

Equilibrium Provider Networks: Bargaining and Exclusion in Health Care Markets*

Kate Ho[†] Robin S. Lee[‡]

August 2017

(Please click [here](#) for most recent version)

Abstract

Why do insurers choose to exclude medical providers, and when would this be socially desirable? We examine network design from the perspective of a profit-maximizing insurer and a social planner to evaluate the welfare effects of narrow networks and restrictions on their use. An insurer may engage in exclusion to steer patients to less expensive providers, cream-skin enrollees, and negotiate lower reimbursement rates. Private incentives for exclusion may diverge from social incentives: in addition to the standard quality distortion arising from market power, there is a “pecuniary” distortion introduced when insurers commit to restricted networks in order to negotiate lower rates. We introduce a new bargaining solution concept for bilateral oligopoly, *Nash-in-Nash with Threat of Replacement*, that captures such bargaining incentives and rationalizes observed levels of exclusion. Pairing our framework with hospital and insurance demand estimates from Ho and Lee (2017), we compare social, consumer, and insurer-optimal hospital networks for the largest non-integrated HMO carrier in California across several geographic markets. We find that both an insurer and consumers prefer narrower networks than the social planner in most markets. The insurer benefits from lower negotiated reimbursement rates (up to 30% in some markets), and consumers benefit when savings are passed along in the form of lower premiums. A social planner may prefer a broader network if it encourages the utilization of more efficient insurers or providers. We predict that, on average, network regulation prohibiting exclusion has no significant effect on social surplus but increases hospital prices and premiums and lowers consumer surplus. However, there are distributional effects, and regulation may prevent harm to consumers living close to excluded hospitals.

Keywords: health insurance, narrow networks, selective contracting, hospital prices, bargaining, bilateral oligopoly
JEL: C78, I11, L13

*We thank Liran Einav, Glenn Ellison, Gautam Gowrisankaran, Paul Grieco, Phil Haile, Barry Nalebuff, Ariel Pakes, Mike Riordan, Bill Rogerson, Mark Shepard, Bob Town, Mike Whinston, Alexander Wolitzky, Ali Yurukoglu, and numerous conference and seminar participants for helpful discussion. Lee gratefully acknowledges support from the National Science Foundation (SES-1730063). All errors are our own.

[†]Columbia University and NBER, kh2214@columbia.edu.

[‡]Harvard University and NBER, robinlee@fas.harvard.edu.

1 Introduction

Since the passage of the Affordable Care Act (2010) there has been growing concern among policy-makers 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 towards 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, 2000), 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 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 effect of such regulation depends 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. Features of the U.S. 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 bilaterally 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 U.S. commercial (employer-sponsored) health insurance market. We begin with a simple framework that isolates the fundamental economic trade-offs when deciding whether or not 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 U.S. commercial health care market developed and estimated in Ho and Lee (2017)—which incorporates insurer-employer bargaining over premiums and consumer demand for hospitals

¹An Associated Press survey in March 2014 found, for example, 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, TX 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, available at <http://www.usnews.com/news/articles/2014/03/19/concerns-about-cancer-centers-under-health-law>.

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 Giovanelli, Lucia and Corlette (2016) for additional examples and discussion.

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 predict 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 empirical component of our analysis is an estimated model of insurer and hospital demand from Ho and Lee (2017) that leverages detailed admissions, claims and enrollment data from the California Public Employees’ Retirement System (CalPERS), a large benefits manager. This estimated model enables us to predict how consumers’ insurance enrollment and hospital utilization decisions—inputs into insurers’ revenues and costs—are affected by counterfactual changes in insurer networks. Importantly, our demand estimates condition on an individual’s age, gender, zipcode, and diagnosis, and ensure that our analysis is able to capture aspects of insurers’ incentives for cream skimming and selection.

Our demand estimates and simulations are based on data from 2004. During this period, CalPERS provided access to three large insurance plans for over a 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; 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 across twelve distinct geographic markets in California. Motivated by our empirical setting, we hold fixed the hospital networks offered by Blue Shield’s competitors as they are either complete (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.⁴ We

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

⁴There are additional institutional constraints governing premium setting and cost-sharing, 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. These requirements exacerbate insurers’ incentives to exclude high-priced hospitals since premiums cannot easily be increased solely for the consumers that most value the hospital. Consumer cost-sharing at the point

view these simulations as predicting the likely effects of removing network constraints imposed by the DMHC on Blue Shield, and 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. 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.

The key methodological contribution of this paper is the development of a new bargaining concept that extends one which has been used in previous empirical work on insurer-hospital negotiations (e.g., Gowrisankaran, Nevo and Town (2015) and 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, 2016), the bargaining concept that we build upon predicts that each hospital is paid a fraction of its marginal contribution to an insurer’s network; 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 upon a bargaining disagreement with another hospital outside the current network. For example, if an insurer negotiates with only one of two children’s hospitals in a market, its disagreement point under *Nash-in-Nash* bargaining typically involves having *no* children’s hospital in its network; instead, it may be more plausible that the insurer is able to form a contract with the other hospital upon disagreement.

Our new bargaining concept, *Nash-in-Nash with Threat of Replacement* (NNTR), relaxes these restrictions. It both endogenizes the choice of an insurer’s network, and allows the insurer to replace an included hospital with an excluded alternative. Importantly, we show that while the *Nash-in-Nash* concept has difficulty rationalizing *any* exclusion introduced by insurers in the year after our data, our NNTR concept does not. We prove that our NNTR solution always exists in our setting, and provide a non-cooperative extensive form and conditions under which a single network and set of negotiated prices, governed by this solution concept, emerge as the unique equilibrium outcome. The NNTR solution is not limited specifically to health care settings, and thus we believe that this concept and others based on it may prove useful in examining network formation and selective contracting in other industrial organization settings.⁵

of care is also very limited: in particular, Blue Shield charges a co-payment to consumers for hospital episodes that is fixed across hospitals and therefore has no effect on hospital choice. Again this generates an incentive to exclude, since the insurer bears the cost of adding a high-priced hospital to the network.

⁵Our model focuses on a setting where a single “downstream” firm (an insurer, Blue Shield) is able to leverage other “upstream” firms (hospitals) in its negotiations; in our setting, hospitals do not have alternative insurers to replace Blue Shield with, as they already contract with Blue Cross and cannot contract with Kaiser. Our framework most directly applies to environments in which a single firm can credibly negotiate with a subset of potential contracting partners; potential examples include a large retailer (e.g., Amazon) negotiating with upstream suppliers, or a monopolist content provider (e.g., a sports team selling distribution rights) with multiple downstream distributors.

