdot) \forall i \in G$.³⁵ Our NNTR solution extends Nash-in-Nash by allowing a firm's outside option to involve replacing their current bargaining partner with another that they are not negotiating with. If there is no alternative bargaining partner, which is the case when the network G in question is complete and includes all hospitals, then NNTR and Nash-in-Nash coincide.

The second complication arises when determining MCO j 's outside option when bargaining with hospital $i \in G$. Under the NNTR solution, MCO j can threaten hospital i with replacing it with some hospital $k \notin G$ at k 's reservation price. Since $p_{ij}^*(\cdot) = \min\{p_{ij}^{Nash}(\cdot), p_{ij}^{OO}(\cdot)\}$, such a threat serves as a constraint on the bilateral Nash bargaining solution that would emerge between i and j only if it is binding. It may not be obvious, however, that MCO j can credibly threaten to pay hospital k its reservation price when not contracting with i . That is, why wouldn't k demand more? As we discuss further in the next subsection, in our noncooperative extensive form that generates the NNTR solution, the ability for the MCO to commit to negotiating with a particular subset of hospitals and go back and forth between hospitals when negotiating bestows upon it the ability to play off hospitals that are included against those excluded from its network, and hence credibly use paying k its reservation price as a threat. Indeed, it is important not to think of an MCO *permanently* disagreeing with hospital i when it exercises its outside option to negotiate with hospital k in our setting; as a matter of corresponding to both reality and our alternating-offers extensive form game presented below, it is natural for an insurer to be able to renegotiate with a hospital with which it may have previously disagreed. This is in contrast to alternative bargaining models where negotiations can only be among agents who have never reached a disagreement in the past (e.g., Stole and Zwiebel 1996).

We stress that these insights are predicated on an MCO's outside option when bargaining with any hospital not exceeding the surplus the MCO generates with that hospital. If this were not the case, then the MCO would rather exercise its outside option than reach an agreement; furthermore, the logic underlying why the outside option involves paying a replacement hospital its reservation price would no longer be valid. Focusing on stable networks, thus, not only ensures that included hospitals would accept their NNTR prices, but also that an insurer would credibly be able to negotiate such prices in an appropriately specified bargaining game.

³⁵ Collard-Wexler, Gowrisankaran, and Lee (forthcoming) provide a noncooperative foundation for this solution concept based on an alternating offers bargaining game where agents negotiate fixed fee transfers and can engage in multilateral deviations.

D. Microfoundation for Nash-in-Nash with Threat of Replacement

We now provide and analyze a noncooperative extensive form game, in the spirit of the Nash program (cf. Binmore 1987, Serrano 2005), that yields the insurer-optimal stable network at NNTR prices as the unique equilibrium outcome when firms become sufficiently patient.

Consider the following game among a single MCO j and multiple hospitals:

- At period 0, MCO j publicly announces a network $G \subseteq \mathcal{G}_j^S$, and assigns a separate representative r_i to each hospital $i \in G$. By announcing G , the MCO (i) determines the number of representatives that are active for the rest of the game, which in turn restricts the MCO's network to contain no more than the number of hospitals in G ; and (ii) commits each representative r_i to only negotiate with his assigned hospital i or with any hospital that no other representative has been assigned to.³⁶
- At the beginning of each subsequent period $t > 0$, all representatives who have not yet reached an agreement with any hospital simultaneously and privately choose to “engage” with either the hospital to which they were assigned, or with some other hospital $k \notin G$. For each representative and hospital pair that engage, the representative is selected by nature with probability $\tau \in (0, 1)$ to make a take-it-or-leave-it (TIOLI) offer to the hospital, and with probability $1 - \tau$ the hospital is selected to make a TIOLI offer to the MCO's representative. The party receiving the offer can choose to either accept or reject the offer.
- If an offer between a hospital and some representative r_i is accepted, that agreement between MCO j and the hospital is formed immediately at the agreed upon price, and representative r_i is removed from the game. If an offer is rejected, then in the following period the representative is again able to choose whether to engage with hospital i or some other hospital $k \notin G$, and the game continues.

For establishing the main results in this subsection, we make the following assumptions. We assume that all agents share a common discount factor δ , and representatives (when bargaining) are only aware of the initial network G that is announced and do not observe offers and decisions not involving them in subsequent periods.³⁷ We assume that payments take the form of lump-sum transfers: i.e., if an MCO's representative comes to agreement with hospital h in a given period, a transfer P_{hj} is immediately made for that agreement. At the end of each period, we assume that the

³⁶We exogenously restrict the MCO to announce only *stable* networks for analytical convenience. If the MCO in period j is allowed to announce an unstable network G at period 0 (so that there is no equilibrium of the resulting subgame in which network G forms), there is an issue regarding the specification of representatives' beliefs over which network will form (see also Lee and Fong 2013); an alternative approach would be to allow for MCOs to announce any network $G \in \mathcal{G}_j$, but restrict attention to equilibria in which representatives are publicly assigned only to hospitals in some stable network $G' \subset G$ (with similar restrictions on representatives' behavior).

³⁷In the event that two (or more) representatives choose simultaneously in a given period to engage with the same alternative hospital $k \notin G$, we will assume that this alternative hospital engages in separate negotiations with each representative (without being aware of the other negotiations). In the equilibria that we analyze, the probability that such an event will occur converges to zero as $\delta \rightarrow 1$.

MCO receives $(1 - \delta) \pi_j^M(G^t)$ and each hospital i receives $(1 - \delta) \pi_i^H(G^t)$, where G^t are the set of agreements that have been formed by the end of period t , and all linear prices \mathbf{p} are assumed to be 0; as noted above, all profit functions condition on insurer premiums and subsequent consumer enrollment and utilization decisions.³⁸ We discuss how our results extend to linear prices at the end of this subsection. Last, we assume that the profit for any hospital which does not contract with MCO j is not affected by the MCO's choice of contracting partners (which is typically satisfied when there is a single MCO in any market).

Critically, our noncooperative game relies on two elements which are important for generating the NNTR solution. First, we assume that MCO j is able to commit in period 0 to a set of hospitals to negotiate with by announcing a particular stable network. This will provide the MCO with additional bargaining leverage, as it can credibly exclude hospitals from its network by limiting the total number of contracts that can be formed.³⁹ Second, we allow for the MCO's representatives to be able to make or receive offers repeatedly from the same hospital even when they have all been rejected in the past. This implies that any representative can essentially go back and forth between hospitals across rounds, and effectively play hospitals off against one another, until an agreement is finally reached.

This second feature is shared with the protocol introduced in Manea (2018), which we adapt to a setting where an MCO can potentially contract with multiple hospitals. In the simplest version of Manea's model, a single seller attempts to sell a single unit of a good to two buyers, each with potentially different valuations. At the beginning of each period, the seller can select any potential buyer to engage with. With probability τ the seller then makes a TIOLI offer for the sale of the good to the selected buyer, and with probability $1 - \tau$ the buyer makes a TIOLI offer to the seller. If the offer is rejected, the next period the seller can again choose any potential buyer to negotiate with and the game continues; if the offer is accepted, payment is made and the game ends. Denote the valuations among buyers as v_1, v_2 , with $v_1 \geq v_2$. An application of Proposition 1 in Manea (2018) implies that as $\delta \rightarrow 1$, all Markov perfect equilibria (in which the seller uses the same strategy in any period in which it still has the good) are outcome equivalent: seller expected payoffs converge to $\max\{\tau \times v_1, v_2\}$, and the seller trades with the highest valuation buyer with probability converging to 1 (and equal to 1 if $\tau \times v_1 > v_2$).⁴⁰ In other

³⁸We follow Collard-Wexler, Gowrisankaran, and Lee (forthcoming) in our treatment of profit flows and restriction to lump-sum transfers; these assumptions imply that flow profits do not depend on previously negotiated payments. They also discuss how models relying on separate representatives have been used to motivate the Nash-in-Nash bargaining solution; in our setting, we adopt this approach to avoid issues that might arise if an MCO were able to deviate across multiple bargains (e.g., by simultaneously replacing multiple hospitals with alternatives).

³⁹The insight that an insurer may find it beneficial to commit to a limited set of counterparties with which to bargain is related to Stole and Zwiebel (1996), who show that under an alternative bargaining protocol, a firm may choose to employ more workers than is socially optimal in order to negotiate more favorable wages.

⁴⁰In the equilibria that Manea (2018) analyzes, for high enough δ , if $v_2 > \tau v_1$ then the seller randomizes between selecting buyer 1 and buyer 2 with positive probability and comes to agreement with the buyer that is chosen. Otherwise, the seller only selects and comes to agreement with buyer 1. In the event that buyer 2 is chosen in equilibrium (which happens with probability converging to 0 as $\delta \rightarrow 1$), the seller extracts a payment that converges to v_2 ; this follows even when buyer 2 is chosen to make a TIOLI offer, as a high payment is necessary to ensure that the seller doesn't reject the offer and (with probability approaching 1) negotiate with buyer 1 in the following period.

words, the outcome is equivalent to one in which the seller can determine the price by choosing either to Nash bargain with the highest valuation buyer, or by running a second-price auction among the two buyers. Note that this also corresponds to the “deal-me-out” outcome of Binmore, Shaked, and Sutton (1989) if obtaining v_2 was the seller’s outside option when bargaining with buyer 1.⁴¹

Equilibrium.—We restrict attention to stationary Markov perfect equilibria (MPE) in which for a given announcement G , the strategies for each of the insurer’s representatives r_i for $i \in G$, comprising the potentially random choice of hospital to engage with, and the bargaining actions for a particular hospital choice given nature’s choice of proposer, are the same in all periods in which an agreement for r_i has not yet been reached. We focus on limiting equilibrium outcomes as $\delta \rightarrow 1$, defined for a “family of MPE” containing one MPE for every $\delta \in (0, 1)$.

We first focus on describing behavior following the announcement of a network G in period 0 by MCO j (which we refer to as a period-1 subgame).