Overview of Results. For each of our twelve geographic markets, 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.⁶ 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 be full; 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. 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% 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% in some markets). Under the consumer-optimal network, average rate reductions are even larger (20%) 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 full 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 with excluded suppliers, enables the firm to obtain substantial reductions in negotiated input prices.⁷ 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, this “pecuniary incentive” to distort a network away from the social optimum in order to negotiate better rates is likely to be present in other settings, including environments where firms commit *ex ante* to the number of agents with which they will negotiate (e.g., in pharmaceutical formulary design). As we show, other bargaining concepts may be ill-suited to capture this dynamic.

Finally, we use our results to inform the impact of network regulation in health care markets.

⁶Our 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.

⁷This mechanism is similar to that in Bolton and Whinston (1993), where limited supply can advantage a supplier in negotiations with multiple buyers.

Both the magnitude of network (or quality) distortions in the absence of regulation, and the effect of regulation on prices, are empirically substantive. We find that a requirement that Blue Shield contract with all hospitals, 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% 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 from *any* form of exclusion. Second, as is generally the case with complicated interventions, averages mask considerable heterogeneity. Significant distributional effects of regulation are likely when consumers differ in their preferences over product attributes. This is the case in our setting because consumers are spatially distributed and 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.

Prior Literature. We contribute to a nascent but growing literature examining narrow health care networks. Papers including Gruber and McKnight (2014) and Dafny, Hendel and Wilson (2016) study the relation between network breadth of observed plans and utilization choices, costs and premiums. We focus on the emergence of narrow networks and potential welfare consequences of counterfactual regulatory schemes by developing a model that allows us to predict equilibrium hospital networks and negotiated prices and premiums.

The development of our Nash-in-Nash with Threat of Replacement bargaining concept relies on results from Manea (forthcoming), who studies the resale of a single good through a network of intermediaries, to our setting where firms can form agreements with multiple partners and there

exist contracting externalities. Related to our analysis is Lee and Fong (2013), which posits a dynamic formation network game with bargaining in bilateral oligopoly. 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 adjustment costs and frictions (which we abstract away in this paper). There are additional papers that examine variants of the static Nash-in-Nash bargaining protocol in the hospital-insurer setting. Many of them incorporate incentives for insurers to use provider exclusion to select enrollees based on both probability of illness and preferences for high cost providers, as in Shepard (2015).⁸ Ghili (2016) and Liebman (2016) allow excluded hospitals to affect insurers’ negotiated rates with included hospitals, as in this paper.⁹ These two papers incorporate their amended bargaining frameworks into an estimated model of the insurer-hospital health care market similar to that in Ho and Lee (2017), with the primary objective of quantifying the impact of narrow networks on negotiated prices. Our focus is on the broader issue of network regulation and its welfare and distributional implications. The impact of regulation on insurer-hospital negotiations is one input into the welfare and efficiency considerations that we analyze.

Finally, the nature of our 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 a competitive health insurance exchange market with empirical estimates of the joint distribution of risk preferences and health status in order to simulate equilibria under different hypothetical exchange designs.

2 Network Design in U.S. Health Care Markets

While the concept of selective contracting has been present in the market 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% 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 (2016) document that only 57% of potential links were formed by HMO plans on the 2014 Texas exchange. The 2015 Employer Health Benefits Survey, released by the Kaiser Family Foundation, suggests narrowing is occurring also in the employer-sponsored health insurance market: for seventeen 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

⁸See also Prager (2016), which shows that similar incentives exist when insurers offer tiered hospital networks in which some hospitals are available at lower co-insurance rates than others; and Arie, Grieco and Rachmilevitch (2016), which incorporates repeated interaction and limits on the number of simultaneous negotiations by the same insurer.

⁹There are differences in the particular bargaining concepts that are used. E.g., Ghili (2016) posits conditions for price and network stability, providing a non-cooperative implementation only for the case for two hospitals and a single insurer. Liebman (2016) examines a bargaining protocol adapted from Collard-Wexler, Gowrisankaran and Lee (2016), allows for an insurer to commit to the maximum number of hospitals that it will contract with, and allows for random sets of hospitals to make offers to the insurer in the case of disagreement.

selected providers. Nine percent of employers reported that their plan eliminated hospitals or a health system to reduce costs, and seven percent offered a plan considered to be a narrow network plan.¹⁰

2.1 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., 2015). 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 towards 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.

2.1.1 Baseline Analysis

To build intuition, consider the incentives facing a monopolist insurer choosing the set of hospitals to include in its network.¹¹ 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.

¹⁰This survey was released jointly by the Kaiser Family Foundation and the Health Research & Educational Trust. Survey results available at <http://kff.org/health-costs/report/2015-employer-health-benefits-survey/>.

¹¹ Such an insurer may be thought of as optimizing relative to a non-strategic 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.

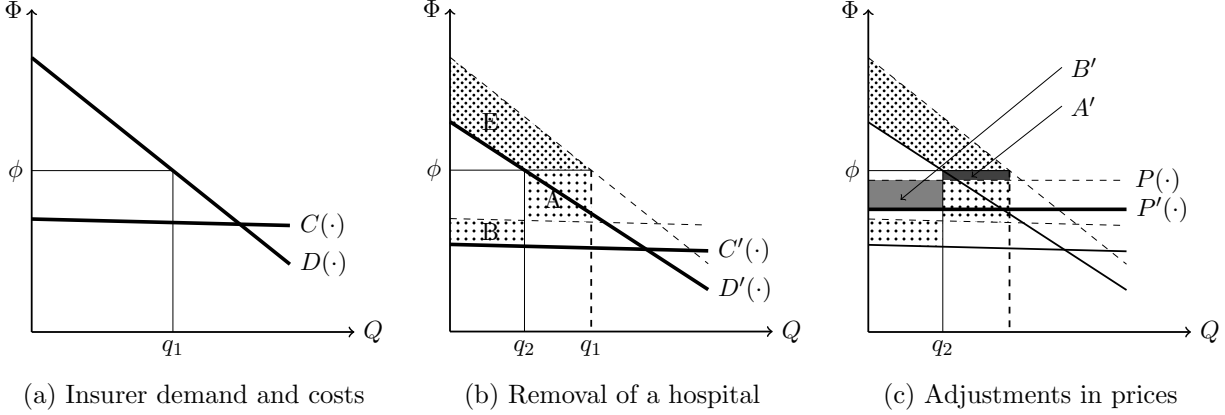


Figure 1: (a) depicts demand $D(\cdot)$ and costs $C(\cdot)$ for a hypothetical monopolist insurer; q_1 is realized demand at a fixed premium ϕ . (b) depicts new demand $D'(\cdot)$ and costs $C'(\cdot)$ upon removal of a hospital, resulting in new quality q_2 : if the insurer reimburses providers at cost, areas A is the reduction in premium revenues, B is the savings in costs, and E is the reduction in consumer surplus. (c) depicts potential adjustments in reimbursement payments $P(\cdot)$ to $P'(\cdot)$ upon removal of a hospital: A' is the reduction in insurer premium revenues, and B' is the savings in payments to hospitals.

Figure 1a 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.¹²

Consider what might occur if the insurer drops a hospital from its network, depicted by moving from Figure 1a to Figure 1b. 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 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 1b. 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 skinning.

However, a social planner would also consider the change in inframarginal consumer surplus for current enrollees if the hospital were removed—a consideration ignored by the profit-maximizing

¹²Note that the marginal cost curve may be downward sloping in the presence of adverse selection.

monopolist optimizing over quality (Spence, 1975)—as well as the 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).

2.1.2 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.

Figure 1c 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.¹³ 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 Figure 1b, $A > B$ 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 Figure 1c, $A' < B'$, 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.

¹³There 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.

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.

2.2 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, 2015). 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.

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 depends 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 towards 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.

3 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 comprise 10% of the total commercially insured population of the state. We observe the set of insurance plans offered to CalPERS enrollees, their enrollment choices, and medical claims and admissions information in 2004 (detailed further in Section 6.1).

3.1 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 (BS), 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; BS's network was somewhat narrower, containing approximately 85% of the hospitals in Blue Cross's

network;¹⁴ 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.

In June 2004, Blue Shield filed a proposal with the California Department of Managed Health Care (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 Health Care Service Plan Act (1975), a piece of legislation that regulates California’s managed care insurance plans. The resulting DMHC “Report on the Analysis of the CalPERS/Blue Shield Narrow Network” (Zaretsky and pmpm Consulting Group Inc. (2005)) describes the Blue Shield proposal as offering “a vastly different approach to cost savings” compared to other employers’ use of co-payments, deductibles or cost sharing. 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.¹⁵ Some of the hospitals that BS 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 Appendix Table A1 (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 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 can threaten replacement. 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 ideal for studying network design for several reasons. First, we have sufficiently detailed data (from 2004, the year before the BS 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

¹⁴The number of in-network hospitals only count those with least 10 admissions in our data for a particular insurer. We obtain BC hospital network information directly from the insurer; for BS, we infer the hospital network by including all hospitals had claims data indicating that the hospital was a “network provider.”

¹⁵The 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.

hospitals than requested.¹⁶ Our simulations therefore capture some 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 exclusion incentives for some but not others. We use the list of hospitals in the Blue Shield proposal as a check on the ability of our model to make predictions that are close to the data. 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.

3.2 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 U.S. commercial health care market. We rely on an estimated 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—also 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 whether they are included in their network, and if so the reimbursement rates that are paid.
- 1b. Premium Setting: Simultaneously with the determination of hospital networks and negotiated rates, the employer and the set of MCOs bargain over per-household premiums.
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 that 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

¹⁶As noted below, the fact that Blue Shield’s networks were constrained to be essentially complete in the year of our data, 2004, is useful for our simulations. The reason is that, when networks are complete, the Nash-in-Nash model assumed in Ho and Lee (2017) has the same predictions as the NNTR model developed in this paper; see also Section 6.2.

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 concept 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 5.

4 Equilibrium Hospital Networks and Reimbursement Rates

We now extend the framework of the U.S. commercial health care market developed in Ho and Lee (2017) by (i) incorporating a bargaining solution that allows an insurer to “play off” hospitals with those excluded from its network in order to negotiate more advantageous rates, and (ii) providing a way to predict an insurer’s hospital network.

4.1 Intuition

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, 2016), 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., 2015) and health care (e.g., Grennan, 2013; Gowrisankaran, Nevo and Town, 2015; Ho and Lee, 2017) settings.¹⁷ Defined for a particular network, the Nash-in-Nash solution in our health care context specifies that the reimbursement price negotiated between each hospital and MCO solves that pair’s Nash bargaining problem given that 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 (see also Lee and Fong, 2013). First, as the Nash-in-Nash solution is often characterized as a surplus division rule for a given network, it provides limited guidance as to which network(s) might arise in equilibrium (i.e., only that

¹⁷This 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 respective Nash bargaining product given the actions of other pairs.

bilateral agreements that are observed to form generate positive “gains-from-trade”). Second, the Nash-in-Nash solution typically does not allow parties to adjust their contracting decisions when evaluating disagreement points from any particular bilateral bargain. 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. In general, the Nash-in-Nash solution does not typically allow firms to form new contracts, or adjust other contracts, in the event of a bargaining disagreement.

This issue is not innocuous. Failing to accurately account for parties’ true outside options 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 with static disagreement points implies that the presence of any hospitals that are excluded from the network has no effect on the negotiated rates with in-network providers. Hospitals are reimbursed based 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 value it generates when contracting with an insurer.¹⁸ 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 towards broader and more inclusive hospitals networks partly explains why, as we will show, the Nash-in-Nash solution has difficulty rationalizing observed levels of exclusion in our empirical application.

We thus depart from a direct application of Nash-in-Nash in order to capture MCOs’ exclusionary incentives, and develop a new bargaining solution that we refer to as the *Nash-in-Nash with Threat of Replacement* (NNTR) solution. 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, NNTR effectively 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. Importantly, 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 does not require 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.

Though the Nash-in-Nash concept may understate the extent to which an insurer can form selective networks and play hospitals off one another, it is important to note that 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

¹⁸This value can be summarized by the higher premium revenues that the insurer obtains as a result of having the hospital in its network, net of any increases in costs borne by the insurer and hospital.

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.¹⁹

Finally, we emphasize that 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 the hospital networks offered by Blue Shield’s competitors are either complete (Blue Cross) or integrated (Kaiser) 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.

The rest of this section proceeds as follows. 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. As we will discuss, such a solution will only be defined 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. Second, we provide a non-cooperative extensive form that, under certain conditions that we specify, admits a unique equilibrium network and set of negotiated reimbursement rates. As firms become patient, the network that emerges coincides with what we refer to as the *insurer optimal stable network*, and the negotiated rates converge to the NNTR bargaining solution. This exercise provides support for the reasonableness of the NNTR bargaining solution, how it might arise in practice, and why—in our empirical application—it is plausible that Blue Shield can commit to and eventually form the network that maximizes its equilibrium profits under NNTR bargaining.

A reader interested in our empirical application and results can skip to Section 5.

4.2 Nash-in-Nash with Threat of Replacement

Setup and Notation. To begin, we introduce notation and delineate objects that are assumed to be primitives of our analysis.