PROPOSITION 3: *With lump-sum transfers, there exists a unique vector of NNTR payments for any network G .*

Given this result, it is straightforward to characterize the set of stable networks using Proposition 2. In equilibrium, representatives for both the MCO and hospitals have consistent expectations over the set of agreements that will form following such an announcement.

PROPOSITION 4: *Consider any period-1 subgame in which stable network G is announced in period 0. For any $\varepsilon_1, \varepsilon_2 > 0$, there exists $\underline{\delta} < 1$ such that for $\delta > \underline{\delta}$, there exists an MPE where G forms with probability greater than $1 - \varepsilon_1$ at prices within ε_2 of NNTR payments.*

The intuition for the proof of Proposition 4 is straightforward. First, consider any period-1 subgame where MCO j announces a network G that contains a single hospital. By Proposition 2, for G to be stable, G must contain the hospital i that maximizes the bilateral surplus between MCO j and any hospital. In this case, we can directly apply the results of Manea (2018): in any family of MPEs, in the limit as $\delta \rightarrow 1$, either the MCO obtains τ -share of the total gains-from-trade $[\Delta_{ij}\Pi_{ij}(G, \cdot)]$ (which corresponds to paying hospital i the lump-sum equivalent of the Nash-in-Nash price $p_{ij}^{Nash}(G)$), or—if the outside option is binding so that $[\Delta_{kj}\Pi_{kj}((G \setminus i) \cup k, \cdot)] > \tau[\Delta_{ij}\Pi_{ij}(G, \cdot)]$, where hospital k is the second-highest-surplus creating hospital—the MCO obtains the equivalent of $[\Delta_{kj}\Pi_{kj}((G \setminus i) \cup k, \cdot)]$ by paying only the lump-sum equivalent of $p_{ij}^{OO}(G)$ to hospital i .

Next, consider any period-1 subgame where MCO j announces a network G that contains two or more hospitals. As long as each representative r_i when negotiating

⁴¹ Bolton and Whinston (1993) shows that a similar result obtains (using a different bargaining protocol) when an upstream firm with a single unit to sell bargains with two downstream firms. See also the discussion in Chapter 9 of Osborne and Rubinstein (1990).

for MCO j believes that all other agreements $(G \setminus i)$ will form with sufficiently high probability, it too will reach an agreement with its assigned hospital i in a manner similar to how it would behave if it were the only bargain being conducted. Here, we leverage the assumption that the MCO's representatives cannot coordinate with one another across bargains. To establish existence, we rely on and adapt results from Manea (2018).

Having analyzed behavior in each period-1 subgame, we establish our main result.

PROPOSITION 5: *For any $\varepsilon_1, \varepsilon_2 > 0$, there exists $\underline{\delta} < 1$ such that for $\delta > \underline{\delta}$:*

- (i) *there exists an MPE where the insurer-optimal stable network $G^{*,ins}$ is announced in period 0, and $G^{*,ins}$ forms with probability greater than $1 - \varepsilon_1$ at prices within ε_2 of NNTR payments;*
- (ii) *there exists $\underline{\Lambda} < 1$ such that in any MPE where the announced period-0 network forms with probability $\Lambda > \underline{\Lambda}$, the insurer-optimal stable network $G^{*,ins}$ is announced in period 0, and forms with probability greater than $1 - \varepsilon_1$ at prices within ε_2 of NNTR payments.*

Thus, as long as agents are sufficiently patient, there exists an equilibrium in which the insurer-optimal stable network $G^{*,ins}$ is announced and formed with probability arbitrarily close to 1 at prices arbitrarily close to NNTR prices. Furthermore, restricting attention to MPE in which the announced period-0 network is formed with sufficiently high probability (which rules out equilibria in which one network is announced in period 0, but another network is negotiated over), this outcome is essentially unique.

Linear Prices.—We are able to extend certain noncooperative results, so far presented in this subsection when payments are lump-sum transfers, to settings with linear prices.⁴² In the proof of Proposition 4 (see online Appendix A.3), we show that in any period-1 subgame where some stable network G is announced, if for any $i \in G$ all other agreements in $(G \setminus i)$ are formed at (linear) NNTR prices \mathbf{p}_{-ij}^* , then in any family of MPEs the negotiated price for hospital i upon agreement converges to the NNTR price $p_{ij}^*(G, \mathbf{p}_{-ij}^*)$ as $\delta \rightarrow 1$.

This result provides one interpretation for why $G^{*,ins}$ and associated NNTR prices \mathbf{p}^* may be perceived as reasonable outcomes. Assume that MCO j had previously negotiated its network $G^{*,ins}$ at prices \mathbf{p}^* ; further assume that contracts for individual hospitals come up for renewal one at a time, and firms do not

⁴² Several complications arise with linear prices. In particular, the gains-from-trade between any MCO and hospital depend not only on the set of agreements that other hospitals have reached, but also on their terms; this implies that the (linear) NNTR price for any MCO-hospital pair is a function of both the network and the set of reimbursement rates for other hospitals (as in (1)). Collard-Wexler, Gowrisankaran, and Lee (forthcoming) discuss how similar complications arise when contracts are not restricted to lump-sum transfers in their noncooperative foundation for the Nash-in-Nash solution. These may be possible to address in certain settings: for example, Crawford et al. (2018) prove in their application that the Nash-in-Nash solution with linear fees emerges in any equilibrium of the Collard-Wexler, Gowrisankaran, and Lee (forthcoming) alternating offers bargaining game where all agreements are formed immediately following any history of play.

anticipate future negotiations following the conclusion of any bargain. If a contract renewal for some hospital i was conducted with the MCO according to the previously described protocol based on Manea (2018) (i.e., the MCO could engage with either hospital i or some hospital k where $k \notin G^{*,ins}$ in any period, and nature would then choose either the MCO or hospital selected to make a take-it-or-leave-it offer) and the MCO could not sign up more than one hospital during these negotiations, then a direct application of our results implies that the unique limiting MPE outcome would be the MCO renewing its contract with hospital i at the NNTR price $p_{ij}^*(G, \mathbf{p}_{-ij}^*)$. Given this interpretation, it is also natural, as the NNTR solution does, to treat hospitals not on the network (for which new contracts can be signed at any point in time) differently from hospitals already on the network (which have contracts currently in place that cannot be adjusted).

E. Extensions

Premium Setting.—In our application, we allow for all MCO premiums to adjust in response to potential network changes and rate negotiations. As we noted in Section IIB, we assume in the timing of our model that premiums are negotiated simultaneously (and separately) with networks and reimbursement rates. Our timing assumption enables us to condition on MCO premiums in the preceding analysis (where we have omitted premiums as arguments in each firms' profit function). An equilibrium of our complete model is characterized by a network, set of reimbursement prices, and premiums.

Hospital System Bargaining.—Our discussion to this point has assumed that MCO j negotiates with hospitals individually. In our empirical analysis, we adopt methods developed in Ho and Lee (2017) that allow for insurers to bargain with hospital systems jointly on an all-or-nothing basis. We modify our NNTR protocol in a straightforward fashion, and define NNTR prices over hospital systems as opposed to individual hospitals.⁴³

Multiple Strategic Insurers.—We have defined and motivated the NNTR solution in an environment where a single MCO can adjust its network and negotiate prices, conditioning on the networks and reimbursement rates for other MCOs. Though formally extending and applying our results to admit multiple strategic intermediaries is outside the scope of this paper and the subject of future work, we briefly describe potential approaches and associated challenges here.

The most direct extension would be to assume that each MCO simultaneously announces the set of hospitals that it negotiates with, where prices for all announced MCO-hospital pairs (and not just those involving a single MCO) are subsequently determined by the NNTR solution given by (1)–(4). Alongside appropriate extensions for notions of stability, such an approach would enable preliminary

⁴³ We assume that the MCO and each hospital system negotiate over a single linear price. Each hospital within the system receives a multiplier of this price, where the multiplier is chosen so that the ratio of prices of any two hospitals within the same system is equal to that observed in the data. See Ho and Lee (2017) for further details.

analyses of competitive insurer behavior with both endogenous networks and rates; the intuition for how insurers leverage excluded hospitals to obtain advantageous rates would carry over to this more complex setting.

This extension highlights at least two sets of challenges. The first set is computational: e.g., with a single strategic MCO, only $2^{|\mathcal{H}|}$ networks need be analyzed; if all MCOs can adjust their networks, this increases to $2^{|\mathcal{M}| \times |\mathcal{H}|}$, which may be burdensome in environments with many MCOs and hospitals. The second is conceptual: e.g., the NNTR solution, as defined, only allows one side of the market (MCOs) to exercise outside options involving replacement. In a setting with multiple MCOs, it may be unreasonable to ignore the possibility that hospitals can do the same and leverage excluded *insurers* in negotiations.⁴⁴ Although several bilateral bargaining models where both parties possess exogenous outside options exist (cf. Osborne and Rubinstein 1990, Muthoo 1999), endogenizing such options for both MCOs and hospitals introduces several complications.⁴⁵ One promising route forward is detailed in Lee and Fong (2013), which, though sharing many similarities to the model developed here, takes an alternative approach by studying a dynamic network formation game where agents anticipate future (as opposed to contemporaneous) adjustments to the network while bargaining. Developing and applying theoretically grounded and tractable bargaining solutions that allow all firms to threaten to replace counterparties remain fertile areas for research.

IV. Empirical Analysis

The previous section introduced and defined the Nash-in-Nash with Threat of Replacement (NNTR) bargaining solution and proved that it emerges as the equilibrium of a noncooperative game. The NNTR solution conditions on the hospital networks in a market and each firm's profits given these networks. In turn, these profits condition on insurer premiums and the enrollment and hospital utilization decisions of consumers.

In this section we provide explicit parameterizations for MCO and hospital profits that we employ in our empirical analysis. We then discuss Stages 1b, 2, and 3 of the model outlined in Section IIB, which determine premiums and consumer demand for MCOs and hospitals, and the estimation of their structural parameters. All model components discussed in this section are developed in Ho and Lee (2017); further details are contained there and summarized in online Appendix B.