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.

¹⁹Even 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).

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 $\pi_j^M(G, \mathbf{p}) \equiv \tilde{\pi}_j^M(G) - \sum_{i \in G} D_{ij}^H(G) p_{ij}$ and $\pi_i^H(G, \mathbf{p}) \equiv \tilde{\pi}_i^H(G) + \sum_{n \neq j} D_{in}^H(G) p_{in}$ 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 derive from 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 non-payment 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 , 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 are formed at prices \mathbf{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 $\mathbf{p}^*(G) \equiv \{p_{ij}^*(G, \mathbf{p}_{-ij}^*)\}$, where $\forall i \in G$, each p_{ij}^* can be interpreted as the outcome of a Nash bargain between the two parties (MCO j and hospital i) where: (i) 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 in $G \setminus i$), and (ii) MCO j has an outside option of being able replace hospital i with some hospital k not in network G at the minimal price k would be willing to accept.²⁰

Formally, each individual NNTR price paid by MCO j satisfies

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

²⁰As noted above, we hold fixed the hospital networks and reimbursement prices $\{p_{in}\}$ for other MCOs $n \neq j$.

where

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

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

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

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:

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

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 $\mathbf{p}_{-ij}^*(\cdot)$; 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.²¹ 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.²²

Example. Assume 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 $\tau_j = 1/2$.

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 non-negative payment for inclusion in the MCO's network; thus hospital k 's reservation price

²¹The concept 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.

²² Hospital k may earn profits even if excluded from MCO j 's network as it may contract with other MCOs $-j$.

$p_{kj}^{res} = 0$. 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). In a sense, the option of contracting with k is not a ‘‘credible threat’’; 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 . (Later, we provide explanations for why paying hospital k its reservation price may be a credible threat for the MCO).

Existence. To guarantee that NNTR prices exist for a given network G , we restrict each NNTR price to lie on a compact interval of the real line: i.e., we adjust (1) so that:

$$p_{ij}^*(\cdot) = \max \left\{ -\bar{p}, \min\{p_{ij}^{Nash}(G, \mathbf{p}_{-ij}^*), p_{ij}^{OO}(G, \mathbf{p}_{-ij}^*), \bar{p}\} \right\} \forall i \in G, \quad (5)$$

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

Given this restriction, we establish the following result:

Proposition 4.1. *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 for MCO j , $\mathbf{p}_j^*(G) \equiv \{p_{ij}^*(G)\}_{i \in G}$, that satisfy (5).*

(All proofs in the 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 $\mathbf{p}^*(G)$ have not been found for any network.

Stability. 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. Stability of an agreement $i \in G$ at prices \mathbf{p} implies that neither party has a unilateral incentive to terminate their agreement, holding fixed all other agreements $G \setminus i$.

A particular agreement $i \in G$ can be unstable at p_{ij}^* if, given G and other prices \mathbf{p}_{-ij}^* , (i) the Nash bargaining problem represented by (2) has no solution, as there is no price p_{ij}^{Nash} for which both parties wish to come to agreement (given other agreements $G \setminus i$ have been formed at \mathbf{p}_{-ij}^*); or (ii) at $p_{ij}^{OO}(G)$, hospital i would rather not come to agreement with MCO j (i.e.,

$[\Delta_{ij}\pi_i^H(G, \{p_{ij}^{OO}(\cdot), \mathbf{p}_{-ij}^*\})] < 0$). Case (i) is typically ruled out for observed agreements in applications of Nash bargaining to bilateral oligopoly: if there are no bilateral gains from trade, an agreement is not typically expected to form. However, negative (total) bilateral gains-from-trade can arise in settings where there are contracting externalities: e.g., as discussed previously, one way in which this can occur in our setting is if hospital i is a high cost hospital that delivers very little incremental value to consumers; an MCO j thus may thus be better off excluding hospital i than including it at cost. Case (ii) is a new source of instability in our setting, and emerges due to MCO j 's ability to replace i with a hospital k outside of a given network.

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

Proposition 4.2. *Network G is stable at NNTR prices \mathbf{p}^* iff, for all $i \in G$,*

$$[\Delta_{ij}\Pi_{ij}(G, \mathbf{p}^*)] \geq [\Delta_{kj}\Pi_{kj}(G \setminus i \cup k, \mathbf{p}^*)] \quad \forall k \in (\mathcal{H} \setminus G) \cup \emptyset.$$

Thus, case (i) above, under which the Nash bargaining problem has no solution for some $i \in G$, would lead to $[\Delta_{ij}\Pi_{ij}(G, \mathbf{p}^*)] < 0$. Further, if some agreement $i \in G$ is unstable because $p_{ij}^{OO}(\cdot)$ is low enough so that hospital i would rather reject than accept the payment, it means that there is some other hospital $k \notin G$ that generates higher bilateral surplus with the MCO than i . As a result, MCO j may not be able to credibly exclude hospital k from its network.

Note that stability only tests whether any agreement $i \in G$ at a given set of prices \mathbf{p} does not wish to be terminated by either party involved. It may be the case that agreements not contained in G would be profitable to form if all agreements in G remained fixed at the same set of prices; since the formation of a new link may be seen as a bilateral deviation, we do not impose this condition as a requirement for stability.²³ In our extensive form representation provided in the next subsection, we allow the insurer to choose a stable network that it wishes to form, and provide conditions under which any equilibrium results in that network forming at NNTR prices. Thus, allowing the insurer to commit to a network is the manner by which the most profitable set of agreements for the insurer is determined (at prices that can credibly be negotiated given all agents have consistent expectations over the network that forms).

Finally, let \mathcal{G}_j^S be the set of all networks $G \in \mathcal{G}_j$ that are stable at NNTR prices $\mathbf{p}^*(G)$. We define the *insurer optimal stable network* to be the stable network that maximizes insurer's profits at NNTR prices: i.e., $G^{*,ins} = \arg \max_{G \in \mathcal{G}_j^S} \pi_j^M(G, \mathbf{p}^*(G))$.²⁴

Discussion. 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 currently in j 's network k 's reser-

²³ The profitability of unilateral deviations is often determined by holding fixed the actions of other agents. We believe that this assumption is less reasonable under a multilateral deviation such as that where a new hospital is added to the network.

²⁴If 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.

vation price, then $p_{ij}^*(\cdot)$ would emerge as the outcome of certain non-cooperative implementations of the Nash bargaining solution. For example, Binmore, Shaken 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; they show that the unique subgame equilibrium outcome converges to the Nash bargaining solution as the discount factor approaches one *unless* the outside option is binding, in which case the outside option acts as a constraint on the bargaining solution and the side for which the outside option is binding obtains exactly its outside option. This insight, which they refer to as the “outside option principle,” highlights the different roles that outside options and disagreement points play. In many circumstances, outside options only affect bargaining outcomes if they are “credible”: i.e., they would deliver payoffs greater than would be achievable in the bargaining game without outside options (see also Muthoo, 1999).

There are two important complications that arise in 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 option is—in particular, when bargaining with hospital i , why an MCO can credibly (threaten to) 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 agreement. 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 $p_{ij}^*(\cdot) = p_{ij}^{Nash}(\cdot) \forall i \in G$.²⁵ 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 , $i \in G$. In (3), MCO j can threaten hospital i with replacing it with some hospital k 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 —if it were to exercise its outside option and replace i with k —its reservation price. That is, why wouldn’t k demand more? As we will discuss further in the next subsection when providing a non-cooperative extensive form that generates the NNTR solution, the ability for the MCO to *commit* to negotiating with any stable network of hospitals bestows upon it the ability to effectively “play off” hospitals

²⁵Collard-Wexler, Gowrisankaran and Lee (2016) provide a non-cooperative foundation for this solution concept based on an alternating offers bargaining game where agents negotiate fixed fee transfers and can engage in multilateral deviations.

that are included and those excluded from its network. It will also be clearer why the reservation price that MCO j must pay k , $p_{kj}^{res}(G \setminus i, \mathbf{p}_{-ij}^*)$, is given by (4).

4.3 Microfoundation for Nash-in-Nash with Threat of Replacement

We now provide and analyze a non-cooperative 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 extensive form game:

- At period 0, MCO j publicly announces a network $G \subseteq \mathcal{G}_j^S$, and sends separate representatives r_i to negotiate with each hospital i s.t. $i \in G$;²⁶
- At the beginning of each subsequent period $t > 0$, each representative r_i that has not yet reached an agreement simultaneously and privately chooses to “engage” with either hospital i or with some other hospital k , $k \notin G$. 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 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.

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.²⁷ For establishing the main result for this section, 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 MCO receives $(1 - \delta)\pi_j^M(G^t)$ and

²⁶The restriction to announcing only *stable* networks is made for expositional 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 representatives’ beliefs at the beginning of subsequent negotiations (see also Lee and Fong, 2013). In such a case, we can extend our analysis to allow for the MCO to announce some unstable network G , but restrict attention to equilibria in which representatives are then subsequently sent only to hospitals contained in some stable network $G' \subset G$ that is publicly observable.

²⁷In 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 $\delta \rightarrow 1$. Adopting an alternative specification—in which representatives engaging with the same hospital can coordinate their negotiations for that period only, and all representatives negotiating with that hospital are removed if an agreement is reached—will not affect our main result.

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 0.²⁸ (As discussed above, these profit functions condition on insurer premiums and subsequent consumer enrollment and utilization decisions).

Our extensive form adapts part of the protocol introduced in Manea (forthcoming) 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 (forthcoming) 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 $\tau \times v_1 > v_2$). In other words, the outcome is interpretable as one in which the seller engages in Nash bargaining with the highest valuation buyer, and has an outside option of extracting full surplus from the second highest valuation buyer; i.e., the “outside option” principle of Binmore, Shaken and Sutton (1989) applies here.²⁹

Our extensive form also makes clear the role of *commitment* in generating the NNTR solution. In our particular extensive form, MCO j is able to commit in period 0 to negotiating with a particular stable network and to the maximum number of hospitals that it contracts with. This will provide it with additional bargaining leverage as it can credibly exclude certain hospitals from its network when negotiating with hospitals.³⁰

Equilibrium. As in Manea (forthcoming), we restrict attention to stationary Markov perfect equilibria (henceforth, 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

²⁸Using separate representatives has also been used to motivate the Nash-in-Nash bargaining solution (see also the discussion in Collard-Wexler, Gowrisankaran and Lee, 2016). We follow Collard-Wexler, Gowrisankaran and Lee (2016) in our treatment of profit flows and focus on lump-sum transfers.

²⁹In the equilibria that Manea (forthcoming) 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. See also Bolton and Whinston (1993) who examine a related setting with a single upstream firm bargaining over input prices with two downstream firms (using a different protocol), and obtain similar results when the upstream firm has only a single unit to sell.

³⁰The insight that an insurer may find it beneficial to commit to a certain set of counterparties with which to bargain is related to Stole and Zweibel (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.

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. Similarly, 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 subgame).

Proposition 4.3. *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 4.2. In equilibrium, representatives for both the MCO and hospitals must have consistent expectations over the set of agreements that will form following such an announcement.

Proposition 4.4. *Consider any 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 4.4 is straightforward. First, consider any subgame where MCO j announces a network G that contains a single hospital. By Proposition 4.2, 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 (forthcoming): in any MPE, in the limit either the MCO obtains τ of the total gains from trade $[\Delta_{ij}\Pi_{ij}(G)]$ (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 $[\Delta_{kj}\Pi_{kj}(\{kj\})] > \tau[\Delta_{ij}\Pi_{ij}(G)]$, where hospital k is the second-highest-surplus creating hospital—the MCO obtains the equivalent of $[\Delta_{kj}\Pi_{kj}(\{kj\})]$ by paying only the lump-sum equivalent of $p_{ij}^{O}(G)$ to hospital i .

Next, consider any subgame where MCO j announces a network G that contains two or more hospitals. As long as each representative r_i when negotiating 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 upon and adapt the techniques used in Manea (forthcoming) to our setting.

Having analyzed behavior in each subgame, we then can establish our main result.

Proposition 4.5. *For any $\varepsilon_1, \varepsilon_2 > 0$, there exists $\underline{\delta} < 1$ and $\underline{\Lambda} < 1$ such that for $\delta > \underline{\delta}$, 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;
- the probability that $G^{*,ins}$ is formed is greater than $1 - \varepsilon_1$;
- all agreements $i \in G^{*,ins}$ are formed with ε_2 of NNTR prices.

Such an equilibrium exists.

Thus, if we restrict attention to MPE in which the announced period-0 network is formed with sufficiently high probability, the unique MPE outcome as $\delta \rightarrow 1$ will be the insurer optimal stable network $G^{*,ins}$ being announced and formed with probability arbitrarily close to 1 at prices arbitrarily close to NNTR prices.