⁴⁴In addition, agents' perceived disagreement points (which also affect reservation prices) may depend on their beliefs over the networks formed by all MCOs.

⁴⁵For instance, consider a hospital i negotiating with MCO j . If hospital i replaces MCO j with an alternative MCO (thereby exercising its outside option), does it anticipate that MCO j is also able to replace i with an alternative hospital? As there are contracting externalities, the answer may influence the identity of hospital i 's optimal replacement for MCO j . Furthermore, it will be critical to delineate the circumstances under which hospital i and MCO j can credibly replace each other with alternatives at their reservation prices; in particular, under certain information structures or timing assumptions, it may be that an alternative contracting partner can maintain some bargaining leverage rather than being forced down to its reservation price.

A. Firm Profits

We assume that MCO j 's profits are

$$(8) \quad \pi_j^M(G, \mathbf{p}_j, \phi) = \sum_m \left(\underbrace{\phi_j \times \Phi' \times \mathbf{D}_{jm}(\cdot)}_{\text{premium revenues}} - \underbrace{D_{jm}^E(\cdot) \times \eta_j}_{\text{non-inpatient hospital costs}} - \underbrace{\sum_{h \in G_{jm}} (D_{hj}^H(\cdot) \times p_{hj})}_{\text{hospital payments}} \right),$$

where G now represents the set of *all* MCOs' hospital networks in all markets; \mathbf{p}_j is the set of hospital prices MCO j has negotiated with its in-network hospitals; and $\phi \equiv \{\phi_k\}$ is the set of all MCOs' single household premiums. The profit function sums across markets (indexed by m) and conditions on demand functions $\mathbf{D}_{jm}(\cdot)$ and $D_{jm}^E(\cdot)$, representing household and individual demand for MCO j in market m , and $D_{hj}^H(\cdot)$, the number of enrollees in MCO j who visit hospital i . These demand functions arise from the consumer demand systems that are outlined in Section IVB.

The first term on the right-hand side of (8) represents total premium revenues obtained by MCO j . It accounts for the different premiums charged to different household types: ϕ_j is MCO j 's premium charged to single households, $\Phi = [1, 2, 2.6]$ is the observed vector of premium multipliers for each household "type" (i.e., single, two-party, or family), and $\mathbf{D}_{jm}(\cdot)$ is a vector containing the number of households of each type enrolled in MCO j . The second term makes the distinction that an MCO's non-inpatient hospital costs η_j are incurred on an *individual* and not a household basis, and thus are multiplied by $D_{jm}^E(\cdot)$. The third term represents payments made to hospitals in MCO j 's network for inpatient hospital services; it sums, over all in-network hospitals for MCO j in market m (denoted G_{jm}), the price per admission (p_{hj}) negotiated with the hospital multiplied by the number of patients admitted (D_{hj}^H).

We assume that profits for a hospital i (active in only a single market m) are

$$(9) \quad \pi_i^H(G, \mathbf{p}_i, \phi) = \sum_{n \in G_i} D_{in}^H(G, \phi) \times (p_{in} - c_i),$$

which sums, over all MCOs n with which hospital i contracts (denoted G_i), the number of patients it receives multiplied by an average margin per admission (where c_i is hospital i 's average cost per admission for a patient). Consistent with the presence of limited cost sharing, we assume that utilization of a particular hospital h on MCO j , represented by $D_{hj}^H(\cdot)$, does not depend on the set of negotiated reimbursement rates \mathbf{p} .

Any components of firm profits that are not network varying (such as fixed or sunk costs) are omitted from notation, and do not affect firms' decisions that we model. Nevertheless, they may affect the interpretation of results concerning social surplus; we discuss this further in Section V.

B. Premium Bargaining, Consumer Demand, and Estimation

We now summarize Stages 1b, 2, and 3 of our model in reverse order.

Stages 2 and 3: Insurer and Hospital Demand.—Our model of household demand for insurers and consumer demand for hospitals is based on previous applications of discrete choice models to health care settings (Town and Vistnes 2001; Capps, Dranove, and Satterthwaite 2003; Ho 2006). We use this demand system to predict insurance enrollment $\{\mathbf{D}_{jm}(\cdot), D_{jm}^E(\cdot)\}$ and hospital utilization $D_{hj}^H(\cdot)$ across MCOs and hospitals for any set of hospital networks and insurance premiums.

In Stage 3 of our model, with some age-gender-specific probability, each individual requires admission to a hospital for one of six diagnoses and chooses the hospital in her insurance plan's network that maximizes her expected utility. We allow preferences for hospitals (and their associated characteristics) to differ by diagnosis, income, and location. Given anticipated hospitalization needs in Stage 3, households in Stage 2 of our model choose to enroll in the insurance plan that maximizes their members' expected utilities. Our specification allows for household premium sensitivities to differ by income, and for households to weight the utility that each member expects to derive from a plan's network differently by age and gender (in addition to allowing the probability of admission to a hospital, and of particular diagnoses conditional on admission, to vary along these dimensions).

Our demand specifications capture selection of consumers across both hospitals and MCOs based on age, gender, diagnosis, income, and zip code (all of which are observed in the data). The hospital demand model allows different types of enrollees to be differentially sensitive to proximity to particular included hospitals, and these estimates feed into the household's expected utility from each MCO's network (which differs by zip code and family type) in insurer demand. The richness of this model allows us to capture the sensitivity of each household type's enrollment choices to network breadth; this will be an important input into our simulations. Thus, while we follow most of the previous literature by assuming there is no selection across insurance plans or hospitals based on unobservable consumer preferences, there are multiple paths by which differential responses by consumers to changes in an insurer's hospital network can affect its incentives to exclude.⁴⁶

Stage 1b: Premium Bargaining.—Each MCO negotiates its premium with the employer via simultaneous bilateral Nash bargaining, where the employer maximizes its employees' welfare minus its total premium payments, and MCOs maximize profits. Negotiations take place simultaneously with the determination of hospital networks and negotiated rates (Stage 1). When bargaining, both employer and MCO perceive that upon disagreement, the MCO would no longer be available to the employer's members, and members' enrollment and hospital utilization decisions would adjust accordingly.

⁴⁶ An exception is Ho and Pakes (2014), which allows unobserved patient severity to affect preferences for hospitals.

C. Data, Estimation, and the Internal Consistency of Parameter Estimates

In online Appendices B.2–B.3, we summarize the detailed data (covering 2004 enrollment, claims, and hospital admissions for all CalPERS enrollees) and the estimation strategy used in Ho and Lee (2017) to recover structural parameters governing hospital choice, insurance demand, and bargaining over premiums and reimbursement rates. We also summarize the variation in determinants of insurer costs that generate incentives to exclude hospitals. The standard deviation in hospital costs per admission is \$552, compared to a mean of \$1,693. This cross-hospital variation generates steering incentives for the insurer. The level of hospital prices on average \$6,624 per admission for Blue Shield, suggests that hospital margins are large enough to leave room for rate reductions due to exclusion of some providers. Last, we summarize the variation in admission probabilities and treatment intensity across age and gender groups that generates cream-skimming incentives.

Estimated outputs of our model align closely with outside sources. For example, we estimate insurers' non-inpatient hospital costs per enrollee per year in our sample to be \$1,690 for Blue Shield and \$1,950 for BC. By comparison, the Kaiser Family Foundation reports a cross-insurer average of \$1,836 spending per person per year on physician and clinical services for California in 2014, and 2010–2012 data from the three largest insurers in Massachusetts indicate average spending of \$1,644 per person per year on professional services.⁴⁷ Similarly, our recovered own-premium elasticities for each insurer, which range from -1.23 for single-person households for Kaiser to -2.95 for families with children for BC, are well within the range estimated in the previous literature.⁴⁸

As discussed in Section IIIA, while carriers were not precluded from excluding a hospital system in 2004, they were expected to negotiate broad networks. This led to all five of the largest systems being included in the Blue Shield 2004 network for 11 out of the 12 markets that we focus on. Recall that the Nash-in-Nash and NNTR bargaining outcomes are identical, and hence the two models would generate the same parameter estimates, in settings where networks are complete.⁴⁹ Hence, even though bargaining parameters are estimated in Ho and Lee (2017) under Nash-in-Nash bargaining conditional on the observed (almost complete) networks, a natural assumption given the institutional details previously described, we

⁴⁷ Kaiser data accessed from <http://kff.org/other/state-indicator/health-spending-per-capita-by-service/> on February 25, 2015. Massachusetts data obtained from the report "Massachusetts Commercial Medical Care Spending: Findings from the All-Payer Claims Database 2010–12," published by the Center for Health Information and Analysis in partnership with the Health Policy Commission. Both of these figures include member out-of-pocket spending, which is excluded from our estimates; the California data also include the higher-cost Medicare population in addition to the commercially insured enrollees in our sample.

⁴⁸ We report elasticities based on the total premium rather than the out-of-pocket prices faced by enrollees; they are referred to in the previous health insurance literature as "insurer-perspective" elasticities. Ho (2006) uses a similar model (although a different dataset) to generate an estimated elasticity of -1.24 . Cutler and Reber (1998) and Royalty and Solomon (1999) use panel data on enrollee responses to observed plan premium changes in employer-sponsored large group settings to estimate elasticities of -2 , and between -1.02 and -3.5 , respectively.

⁴⁹ The NNTR model can be used for estimation even in settings with incomplete networks, i.e., when its predicted outcomes do not coincide with Nash-in-Nash. One approach would be to adapt the estimation procedure in Ho and Lee (2017) to use moments that match observed negotiated reimbursement rates to those predicted by the NNTR solution (as opposed to the Nash-in-Nash solution).

hold these estimated parameters fixed in our simulations without loss of internal consistency.⁵⁰

V. The Impact of Narrow Networks

We now use our estimated model, paired with our Nash-in-Nash with Threat of Replacement (NNTR) bargaining solution, to simulate market and welfare outcomes as Blue Shield adjusts its hospital network across 12 health service areas (HSAs) in California.⁵¹ In every market, we determine the stable network under NNTR bargaining that maximizes three different objectives—social surplus (which we compute as the sum of insurer and hospital profits and consumer welfare), consumer welfare alone, and Blue Shield's profits—and compute market outcomes and welfare measures under these networks. We refer to these networks as the social, consumer, and insurer-optimal networks. We repeat the same exercise for the network that maximizes Blue Shield's profits under Nash-in-Nash bargaining, and under the scenario where Blue Shield is required to engage in negotiations with all hospital systems (our “full-network regulation” scenario).