An Alternative Interpretation: Sequential Negotiation. An alternative interpretation for why $G^{*,ins}$ and \mathbf{p}^* may be perceived as reasonable outcomes is if contracts for individual hospitals were assumed to come up for renewal one at a time, and firms did not anticipate future renegotiation following the conclusion of any bargain. For example, assume that MCO j 's network was $G^{*,ins}$ and negotiated prices were \mathbf{p}^* . Assume some hospital contract $i \in G^{*,ins}$ came up for renewal, but all other contracts were fixed. If new negotiations over ij occurred according to the previously described protocol (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 choose either the MCO or hospital selected to make an offer) and the MCO could not sign up more than one hospital (including i), then a direct application of Proposition 4.4 implies that the unique limiting MPE outcome would be the MCO renewing its contract with hospital i at price p_{ij}^* . Given this interpretation, we believe that treating hospitals not on the network (for which new contracts can be signed at any point in time) asymmetrically from hospitals already on the network (which plausibly have contracts currently in place that cannot be adjusted) is an appealing feature of the NNTR solution.

4.4 Extensions

Premium Setting. In our application, we will allow for all MCO premiums to adjust in response to potential network changes and rate negotiations. As we noted in Section 3.2, 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 firm's 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.³¹

³¹We 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.

Synthesizing Nash-in-Nash and NNTR. Finally, our NNTR concept may be extended to allow for the possibility that the insurer may not be able to perfectly act upon its “threat of replacement.” Such a modification allows our bargaining concept to nest both the the Nash-in-Nash concept and the previously defined NNTR concept via the addition of a single parameter $\alpha \in [0, 1]$.³² Formally, for a given network G , define the NNTR- α price $p_{ij}^*(\cdot; \alpha) = \min\{p_{ij}^{Nash}(G, \mathbf{p}_{-ij}^*), (1 - \alpha)p_{ij}^{Nash}(G, \mathbf{p}_{-ij}^*) + \alpha p_{ij}^{OO}(G, \mathbf{p}_{-ij}^*)\} \forall i \in G$ where $p_{ij}^{Nash}(\cdot)$ and $p_{ij}^{OO}(\cdot)$ are defined as in (2)-(3). Note that NNTR-0 prices ($\alpha = 0$) correspond to Nash-in-Nash prices, and NNTR-1 prices ($\alpha = 1$) correspond to the NNTR prices defined in (1).

5 Hospital and Insurance Demand and Premium Setting

The previous section introduced and defined the Nash-in-Nash with Threat of Replacement (NNTR) bargaining solution, and provided a justification for it as the equilibrium of a non-cooperative game. The NNTR solution conditions on the hospital networks in a market, and profits for each firm conditional on 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 also summarize Stages 1b, 2, and 3 of our model, which generate equilibrium predictions for insurer premiums and consumer demand for insurance plans and hospital services. These stages follow Ho and Lee (2017); additional details are contained therein.

5.1 Firm Profits

In our empirical analysis, we assume that MCO j 's profits are

$$\pi_j^M(G, \mathbf{p}_j, \phi) = \sum_m \left(\underbrace{\phi_j \Phi' \mathbf{D}_{jm}(\cdot)}_{\text{premium revenues}} - \underbrace{D_{jm}^E(\cdot) \eta_j}_{\text{non-inpatient hospital costs}} - \underbrace{\sum_{h \in G_{jm}} D_{hj}^H(\cdot) p_{hj}}_{\text{hospital payments}} \right). \quad (6)$$

where G now represents the set of *all* MCOs' hospital networks in all markets; $\mathbf{p}_j \equiv \{p_{ij}\}_{ij \in G}$ 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

³²Such an extension may be useful for applied work seeking to use information on observed prices to infer the extent to which firms are able to leverage outside options in negotiations. In our setting, observed prices are negotiated for (essentially) complete networks; thus, we primarily rely on our NNTR concept (for $\alpha = 1$) to predict how prices adjust when networks are incomplete.

i. These demand functions arise from the consumer demand systems that are outlined below in Section 5.2.

The first term on the right-hand side of (6) 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, $\Phi \equiv [1, 2, 2.6]$ is the observed vector of premium multipliers for each household type $\lambda \in \{\text{single, two-party, family}\}$, 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 $D_{jm}^E(\cdot)$.³³ 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).³⁴

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

$$\pi_i^H(G, \mathbf{p}_i, \phi) = \sum_{n \in G_i} D_{in}^H(G, \phi) \times (p_{in} - c_i), \quad (7)$$

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 7.

5.2 Premium Bargaining and Consumer Demand

We now detail Stages 1b, 2, and 3 of the model outlined in Section 3.2 in reverse order.

5.2.1 Stages 2 and 3: Insurer and Hospital Demand

Household demand for insurers, and consumer demand for hospitals, are determined by the following utility equations which build on a previous literature (Town and Vistnes, 2001; Capps, Dranove and Satterthwaite, 2003; Ho, 2006).

³³For Blue Shield and Blue Cross, as we are explicitly controlling for prices paid to hospitals, the estimated cost parameters $\{\eta_j\}_{j \in \{BS, BC\}}$ represent non-inpatient hospital marginal costs per enrollee, which may include physician, pharmaceutical, and other fees. Since we do not observe hospital prices for Kaiser, η_{Kaiser} will also include Kaiser's inpatient hospital costs.

³⁴In our empirical application, we account for variation in disease severity across admissions by scaling D_{hj}^H by the expected Medicare diagnosis-related group (DRG) weight (which varies based on age and gender) for the relevant patient, and assume that prices negotiated between each hospital and insurer represent DRG-adjusted admission prices. See Section 6 for more details.

In stage 3 of our model, we assume that an individual of type κ (representing one of 10 age-sex categories) requires admission to a hospital with probability γ_{κ}^a . Conditional on admission, the individual receives one of six diagnoses l with probability $\gamma_{\kappa,l}$. Individual k of type $\kappa(k)$ with diagnosis l derives the following utility from hospital i in market m :

$$u_{k,i,l,m}^H = \hat{\delta}_i + z_i v_{k,l} \hat{\beta}^z + d_{i,k} \hat{\beta}_m^d + \varepsilon_{k,i,l,m}^H \quad (8)$$

where z_i are observed hospital characteristics (e.g. teaching status, and services provided by the hospital), $v_{k,l}$ are characteristics of the consumer (including diagnosis), $d_{i,k}$ represents the distance between hospital i and individual k 's zip code of residence (and has a market-specific coefficient), and $\varepsilon_{k,i,l,m}^H$ is an idiosyncratic error term assumed to be i.i.d. Type 1 extreme value. Coefficients with $\hat{\cdot}$ s, including hospital fixed effects $\hat{\delta}_i$, represent estimated parameters.

This utility equation allows for the construction of a measure of consumers' ex-ante expected utility ("willingness-to-pay" or *WTP*) for an insurer's hospital network. Given the assumption on the distribution of $\varepsilon_{k,i,l,m}^H$, individual k 's *WTP* for the hospital network offered by plan j in market m is

$$WTP_{k,j,m}(G_{j,m}) = \gamma_{\kappa(k)}^a \sum_{l \in \mathcal{L}} \gamma_{\kappa(k),l} \log \left(\sum_{h \in G_{j,m}} \exp(\hat{\delta}_h + z_h v_{k,l} \hat{\beta}^z + d_{h,k} \hat{\beta}^d) \right),$$

This object varies explicitly by age and gender. The model therefore accounts for differential responses by particular types of patients (i.e., selection) across insurers and hospitals when an insurer's hospital network changes.

Finally we assume that the utility a household or family f receives from choosing insurance plan j in market m is given by:

$$u_{f,j,m}^M = \hat{\delta}_{j,m} + \hat{\alpha}_f^{\phi} \phi_j + \sum_{\forall \kappa} \hat{\alpha}_{\kappa}^W \sum_{k \in f, \kappa(k)=\kappa} WTP_{k,j,m} + \varepsilon_{f,j,m}^M \quad (9)$$

where $\hat{\delta}_{j,m}$ is an estimated insurer-market fixed effect and ϕ_j is the premium. The premium coefficient is permitted to vary with the (observed) income of the primary household member. The third term controls for a household's *WTP* for the insurer's hospital network by summing over the value of $WTP_{k,j,m}$ for each member of the household multiplied by an age-sex-category specific coefficient, $\hat{\alpha}_{\kappa}$. Finally $\varepsilon_{f,j,m}^M$ is a Type 1 extreme value error term. This specification is consistent with households choosing an insurance product prior to the realization of their health shocks and aggregating the preferences of members when making the plan decision.

We note that this demand specification captures selection of consumers across both hospitals and MCOs based on age, gender, diagnosis, income and zipcode. Preferences for hospital characteristics are permitted to differ by diagnosis, income and location; the probability of admission to hospital, and of particular diagnoses conditional on admission, differ by age and gender; and the weight placed on the value of the network in the insurer utility equation differs by age and gender. The elasticity with respect to premiums in that equation also differs by income. Thus, while we follow

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 patients to changes in an insurer’s hospital networks (i.e. selection) can affect its incentives to exclude.

The utility equations provided in (8) and (9) are integrated over the population of families and individuals across markets in our sample to predict insurance enrollment $\{D_{jm}(\cdot), D_{jm}^E(\cdot)\}$ and hospital utilization $D_{hj}^H(\cdot)$ across MCOs and hospitals for any set of hospital networks and insurance premiums.

5.2.2 Stage 1b: Premium Bargaining

Premiums for each insurer are assumed to be negotiated with the employer through simultaneous bilateral Nash bargaining, where the employer maximizes its employees’ welfare minus its total premium payments. These negotiations take place simultaneously with the determination of hospital networks and negotiated rates. Negotiated premiums ϕ_j for each MCO j satisfy

$$\phi_j = \arg \max_{\phi} \left[\underbrace{\pi_j^M(G, \mathbf{p}, \{\phi, \phi_{-j}\})}_{GFT_j^M} \right]^{\tau^\phi} \times \left[\underbrace{W(\mathcal{M}, \{\phi, \phi_{-j}\}) - W(\mathcal{M} \setminus j, \phi_{-j})}_{GFT_j^E} \right]^{(1-\tau^\phi)} \quad \forall j \in \mathcal{M}, \quad (10)$$

(where $\phi_{-j} \equiv \{\phi \setminus \phi_j\}$) subject to the constraints that the terms $GFT_j^M \geq 0$ and $GFT_j^E \geq 0$. These terms represent MCO j ’s and the employer’s gains-from-trade from coming to agreement, i.e., from MCO j being included in the choice set that is offered to employees. The MCO’s gains-from-trade are simply its profits from being part of the employer’s choice set (where its outside option from disagreement is assumed to be 0). The employer’s gains-from-trade are represented by the difference between its “objective” $W(\cdot)$, defined as the employer’s total employee welfare net of its premium payments to insurers, when MCO j is and is not offered. The “premium Nash bargaining parameter” is represented by $\tau^\phi \in [0, 1]$.

6 Empirical Analysis

In this section we summarize the data and estimation strategy used in Ho and Lee (2017) to recover the inputs for our simulations.

6.1 Data

The primary dataset includes 2004 enrollment, claims, and admissions information for over 1.2M CalPERS enrollees. The markets we consider are the California Office of Statewide Health Planning and Development (OSHPD) health service area (HSA) definitions. There are 14 HSAs in California, which we use as our market definition. For enrollees in Blue Shield (BS) and Blue Cross (BC) we

observe hospital choice, diagnosis, and total prices paid by each insurer to a given medical provider for the admission.

The claims data are aggregated into hospital admissions and assigned a Medicare diagnosis-related group (DRG) code which we use as a measure of individual sickness level or costliness to the insurer. We categorize individuals into 10 different age-gender groups. For each we compute the average DRG weight for an admission from our admissions data, and compute the probability of admission to a hospital, and of particular diagnoses, using Census data and information on the universe of admissions to California hospitals. We use enrollment data for state employee households in 2004; for each we observe the age, gender and zip code of each household member and salary information for the primary household member.³⁵

Our measure of hospital costs is the average cost associated with the reported “daily hospital services per admission” divided by the the computed average DRG weight of admissions at that hospital (computed using our data). The prices paid to hospitals are constructed as the total amount paid to the hospital across all admissions, divided by the sum of the 2004 Medicare DRG weights associated with these admissions. We assume each hospital-insurer pair negotiates a single price index that is approximated by this DRG-adjusted average. Both this price, and the hospital’s cost per admission, are scaled up by the predicted DRG severity of the relevant admission given age and gender. Finally, we use 2004 financial reports for each of our three insurers from the California Department of Managed Health Care to compute medical loss ratios for each insurer by dividing total medical and hospital costs by total revenues.

Summary statistics are provided in Appendix Tables 2 and 3. In 2004, BS offered 189 hospitals across California, compared to 223 for BC and just 27 for the vertically integrated Kaiser HMO. Annual premiums for single households across BS, BC, and Kaiser were \$3,782, \$4,193, and \$3,665; premiums for 2-party and families across all plans were a strict 2x or 2.6x multiple of single household premiums. There was no variation in premiums across markets within California or across demographic groups. State employees received approximately an 80% contribution by their employer. We use total annual premiums received by insurers when computing firm profits, and household annual contributions (20% of premiums) when analyzing household demand for insurers.