Our measure of social surplus does not account for any fixed or sunk costs that enter into firms' profits, as we are unable to estimate or obtain information on these objects for the insurers and hospitals in our analysis. Maximizing social surplus will correspond to maximizing total welfare in our environment if entry and exit decisions as well as fixed or sunk costs (e.g., capital expenditures) are invariant to changes in Blue Shield's network for CalPERS. If this is not the case, then our measure of social surplus will only capture consumer welfare and marginal cost efficiency gains, and not total cost savings.

We make the following assumptions. First, we assume that the other insurers available to enrollees, Blue Cross (BC) and Kaiser Permanente, do not adjust their hospital networks or renegotiate hospital rates (recall that Blue Shield is the only insurer in our setting that engages in selective contracting with hospitals). However, we do allow premiums for all plans to adjust under our model of premium bargaining with the employer. Second, we assume that premiums and hospital reimbursement rates are determined on a market-by-market level, and thus we conduct our simulations separately for each HSA.⁵² Third, we focus only on Blue Shield's negotiations with the five largest systems (by admissions) in each market.⁵³ We hold fixed Blue Shield's contracts with the remaining hospital systems; these smaller systems receive less than half of observed admissions in any given market (and only in three markets do they receive more than 20 percent of admissions). Last, we

⁵⁰That is, in our setting, our parameter estimates are consistent with *either* bargaining solution when networks are essentially complete, and can be used to compare simulated outcomes across Nash-in-Nash and NNTR (when networks are unconstrained). For robustness, we also re-estimated parameters under Nash-in-Nash bargaining using only data from the 11 markets in which the Blue Shield network was complete; estimates were not statistically different.

⁵¹In two HSAs (Northern California and Mid Coast), Kaiser is not available; we exclude these HSAs from our analysis to maintain comparability across markets.

⁵²In reality CalPERS constrains premiums to be set at the state level, and hospital systems may engage in state-wide bargaining.

⁵³In the San Francisco market, we observe only four hospital systems that have more than ten admissions from either Blue Shield or BC in our data. As a result, we only consider the contracting decisions between Blue Shield and the four largest systems for this market.

maintain that MCOs are limited in their ability to steer consumers via cost-sharing instruments, and that hospitals continue to be paid via linear reimbursement rates.

Our simulations accomplish several goals. First, by contrasting the extent to which hospitals are excluded across social, consumer, and insurer-optimal networks, they assess the empirical importance of both private and social incentives to reallocate admissions to relatively inexpensive hospitals; prompt relatively unhealthy enrollees to select into other insurers; and negotiate lower reimbursement rates. Second, by comparing outcomes to what would be realized under a network that includes all major hospital systems, our simulations illuminate potential effects of network regulations and minimum quality standards. The welfare effects of such requirements will depend on the magnitude of the distortion away from the social optimum in the absence of regulations, and any changes in prices, premiums and other outcomes that are associated with the network changes. Finally, by comparing outcomes under the assumption that hospital reimbursement rates are determined via NNTR or by Nash-in-Nash bargaining, our simulations establish the impact of allowing for endogenous outside options on the propensity for and efficacy of exclusion.

We begin by summarizing the results of our simulations, averaged across all the markets that we consider in California. We then present results for two specific HSAs, Sacramento and Santa Barbara/Ventura. This allows us to highlight differences across markets and to depict differential distributional impacts across different zip codes within a market. Next, we discuss the fit of our predictions to Blue Shield's proposal filed with the California DMHC for the year following our data, and conduct two additional robustness tests.

We provide additional implementation details for our simulations in online Appendix B.4.

A. Average Results across Markets

Table 1 reports our main simulation results, averaged across all 12 of our markets. Each column reports outcomes under the network that maximizes social, consumer, or Blue Shield's surplus, or under the complete network where Blue Shield contracts with all of the major hospital systems in each market. For the network that maximizes Blue Shield's surplus, we report results under the assumption that Blue Shield's negotiated hospital prices are determined either by our NNTR bargaining solution, or the Nash-in-Nash solution; all other columns assume that reimbursement rates are determined according to NNTR. We report: insurer and hospitals profits and costs (reported as \$ per capita); Blue Shield's premiums, payments, and costs (\$ per enrollee); Blue Shield's market share; consumer and total welfare (\$ per capita); and the extent to which major hospital systems are excluded from the chosen network. The rightmost column provides the levels of all figures under the complete network, while the remaining columns show average differences compared to this baseline.

We begin by examining the network that maximizes social surplus (first column of results in Table 1). The social-optimal network tends to be quite broad: it is predicted to be complete in half of our markets (6 out of 12), and when exclusion does occur, only one system is predicted to be excluded. Compared to the complete network, Blue Shield's hospital payments are lower by 6 percent, premiums fall by

TABLE 1—SIMULATION RESULTS FOR ALL MARKETS (*Averages*)

Objective	Social (NNTR)	Consumer (NNTR)	Blue Shield		Complete (NNTR/NN)
			(NNTR)	(NN)	
Surplus (\$ per capita)					
BS profits	1.5% [1.1%, 6.9%]	1.4% [0.9%, 8.0%]	2.6% [1.8%, 8.6%]	0.0% [0.0%, 0.0%]	304.7 [287.5, 312.1]
Hospital profits	−6.4% [−24.9%, −4.9%]	−22.9% [−37.7%, −15.0%]	−14.7% [−33.0%, −12.8%]	0.0% [0.0%, 0.0%]	170.0 [159.4, 209.4]
Total hospital costs	0.2% [0.0%, 1.9%]	0.7% [0.0%, 2.5%]	0.5% [0.4%, 2.0%]	0.0% [0.0%, 0.0%]	95.6 [94.1, 96.3]
Total insurance costs	−0.1% [−0.4%, −0.1%]	0.1% [−0.3%, 0.2%]	−0.1% [−0.5%, −0.1%]	0.0% [0.0%, 0.0%]	2,008.5 [1,990.4, 2,025.7]
Transfer/cost (\$ per enrollee)					
BS premiums	−0.6% [−2.7%, −0.5%]	−2.1% [−4.1%, −1.2%]	−1.2% [−3.6%, −1.0%]	0.0% [0.0%, 0.0%]	2,640.1 [2,615.8, 2,695.1]
BS hospital payments	−5.6% [−22.4%, −4.4%]	−19.9% [−34.1%, −12.7%]	−11.9% [−29.6%, −10.1%]	0.0% [0.0%, 0.0%]	369.3 [347.5, 449.3]
BS hospital costs	−0.3% [−0.3%, 0.1%]	0.9% [0.0%, 1.2%]	0.0% [−0.1%, 0.2%]	0.0% [0.0%, 0.0%]	146.2 [146.1, 146.3]
BS market share	0.4% [0.2%, 1.7%]	−1.8% [−2.0%, 0.5%]	0.2% [−0.2%, 1.7%]	0.0% [0.0%, 0.0%]	0.52 [0.51, 0.53]
Welfare Δ (\$ per capita)					
Consumer	11.7 [8.8, 50.3]	27.8 [17.3, 69.2]	19.9 [15.4, 60.9]	0.0 [0.0, 0.0]	
Total	1.0 [0.5, 4.4]	−11.5 [−12.1, −4.2]	−1.1 [−3.4, 2.0]	0.0 [0.0, 0.0]	
Number of complete network markets (out of 12)	6 [1, 7]	1 [0, 2]	4 [0, 4]	12 [12, 12]	
Number of systems excluded	0.5 [0.4, 1.3]	2.3 [1.8, 2.6]	1.2 [1.2, 1.8]	0.0 [0.0, 0.0]	
Number of systems excluded conditional on exclusion	1.0 [1.0, 1.4]	2.5 [2.1, 2.6]	1.8 [1.8, 2.0]	0.0 [0.0, 0.0]	

Notes: Unweighted averages across markets. First four columns report outcomes for the stable network that maximizes social surplus, consumer welfare, or Blue Shield's (BS) profits, under Nash-in-Nash with Threat of Replacement (NNTR) or Nash-in-Nash (NN) bargaining over hospital reimbursement rates. Percentages and welfare calculations represent changes relative to outcomes under the complete network; outcome levels for the complete network (where all five major hospital systems are included) are presented in right-most column. Ninety-five percent confidence intervals, reported below all figures, are constructed by using 80 bootstrap samples of admissions within each hospital-insurer pair to re-estimate hospital-insurer DRG weighted admission prices, re-estimate insurer marginal costs and Nash bargaining parameters, and re-compute simulations (see Ho and Lee 2017 for further details).

approximately \$17 per enrollee, and Blue Shield's market share increases by 0.4 percent. Since Blue Shield has lower estimated marginal costs per-enrollee than its competitors (see online Appendix Table C.2), the market share increase results in a reduction in total insurer costs. We find that these insurer cost reductions are large enough to offset a slight increase in realized hospital costs, resulting in small average social surplus gains of approximately \$1 per capita compared to the complete network. Furthermore, consumers benefit when hospitals are excluded: consumer surplus increases by \$12 per capita as a result of lower premiums.⁵⁴

We next contrast outcomes under the social-optimal network to those from Blue Shield's optimal stable network under NNTR bargaining (third column of Table 1).

⁵⁴ Consumer welfare gains are less than realized premium reductions because the narrower hospital network offered by Blue Shield's plan represents a reduction in plan quality.

The Blue Shield optimal network is weakly narrower than the social optimum in all but one HSA, and is complete in only four markets.⁵⁵ Furthermore, when exclusion occurs, more hospital systems are dropped under BS's optimal network (1.8 on average).

Why does Blue Shield generally prefer to exclude hospitals from its network? In Section I, we discussed why an insurer might wish to engage in exclusion. The simulations suggest that cream-skimming and steering incentives are not primarily driving Blue Shield's decision to exclude. Blue Shield's average realized hospital costs per enrollee (i.e., marginal costs for hospitals treating its patients, not payments) are essentially unchanged when it moves from the complete network outcome (last column of Table 1) to its optimal network under NNTR bargaining. Its market share (and predicted admission probabilities and DRG weights conditional on admission, not reported) are also not significantly different. On the other hand, Blue Shield's average hospital payments are significantly lower (by 11.9 percent, or \$44 per enrollee per year) under its optimal network relative to the complete network. This strongly suggests that rate-setting incentives, the focus of our NNTR bargaining model, are the primary motive behind exclusion.⁵⁶ Indeed, in markets where exclusion actually occurs, average rate reductions are larger: 18 percent across markets, and up to 30 percent in some markets.

Why is the network that maximizes Blue Shield's profits different from the one that maximizes social surplus? Recall that while Blue Shield's profits depend on the hospital rates that it negotiates, these rates are transfers that a social planner does not incorporate into its welfare calculation. On the other hand, social surplus accounts for inframarginal consumer surplus and hospitals' and other insurers' realized costs which do not affect Blue Shield's profits. In our setting, these differences combine to make the social-optimal network generally broader than that preferred by Blue Shield.

Note finally that consumer surplus is on average higher under the insurer-optimal network than under the social-optimal network (the magnitude of the difference is about \$8 per capita per year). Much of this consumer welfare gain arises from a reduction in insurer premiums which are sufficient to compensate consumers for any potential loss in network utility.⁵⁷ These insights also explain why the network maximizing consumer welfare (column 2 of Table 1) is the most narrow in all markets: it is complete in only a single market, and on average excludes 2.3 of the largest hospital systems in each market. Under the consumer-optimal network, Blue Shield negotiates even lower hospital payments on average than under the insurer-optimal network; these savings are passed along to consumers in the form of even lower

⁵⁵ The insurer-optimal network under NNTR bargaining is broader than the social-optimal network only in Santa Barbara/Ventura; see Section VB.

⁵⁶ We also compute Blue Shield's realized payments to hospitals under the optimal network reported in column 3 if its hospital rates were *fixed* to be the same as rates negotiated under the complete network. This exercise examines whether Blue Shield might still benefit from exclusion by steering patients to lower-*priced* hospitals (as opposed to lower-*cost* hospitals). We find that hospital payments under fixed rates actually increase slightly (by 0.7 percent) on average from the complete network to the narrower network chosen when Blue Shield maximizes its profits under NNTR bargaining. Again, this suggests that bargaining incentives, rather than steering and cream-skimming, provide the dominant motivation for exclusion in our setting.

⁵⁷ Although we report only reductions in Blue Shield's premiums, the premiums for other insurers also fall when Blue Shield's premiums fall.

premiums. However, these premium reductions are not accounted for by a social planner (as they are primarily transfers from firms to consumers), and we find that total welfare falls significantly in the consumer-optimal network relative to the social optimum.

Full-Network Regulation.—Now assume that outcomes under regulations restricting the extent to which insurers can exclude hospitals (e.g., full-network regulation) can be approximated by the rightmost column of Table 1. This reports outcomes when Blue Shield negotiates over a complete network under NNTR (or Nash-in-Nash) bargaining. Under this interpretation, our results suggest that such regulations would significantly increase Blue Shield's payments to hospitals and hospital profits while also generating premium increases and corresponding reductions in consumer welfare. Focusing on the insurer-optimal network, we also find that social surplus would not change significantly on average across our markets from such a regulation.

Thus in our setting, regulations that restrict an insurer's ability to exclude may eliminate an important source of bargaining leverage that insurers can use to reduce hospital payments and benefit consumers through lower premiums. Hospitals and insurers may disagree on the need for such regulation as it directly affects the transfers that are made between the two parties.

NNTR versus Nash-in-Nash.—Last, we contrast predicted outcomes when Blue Shield is able to choose its preferred network, and hospital prices are either negotiated according to our NNTR bargaining solution or under Nash-in-Nash bargaining.

As discussed in Section III, under Nash-in-Nash bargaining the insurer has an incentive to include additional hospitals in order to reduce the marginal contribution of any given hospital and hence its negotiated payments. NNTR, on the other hand, allows an insurer to "play off" included and excluded hospitals when bargaining, and thus does not require that a hospital be included in a network in order to be employed as bargaining leverage.

In our setting, we find that the Blue Shield profit maximizing network under Nash-in-Nash bargaining (column 4 in Table 1) *never* excludes any hospitals. This stands in stark contrast to both Blue Shield's actual proposal to exclude 38 hospitals in 2005, and our predictions for Blue Shield's optimal network under NNTR bargaining. As Nash-in-Nash cannot rationalize observed exclusion on the part of Blue Shield, we report results under NNTR bargaining for the remainder of our results.

B. Simulation Results: Specific Markets

We now provide results for two specific markets: Sacramento, which is located in the northern part of California and includes the state capital; and Santa Barbara/Ventura, located approximately 100 miles from Los Angeles in southern California.

Sacramento.—Our results for Sacramento are reported in Table 2. Here the difference between the social-optimal and insurer-optimal network is stark: the social-optimal network coincides with the complete network, whereas the

TABLE 2—SIMULATION RESULTS FOR SACRAMENTO

Objective	Social	Consumer	Blue Shield	Complete
Surplus (per capita)				
BS profits	0.0% [0.0%, 10.3%]	3.1% [1.7%, 10.3%]	3.1% [1.7%, 10.3%]	316.2 [290.2, 325.9]
Hospital profits	0.0% [−40.1%, 0.0%]	−26.0% [−40.1%, −21.3%]	−26.0% [−40.1%, −21.3%]	115.5 [102.2, 170.7]
Total hospital costs	0.0% [0.0%, 3.6%]	1.6% [1.2%, 3.6%]	1.6% [1.2%, 3.6%]	98.5 [96.1, 99.4]
Total insurance costs	0.0% [−0.6%, 0.0%]	−0.1% [−0.6%, 0.0%]	−0.1% [−0.6%, 0.0%]	2,049.8 [2,032.6, 2,068.5]
Transfers (per enrollee)				
BS premiums	0.0% [−3.5%, 0.0%]	−1.5% [−3.5%, −1.1%]	−1.5% [−3.5%, −1.1%]	2,619.7 [2,593.9, 2,688.7]
BS hospital payents	0.0% [−30.4%, 0.0%]	−16.8% [−30.4%, −12.9%]	−16.8% [−30.4%, −12.9%]	333.8 [307.4, 444.8]
BS hospital costs	0.0% [0.0%, 1.2%]	1.2% [1.1%, 1.3%]	1.2% [1.1%, 1.3%]	165.5 [165.4, 165.7]
Δ Welfare (per capita)				
Consumer	0.0 [0.0, 60.1]	23.3 [15.7, 60.1]	23.3 [15.7, 60.1]	
Total	0.0 [0.0, 5.0]	−3.4 [−5.0, 5.0]	−3.4 [−5.0, 5.0]	
BS market share	0.0% [0.0%, 2.6%]	0.2% [−0.2%, 2.6%]	0.2% [−0.2%, 2.6%]	0.53 [0.52, 0.54]
Network				
Number of systems excluded	0 [0, 3]	3 [3, 3]	3 [3, 3]	
System 1 (Sutter)	1 [1.0]	1 [1.0]	1 [1.0]	
System 2 (Dignity)	1 [1.0]	1 [1.0]	1 [1.0]	
System 3 (UCD)	1 [0.9]	0 [0.0]	0 [0.0]	
System 4 (Rideout)	1 [0.9]	0 [0.0]	0 [0.0]	
System 5 (Marshall)	1 [0.9]	0 [0.0]	0 [0.0]	

Notes: Simulation results from Sacramento HSA. First three columns report outcomes for the stable network that maximizes social surplus, consumer welfare, or Blue Shield's profits, under Nash-in-Nash with Threat of Replacement (NNTR) bargaining over hospital reimbursement rates. Percentages and welfare calculations represent changes relative to outcomes under the complete network; outcome levels for the complete network (where all five major hospital systems are included) are presented in right-most column. Ninety-five percent confidence intervals are reported below all figures (except for individual hospital systems, where the fraction of bootstrap samples under which individual system members are included are reported beneath predictions); see Table 1 for additional details.

insurer-optimal (which is also the consumer-optimal) network includes only the two largest systems in the market. Restricting Blue Shield's network in this fashion allows it to reduce its payments to hospitals by a statistically significant 17 percent (\$57 per enrollee per year); this results in 1.5 percent (\$39 per enrollee) lower premiums.

To understand why Blue Shield chooses to exclude three out of the five major hospital systems (as opposed to just one) in its optimal network under NNTR bargaining, it is important to emphasize that hospitals are differentiated on multiple horizontal and vertical dimensions; thus, some hospitals may serve as more effective threats to replace particular hospitals remaining in the network than others. To illustrate, a map of the Sacramento market is provided in panel A of Figure 2. Hospital systems that are included in both the insurer-optimal NNTR network and the social-optimal network are marked in black, while those excluded by Blue Shield (but included by the social planner) are white. Note that two hospital systems that are excluded by Blue Shield (Rideout and Marshall) both serve more rural zip codes on the outskirts of the Sacramento HSA than the others. The third excluded hospital system, UC Davis, is located closer to the Sacramento city center. From geography alone, one might surmise that while Rideout and Marshall could be effective as threats to replace one another, they might not be as effective threats as UC Davis to the Sutter and Dignity hospital systems.⁵⁸

Indeed, we find that this intuition carries through in our analysis. Figure 3 reports predicted average hospital system admission prices for Blue Shield under NNTR bargaining under the insurer-optimal stable network that excludes exactly zero, one, two, or three major hospital systems. If Blue Shield were restricted to excluding only a single hospital system, it would first choose to exclude Marshall; doing so would reduce its payments to Rideout by 18 percent, but have little effect on those paid to other systems. Restricting attention to stable networks excluding two hospital systems, Blue Shield would exclude both Rideout and Marshall; doing so would have a small (less than 4 percent) effect on other hospital prices. Finally, among all potential stable networks excluding additional hospital systems, Blue Shield would choose to exclude three hospital systems (Rideout, Marshall, and UC Davis); upon excluding these systems, the prices for Sutter and Dignity fall by approximately 17 percent each. We find that as the number of systems being excluded increases from zero to two, Blue Shield's profits are relatively unchanged; however, its profits increase by 3 percent once the third system is excluded as payments to hospitals fall substantially (as reported in Table 2).⁵⁹

Figure 2 also illustrates the distributional consequences of selective contracting and narrow hospital networks. Focusing again on panel A (Sacramento), zip codes in the market are shaded to reflect the average consumer surplus difference between the insurer- and social-optimal network. For zip codes close to hospitals that are still included in the insurer-optimal network (such as those close to the center of Sacramento), consumer surplus increases significantly (denoted by dotted shading) when Blue Shield engages in exclusion as consumers benefit from reduced

⁵⁸ Even though Rideout and Marshall, located in geographically distinct parts of the market, may not be close substitutes to one another from the perspective of a particular individual, they are close substitutes from the perspective of the insurer as they have similar effects on the insurer's enrollment and revenues from inclusion. The distinction between substitutability from the perspective of a consumer versus an insurer also matters when evaluating the price effects of cross-market hospital mergers (cf. Dafny, Ho, and Lee forthcoming).

⁵⁹ Networks that exclude only UC Davis or UC Davis and either Rideout or Marshall are not stable: greater gains from trade would be generated if Blue Shield replaced either Marshall or Rideout with UC Davis (see Proposition 2). Indeed, if UC Davis were excluded, Marshall and Rideout would not wish to accept NNTR prices. Thus, as discussed in Section IV, it would not be credible in our bargaining framework for Blue Shield to exclude UC Davis from its network but not Marshall and Rideout.

Panel A. Sacramento



Panel B. Santa Barbara/Ventura

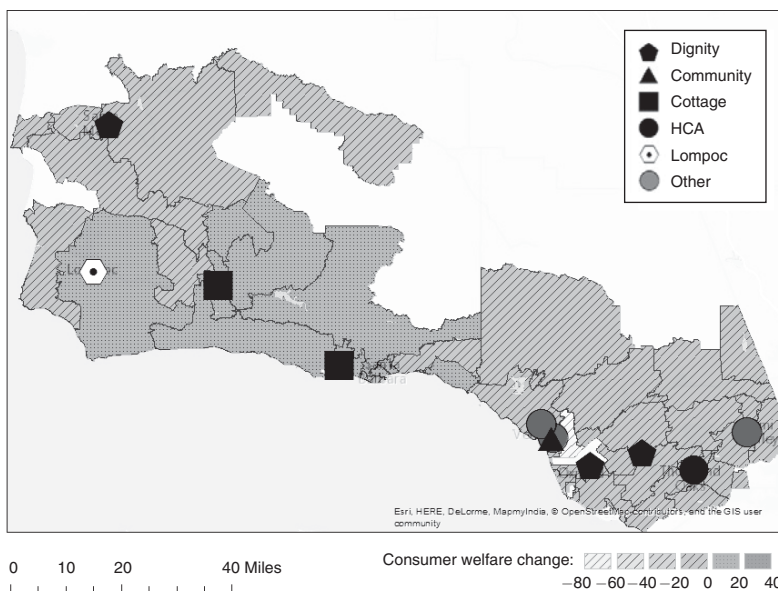


FIGURE 2. CONSUMER WELFARE CHANGES BETWEEN BLUE SHIELD AND SOCIAL-OPTIMAL NETWORKS

Notes: Black shapes indicate hospital systems that are included in both the insurer- and social-optimal networks; white are excluded from the insurer-optimal network but included in the social-optimal network; and white with a black dot indicates a hospital included in the insurer-optimal but excluded from the social-optimal network. Gray shapes indicate other hospitals (Kaiser hospitals and non-Kaiser providers other than the top five systems). Consumer welfare changes (\$ per capita) in each zip code are indicated by shading: darker colors indicating larger changes, with dots (lines) indicating positive (negative) consumer welfare changes between the insurer- and social-optimal network. White areas are zip codes that are either outside of the relevant market or have no residents in our data.




FIGURE 3. SACRAMENTO HOSPITAL SYSTEM PRICES UNDER ALTERNATIVE NETWORKS AND NNTR BARGAINING

Notes: Average severity adjusted hospital system admission prices predicted under NNTR bargaining for the Blue Shield insurer-optimal stable network that excludes exactly zero, one, two, or three major hospital systems. A missing price indicates that the associated system is excluded.

premiums. However, consumers in zip codes close to excluded hospitals experience welfare reductions. There are zip codes experiencing welfare reductions as high as \$70 per capita per year, and others with increases of up to \$40 per capita per year. Thus, although full-network regulation in this industry would actually reduce consumer surplus by \$23 per capita per year on average, there are consumers for whom such regulation would be highly beneficial because hospitals that are close to them would no longer be excluded from the Blue Shield network.

Finally, as discussed earlier, if we instead assumed that hospital rates were determined via Nash-in-Nash bargaining, we would predict that Blue Shield would not wish to exclude any hospital system in this market. This is inconsistent with the observation that a major hospital system in Sacramento was in fact dropped by Blue Shield in 2005 after a lengthy process of negotiation with the Department of Managed Health Care. Our NNTR model, in contrast, is able to rationalize Blue Shield's desire to engage in exclusion.⁶⁰

Santa Barbara/Ventura.—Table 3 presents results for the Santa Barbara/Ventura HSA. In this market (and only this market), the social-optimal network is actually narrower than the one that would be chosen by Blue Shield. Though we employ the NNTR solution, the Blue Shield optimal network is predicted to be complete (thus highlighting that NNTR need not imply exclusion). The social-optimal network excludes a single system primarily for steering reasons: by excluding the smallest but highest-cost of the five largest hospital systems, the social planner reduces hospital cost expenditures per capita by 1.0 percent relative to the complete network. However, since such an adjustment reduces Blue Shield's profits, the insurer does

⁶⁰The NNTR model predicts that Blue Shield should exclude three hospital systems, while in reality only one was dropped. However, the predicted number of excluded *hospitals* is correct. See Section VC.

TABLE 3—SIMULATION RESULTS FOR SANTA BARBARA/VENTURA

Objective	Social	Consumer	Blue Shield	Complete
Surplus (per capita)				
BS profits	−0.3% [−0.3%, 0.1%]	−5.0% [−5.2%, −0.3%]	0.0% [0.0%, 0.1%]	397.7 [382.9, 403.3]
Hospital profits	0.0% [−1.5%, 0.4%]	−1.5% [−15.3%, 0.4%]	0.0% [−1.5%, 0.0%]	240.4 [224.0, 299.9]
Total hospital costs	−1.0% [−1.0%, −0.9%]	−3.5% [−3.6%, −1.0%]	0.0% [−0.9%, 0.0%]	115.8 [115.1, 116.1]
Total insurance costs	0.0% [0.0%, 0.0%]	0.5% [0.0%, 0.6%]	0.0% [0.0%, 0.0%]	1,832.9 [1,815.1, 1,849.7]
Transfers (per enrollee)				
BS premiums	−0.1% [−0.3%, 0.0%]	−0.5% [−2.5%, 0.0%]	0.0% [−0.3%, 0.0%]	2,677.8 [2,646.6, 2,751.6]
BS hospital payments	−0.5% [−2.0%, −0.2%]	−3.1% [−17.0%, −0.2%]	0.0% [−2.0%, 0.0%]	363.9 [338.0, 459.2]
BS hospital costs	−1.4% [−1.4%, −1.4%]	−4.6% [−4.6%, −1.4%]	0.0% [−1.4%, 0.0%]	126.0 [126.0, 126.1]
Δ Welfare (per capita)				
Consumer	1.6 [0.7, 7.0]	7.0 [0.7, 55.7]	0.0 [0.0, 7.0]	
Total	0.5 [0.4, 0.8]	−15.2 [−15.7, 0.5]	0.0 [0.0, 0.8]	
BS market share	−0.2% [−0.2%, −0.1%]	−4.6% [−4.7%, −0.2%]	0.0% [−0.1%, 0.0%]	0.64 [0.63, 0.64]
Network				
Number of systems excluded	1 [1, 1]	3 [1, 3]	0 [0, 1]	
System 1 (Dignity)	1 [1.0]	1 [1.0]	1 [1.0]	
System 2 (Community)	1 [1.0]	1 [1.0]	1 [1.0]	
System 3 (Cottage)	1 [1.0]	0 [0.2]	1 [1.0]	
System 4 (HCA)	1 [1.0]	0 [0.2]	1 [1.0]	
System 5 (Lompoc MC)	0 [0.0]	0 [0.0]	1 [0.9]	

Notes: Simulation results from Santa Barbara/Ventura HSA. See notes from Table 3.

not choose to engage in exclusion. Again, as in Sacramento, we predict a very narrow consumer-optimal network.

Panel B of Figure 2 repeats our exercise depicting the distributional effects of network adjustments using a map of Santa Barbara/Ventura. Systems included under both the insurer- and social-optimal networks are again depicted in black. The system that is excluded by the social planner but included by Blue Shield is in

white with a black dot. By including this system in its optimal network, Blue Shield generates a consumer welfare increase within the same zip code (of \$12 per capita) relative to the social-optimal network. However, there is a slight consumer welfare reduction in other areas due to the resulting premium increase.

C. Fit of the Model

Our information on the 38 hospital systems that Blue Shield proposed to drop in 2005 provides an opportunity to investigate the model's ability to predict out-of-sample outcomes, with the caveat that we should not expect a perfect fit to these data. While it is reasonable to assume that rates are determined under Nash-in-Nash bargaining in 2004 given the institutional details described above, it is difficult to simulate the true data-generating process for Blue Shield's observed 2005 proposal. There are several reasons for this. First, we conduct our simulations at the market (HSA) level, and assume that there are no cross-market linkages when premiums and hospital rates are determined. In reality, premiums are constrained to be constant across the entire state, and decisions to drop hospital systems that are active in multiple markets are likely to be made at the state level. Second, CalPERS and the California DMHC continued to impose external constraints on networks in 2005: Blue Shield proposed to drop 38 hospitals but was permitted to exclude only 24. Such constraints, and the extent to which Blue Shield's original proposal was designed to address them, are not factored into our analysis.

We also find evidence that political considerations may have constrained Blue Shield's proposal. In Sacramento, for example, the NNTR model predicts that Blue Shield should optimally exclude the University of California, Davis medical center. Similarly, in Los Angeles, the University of California, Los Angeles medical center is also predicted to be excluded. Given the close links between CalPERS enrollees (i.e., California state and local employees) and the University of California system, these exclusions might well have been politically difficult to implement. In its actual proposal, Blue Shield did not propose to drop either of the UC medical centers. Instead, in Sacramento it proposed (and was allowed) to drop Sutter hospitals, and in Los Angeles it excluded the hospital belonging to University of Southern California (the private school "rival" of UCLA) among others.

Given these unmodeled factors, we do not use a direct comparison of our predicted networks to those proposed by Blue Shield for 2005 in order to test the empirical relevance of our model.⁶¹ Rather, we focus our attention on the model's ability to predict features of the proposed 2005 network such as the number of systems excluded, the implied number of excluded hospitals, and their characteristics.⁶²

The Blue Shield profit-maximizing network under NNTR bargaining perfectly coincides with the number of excluded systems proposed by Blue Shield in 3 of

⁶¹ We note, however, that the unmodeled complexities that factored into Blue Shield's 2005 proposal do not introduce bias into our estimates since we do not use the 2005 data for estimation. Nor do they prevent us from demonstrating that the NNTR model predicts exclusion when the Nash-in-Nash model cannot, or investigating the consequences of narrow networks in our setting when such complexities are absent.

⁶² Online Appendix Table C.1 lists the hospitals that Blue Shield proposed to exclude in 2005. It covers the same 12 markets that are the focus of our simulations. However it differs from the simulations in that it lists all proposed exclusions, including those for providers outside the five largest systems per market.

the 12 HSAs we consider. In the remaining nine markets, it predicts a broader network than that proposed in five HSAs, and a narrower network in four HSAs. Thus, while the fit is imperfect, the errors do not seem to indicate a systematic bias in our model. The Nash-in-Nash model, in contrast, consistently predicts a broader network than that proposed by Blue Shield. The predicted number of excluded systems is correct in just one market, where in reality Blue Shield did not propose to drop any of the top five systems. Nash-in-Nash predicts a broader network than that proposed in all 11 remaining HSAs. Clearly, based on these simple statistics alone, the NNTR model comes closer to predicting the observed Blue Shield proposal than does Nash-in-Nash.

We further investigate the fit of the NNTR model by considering particular system characteristics. The number of hospitals per system is informative in some markets. For example, in Sacramento, Blue Shield proposed to exclude the Sutter hospital system. This is the largest system in the market, with four main hospitals. The NNTR model predicts instead that three smaller systems should optimally be excluded. The discussion above considers this to be a narrower network than that proposed. However, in reality, those three smaller systems also contain four hospitals in total. Thus, while the model cannot match the predicted identity of the excluded system (which was likely determined in part by cross-market considerations), it does correctly predict the number of hospitals excluded in total.

Our simulations also fit the data quite well with respect to hospital costs and other characteristics. For example, hospitals that Blue Shield proposed to exclude in 2005 had higher average DRG-adjusted costs per admission than included hospitals (\$1,973 compared to \$1,840), which are close to the differences predicted by our simulations (\$2,038 compared to \$1,794). Teaching hospitals were also more likely to be excluded, both in the observed 2005 data (18 percent of excluded hospitals compared to 15 percent of those included) and in the simulations (20 percent compared to 17 percent).

Finally, we also compare the model's predicted social-surplus-maximizing network to the 24 exclusions that were *approved* by the DMHC (as the regulator may have acted to maximize social surplus). We find a reasonable fit in terms of the number of excluded systems: our prediction matches the number of approved exclusions in 6 out of 12 markets. In 5 of the remaining 6, the model predicts a socially optimal network that is broader than that approved by the DMHC. This is consistent with the DMHC's standard for approval being based solely on consumers' distance to in-network hospitals, without concern for other issues such as hospital quality, the match of particular hospitals with surrounding populations' needs, and the underlying costs of insurers. As these issues are accounted for in our social surplus measure, we might expect to observe more approved exclusion in the data than our model predicts would be socially optimal.

D. Robustness

No Blue Cross.—Our main specification assumes that Blue Shield adjusts its network while the networks of its competitors remain fixed. One such competitor is BC, which offers a broad-network PPO plan whose presence may mitigate the consumer harm from a narrow Blue Shield network. We investigate the importance of

this institutional reality by repeating the simulations in the scenario where the BC PPO plan is not included in the CalPERS choice set. Online Appendix Table C.4 reports average results across markets from this specification.

Unsurprisingly, when BC is not available, the consumer-optimal Blue Shield network is broader on average than it was in the case where the BC plan was available. Consumer welfare is maximized when the network excludes 1.7 systems on average (compared to 2.3 in the baseline case). That is, consumer welfare is more negatively affected by Blue Shield offering a narrow network, low premium plan when no comprehensive outside option is offered. Consistent with this, the social-optimal network is slightly more likely to be complete. However, even in this setting, we still find that consumers prefer Blue Shield to engage in exclusion (to negotiate lower hospital rates and pass along savings in the form of lower premiums) in the vast majority of markets. Full-network regulation would again reduce consumer welfare through higher premiums that more than outweigh the benefit of broader provider choice.

Fixed Premiums.—In online Appendix Table C.5, we report average results across markets under the assumption that premiums are fixed at the levels determined when Blue Shield's hospital network is complete. In this case, we predict that the social-optimal network is complete in all but one market. The Blue Shield optimal network under NNTR bargaining involves exclusion in every market. Blue Shield's average hospital rate reductions from exclusion are fairly similar to those when premiums are allowed to adjust (18.8 percent lower compared to 11.9 percent on average), but since it no longer passes savings to consumers in the form of lower premiums, its profits increase by much more than before (8.8 percent as opposed to 2.6 percent).

The consumer optimal network when premiums are fixed no longer tends to be narrow; rather, it is predicted to always be complete. This is unsurprising. When premiums are fixed, consumers can only be harmed when hospitals are excluded in our model, as they cannot be compensated for what is essentially a quality reduction. Clearly, any evaluation of the consumer welfare effects from provider network regulation in health care settings needs to account for resulting premium adjustments by insurers.

E. Caveats and Directions for Future Research

This paper uses estimated primitives governing consumer preferences, firm costs, and Nash bargaining parameters from Ho and Lee (2017) to simulate the equilibrium outcomes that would emerge if one of the insurers in our setting (Blue Shield) could freely adjust its network. Our demand system abstracts away from additional sources of unobservable heterogeneity, risk aversion over provider access, and switching costs and other forms of inertia. Furthermore, we do not account for changes in provider fixed costs, or entry and exit decisions, in response to counterfactual network adjustments. Our welfare results should be interpreted with this in mind.

One promising alternative avenue of research would be to examine environments where prices and premiums could be observed over time as networks narrowed.

This would provide more direct evidence on the impact of exclusion on negotiated rates. In addition, when paired with alternative bargaining solutions such as NNTR, these environments may allow for the estimation of additional supply side primitives such as fixed costs of contracting, or the recovery of additional features of consumer demand that would enrich future work. More challenging extensions include adapting the ideas discussed in Section III E to allow multiple insurers (or competing intermediaries, more generally) to adjust their hospital (or supplier) networks. Moving in this direction would push current theoretical exercises informed by empirical primitives into more full-blown applied analyses.

VI. Conclusion

Narrow provider networks have grown more prevalent in both exchange and employer-sponsored health care markets in recent years. Their presence raises important questions. Are these plans effective at reducing spending, and if so, through what means? And is regulation warranted: are the networks that are introduced too narrow from either a social or consumer welfare perspective?

Our paper addresses these and related questions. We extend the model of the commercial US health care market developed in Ho and Lee (2017) by endogenizing an insurer's hospital network and incorporating a new bargaining solution that explicitly captures an insurer's incentives to exclude. In the employer-sponsored setting that we examine, we find that selective contracting and informed network design can have substantial effects on overall health care spending. Narrow hospital networks are preferred by a profit-maximizing insurer primarily due to their ability to substantially reduce negotiated rates, and not necessarily due to cream-skimming healthier enrollees or steering patients toward lower-cost hospitals. A private insurer tends to engage in exclusion more than is socially optimal, but typically does so to a lesser extent than the average consumer would prefer because consumers often benefit from substantial premium reductions.

These results support the argument that allowing insurers to exclude providers can substantially reduce hospital payments and premiums without significantly affecting social surplus. This tends to benefit consumers (and hence employers) on average, implying that employers and insurers may wish to work together to control spending through exclusion. Our framework may be useful for these and other interested parties to inform network design. It also can be used to address potential distributional consequences of exclusion by identifying affected populations and quantifying the transfers needed to offset harm from reduced access.

REFERENCES

- Arie, Guy, Paul L. E. Grieco, and Shiran Rachmilevitch. 2017. "Generalized Insurer Bargaining." Unpublished.
- Baumgarten, Allan. 2005. *California Health Care Market Report 2005*. Oakland, CA: The California HealthCare Foundation.
- Binmore, Ken. 1987. "Perfect Equilibria in Bargaining Models." In *The Economics of Bargaining*, edited by Ken Binmore and Patha Dasgupta. Oxford: Blackwell Publishers.
- Binmore, Ken, Ariel Rubinstein, and Asher Wolinsky. 1986. "The Nash Bargaining Solution in Economic Modelling." *RAND Journal of Economics* 17 (2): 176–88.

- Binmore, Ken, Avner Shaked, and John Sutton. 1989. "An Outside Option Experiment." *Quarterly Journal of Economics* 104 (4): 753–70.
- Bolton, Patrick, and Michael D. Whinston. 1993. "Incomplete Contracts, Vertical Integration, and Supply Assurance." *Review of Economic Studies* 60 (1): 121–48.
- Capps, Cory, David Dranove, and Mark Satterthwaite. 2003. "Competition and Market Power in Option Demand Markets." *RAND Journal of Economics* 34 (4): 737–63.
- Chipty, Tasneem. 2001. "Vertical Integration, Market Foreclosure, and Consumer Welfare in the Cable Television Industry." *American Economic Review* 91 (3): 428–53.
- Claxton, G., I. Gil, B. Finder, J. Gabel, J. Pickreign, H. Whitmore, and S. Hawkins. 2005. *Employer Health Benefits Annual Survey 2005*. Oakland, CA: Kaiser Family Foundation and the Health Research and Educational Trust.
- Collard-Wexler, Allan, Gautam Gowrisankaran, and Robin S. Lee. Forthcoming. "'Nash-in-Nash' Bargaining: A Microfoundation of Applied Work." *Journal of Political Economy*.
- Cooper, Zack, Stuart V. Craig, Martin Gaynor, and John Van Reenen. 2018. "The Price Ain't Right? Hospital Prices and Health Spending on the Privately Insured." NBER Working Paper 21815.
- Corlette, Sabrina, JoAnn Volk, Robert Berenson, and Judy Feder. 2014. *Narrow Provider Networks in New Health Plans: Balancing Affordability with Access to Quality Care*. Washington, DC: Georgetown University on Health Insurance Reforms and Urban Institute.
- Crawford, Gregory S., Robin S. Lee, Michael D. Whinston, and Ali Yurukoglu. 2018. "The Welfare Effects of Vertical Integration in Multichannel Television Markets." *Econometrica* 86 (3): 891–954.
- Crawford, Gregory S., and Ali Yurukoglu. 2012. "The Welfare Effects of Bundling in Multichannel Television Markets." *American Economic Review* 102 (2): 643–85.
- Cutler, David M., Mark McClellan, and Joseph P. Newhouse. 2000. "How Does Managed Care Do It?" *RAND Journal of Economics* 31 (3): 526–48.
- Cutler, David M., and Sarah J. Reber. 1998. "Paying for Health Insurance: The Trade-Off between Competition and Adverse Selection." *Quarterly Journal of Economics* 113 (2): 433–66.
- Dafny, Leemore S., Igal Hendel, Victoria Marone, and Christopher Ody. 2017. "Narrow Networks on the Health Insurance Marketplaces: Prevalence, Pricing, and the Cost Network Breadth." *Health Affairs* 36 (9): 1606–14.
- Dafny, Leemore, Igal Hendel, and Nathan Wilson. 2015. "Narrow Networks on the Health Insurance Exchanges: What Do They Look Like and How Do They Affect Pricing? A Case Study of Texas." *American Economic Review* 105 (5): 110–14.
- Dafny, Leemore S., Kate Ho, and Robin S. Lee. Forthcoming. "The Price Effects of Cross-Market Hospital Mergers." *RAND Journal of Economics*.
- Draganska, Michaela, Daniel Klapper, Sofia B. Villas-Boas. 2010. "A Larger Slice or a Larger Pie? An Empirical Investigation of Bargaining Power in the Distribution Channel." *Marketing Science* 29 (1): 57–74.
- Gal-Or, Esther. 1997. "Exclusionary Equilibria in Health-Care Markets." *Journal of Economics and Management Strategy* 6 (1): 5–43.
- Ghili, Soheil. 2018. "Network Formation and Bargaining in Vertical Markets: The Case of Narrow Networks in Health Insurance." Unpublished.
- Giovannelli, Justin, Kevin Lucia, and Sabrina Corlette. 2016. "Regulation of Health Plan Provider Networks." *Health Affairs Health Policy Brief*, July 28, 2016.
- Glied, Sherry. 2003. "Health Care Costs: On the Rise Again." *Journal of Economic Perspectives* 17 (2): 125–48.
- Gowrisankaran, Gautam, Aviv Nevo, and Robert Town. 2015. "Mergers When Prices Are Negotiated: Evidence from the Hospital Industry." *American Economic Review* 105 (1): 172–203.
- Grennan, Matthew. 2013. "Price Discrimination and Bargaining: Empirical Evidence from Medical Devices." *American Economic Review* 103 (1): 145–77.
- Gruber, Jonathan, and Robin McKnight. 2016. "Controlling Health Care Costs through Limited Network Insurance Plans: Evidence from Massachusetts State Employees." *American Economic Journal: Economic Policy* 8 (2): 219–50.
- Handel, Ben, Igal Hendel, and Michael D. Whinston. 2015. "Equilibria in Health Exchanges: Adverse Selection versus Reclassification Risk." *Econometrica* 83 (4): 1261–313.
- Ho, Kate. 2006. "The Welfare Effects of Restricted Hospital Choice in the US Medical Care Market." *Journal of Applied Econometrics* 21 (7): 1039–79.
- Ho, Kate, and Robin S. Lee. 2017. "Insurer Competition in Health Care Markets." *Econometrica* 85 (2): 379–417.
- Ho, Kate, and Robin S. Lee. 2019. "Equilibrium Provider Networks: Bargaining and Exclusion in Health Care Markets: Dataset." *American Economic Review*. <https://doi.org/10.1257/aer.20171288>.

- Ho, Kate, and Robin S. Lee.** Forthcoming. "Narrow Medical Provider Networks: Welfare Implications and Approaches to Market Design." In *More Fair By Design: Design Responses to Inequality*, Vol. IV, edited by Scott D. Kominers and Alex Teytelboym. Oxford: Oxford University Press.
- Ho, Kate, and Ariel Pakes.** 2014. "Hospital Choices, Hospital Prices, and Financial Incentives to Physicians." *American Economic Review* 104 (12): 3841–84.
- Horn, Henrick, and Asher Wolinsky.** 1988. "Bilateral Monopolies and Incentives for Merger." *RAND Journal of Economics* 19 (3): 408–19.
- Jackson, Matthew O., and Asher Wolinsky.** 1996. "A Strategic Model of Social and Economic Networks." *Journal of Economic Theory* 71 (1): 44–74.
- Lee, Robin S., and Kyna Fong.** 2013. "Markov-Perfect Network Formation: An Applied Framework for Bilateral Oligopoly and Bargaining in Buyer-Seller Networks." Unpublished.
- Leland, Hayne E.** 1979. "Quacks, Lemons, and Licensing: A Theory of Minimum Quality Standards." *Journal of Political Economy* 87 (6): 1328–46.
- Lieberman, Eli.** 2018. "Bargaining in Markets with Exclusion: An Analysis on Health Insurance Networks." Unpublished.
- Manea, Mihai.** 2018. "Intermediation and Resale in Networks." *Journal of Political Economy* 126 (3): 1250–1301.
- McKinsey and Company.** 2013. *Hospital Networks: Configurations on the Exchanges and Their Impact on Premiums*. Washington, DC: McKinsey Center for US Health System Reform.
- Melnick, Glenn A., Jack Zwanzinger, Anil Bamezai, and Robert Pattison.** 1992. "The Effects of Market Structure and Bargaining Position on Hospital Prices." *Journal of Health Economics* 11 (3): 217–33.
- Muthoo, Abhinay.** 1999. *Bargaining Theory with Applications*. Cambridge, UK: Cambridge University Press.
- Osborne, Martin J., and Ariel Rubinstein.** 1990. *Bargaining and Markets*. San Diego: Academic Press.
- Polisky, Dan, and Janet Weiner.** 2015. *The Skinny on Narrow Networks in Health Insurance Marketplace Plans*. Philadelphia, PA: Leonard Davis Institute of Health Economics.
- Prager, Elena.** 2018. "Consumer Responsiveness to Simple Health Care Prices: Evidence from Tiered Hospital Networks." Unpublished.
- Ronnen, Uri.** 1991. "Minimum Quality Standards, Fixed Costs, and Competition." *RAND Journal of Economics* 22 (4): 490–504.
- Royalty, Anne Beeson, and Neil Solomon.** 1999. "Health Plan Choice: Price Elasticities in a Managed Competition Setting." *Journal of Human Resources* 34 (1): 1–41.
- Rubinstein, Ariel.** 1982. "Perfect Equilibrium in a Bargaining Model." *Econometrica* 50 (1): 97–109.
- Serrano, Robert.** 2005. "Fifty Years of the Nash Program, 1953–2003." *Investigación Económica* 24 (2): 219–58.
- Shapiro, Carl.** 1983. "Premiums for High Quality Products as Returns to Reputations." *Quarterly Journal of Economics* 98 (4): 659–79.
- Shepard, Mark.** 2016. "Hospital Network Competition and Adverse Selection: Evidence from Massachusetts Health Insurance Exchange." NBER Working Paper 22600.
- Spence, Michael.** 1975. "Monopoly, Quality, and Regulation." *Bell Journal of Economics* 6 (2): 417–29.
- Stole, Lars A., and Jeffrey Zwiebel.** 1996. "Intra-Firm Bargaining under Non-Binding Contracts." *Review of Economic Studies* 63 (3): 375–410.
- Town, Robert, and Gregory Vistnes.** 2001. "Hospital Competition in HMO Networks." *Journal of Health Economics* 20 (5): 733–53.
- Villas-Boas, Sofia Berto.** 2007. "Vertical Relationships between Manufacturers and Retailers: Inference with Limited Data." *Review of Economic Studies* 74 (2): 625–52.
- Zaretsky, H., and PMPM Consulting Group Inc.** 2005. *Report on the Analysis of the CalPERS/Blue Shield Narrow Network*. Sacramento, CA: California Department of Managed Health.