The tables 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 large cross-hospital variation in costs 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 fairly large, leaving room for rate reductions due to exclusion of some providers. Finally, Appendix Table 3 summarizes the substantial variation in probabilities of hospital admission and DRG weights by age and gender that generate cream skimming incentives.

³⁵We also use hospital characteristics, including location, from the American Hospital Association (AHA) survey. Hospital costs are taken from the OSHPD Hospital Annual Financial Data for 2004.

6.2 Estimation

The parameters of the hospital demand equation detailed in Section 5.2 (Stage 3) are estimated via maximum likelihood using admissions data under the assumption that individuals, when sick, can go to any in-network hospital in the HSA that is within 100 miles of their zip code. The insurer demand model (Stage 2) is also estimated via maximum likelihood, using household-level data on plan choices, location and family composition, and conditioning on the set of plans available in each zip code.

Insurer non-inpatient hospital costs $\{\eta_j\}$ and the Nash bargaining weights $\{\tau^\phi, \{\tau_j\}\}$ are estimated using the detailed data and the first-order conditions implied by the model of Nash bargaining between insurers and the employer over premiums, and Nash-in-Nash bargaining between insurers and hospitals over hospital prices. A third set of moments is generated from the difference between each insurer’s medical loss ratio (obtained from the 2004 financial reports) and the model’s prediction for this value.

Internal Consistency of Parameter Estimates. During the period of our study, CalPERS and the California Department of Managed Health Care constrained networks to be close to full in reality: in 11 out of the 12 markets that we focus on, all five of the largest systems are included in the BS network. Recall that in the markets where the BS network is full, the Nash-in-Nash and NNTR bargaining outcomes coincide. We have also re-estimated parameters using only data from the 11 markets in which the BS network was full, and recovered parameters that were not statistically different. Thus, we use the estimated values of $\{\eta_j, \tau^\phi, \{\tau_j\}\}$ from Ho and Lee (2017) in our simulations under both bargaining models, and hold these values constant in our simulations.

Estimated Values. Several of the estimates can be compared to outside sources. For example the estimated values of insurers’ non-inpatient hospital costs per enrollee per year are approximately \$1,690 for BS and \$1,950 for BC. Total (including hospital) marginal costs are estimated to be \$2,535 for Kaiser. 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. Data from the Massachusetts Center for Health Information and Analysis indicate average spending of \$1,644 per person per year on professional services for the three largest commercial insurers in the years 2010-12.³⁶ Similarly, the own-premium elasticities for each insurer can be compared to those estimated in previous papers.³⁷ The estimated magnitudes range from -1.23 for single-person households for Kaiser to -2.95 for families with children for BC. These numbers are well within the

³⁶Kaiser data accessed from <http://kff.org/other/state-indicator/health-spending-per-capita-by-service/> on February 25, 2015. Massachusetts data were taken 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.

³⁷We report elasticities based on the full 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.

range estimated in the previous literature. For example 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 (1998) 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.

Overall, we argue that these estimates provide a reliable basis for our simulations and accurately capture insurer and hospital incentives. Individuals' hospital choices are allowed to vary by age, gender and location. Distance from the individual's home to the hospital is a key determinant of hospital choice, as is the fit between the patient's diagnosis and the services offered by the hospital. Expected probabilities of admission, and of particular diagnoses given admission, differ by age and gender. Finally, household preferences over insurance plans reflect each family member's valuation for the hospitals offered.

7 Simulations: 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 twelve health service areas (HSAs) in California.³⁸ 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 offer a network with all hospitals.

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 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

³⁸In two HSAs (Northern California and Mid Coast), Kaiser is not available; we exclude these HSAs from our analysis to maintain comparability across markets.

reimbursement rates are determined on a market-by-market level, and thus we conduct our simulations separately for each HSA.³⁹ Third, we focus only on Blue Shield’s negotiations with the five largest systems (by admissions) in each market.⁴⁰ 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% of admissions).

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 full network, 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 Appendix C.

7.1 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 “full” 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 standard 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);

³⁹In reality CalPERS constrains premiums to be set at the state level, and hospital systems may engage in state-wide bargaining.

⁴⁰In the San Francisco market, we observe only four hospital systems that have more than 10 admissions from either Blue Shield or Blue Cross in our data. As a result, we only consider the contracting decisions between Blue Shield and the four largest systems for this market.

Table 1: Simulation Results for All Markets (Averages)

Objective		Social	Consumer	Blue Shield		Full
		(NNTR)	(NNTR)	(NNTR)	(NN)	(NNTR/NN)
Surplus (\$ per capita)	BS Profits	1.5%	1.4%	2.6%	0.0%	304.7
		[1.1%,6.9%]	[0.9%,8.0%]	[1.8%,8.6%]	[0.0%,0.0%]	[287.5,312.1]
	Hospital Profits	-6.4%	-22.9%	-14.7%	0.0%	170.0
		[-24.9%,-4.9%]	[-37.7%,-15.0%]	[-33.0%,-12.8%]	[0.0%,0.0%]	[159.4,209.4]
	Total Hosp Costs	0.2%	0.7%	0.5%	0.0%	95.6
		[0.0%,1.9%]	[0.0%,2.5%]	[0.4%,2.0%]	[0.0%,0.0%]	[94.1,96.3]
	Total Ins Costs	-0.1%	0.1%	-0.1%	0.0%	2008.5
		[-0.4%,-0.1%]	[-0.3%,0.2%]	[-0.5%,-0.1%]	[0.0%,0.0%]	[1990.4,2025.7]
Transfer / Cost (\$ per enrollee)	BS Premiums	-0.6%	-2.1%	-1.2%	0.0%	2640.1
		[-2.7%,-0.5%]	[-4.1%,-1.2%]	[-3.6%,-1.0%]	[0.0%,0.0%]	[2615.8,2695.1]
	BS Hosp Pmts	-5.6%	-19.9%	-11.9%	0.0%	369.3
		[-22.4%,-4.4%]	[-34.1%,-12.7%]	[-29.6%,-10.1%]	[0.0%,0.0%]	[347.5,449.3]
	BS Hosp Costs	-0.3%	0.9%	0.0%	0.0%	146.2
		[-0.3%,0.1%]	[0.0%,1.2%]	[-0.1%,0.2%]	[0.0%,0.0%]	[146.1,146.3]
BS Market Share		0.4%	-1.8%	0.2%	0.0%	0.52
		[0.2%,1.7%]	[-2.0%,0.5%]	[-0.2%,1.7%]	[0.0%,0.0%]	[0.51,0.53]
Welfare Δ (\$ per capita)	Consumer	11.7	27.8	19.9	0.0	
		[8.8,50.3]	[17.3,69.2]	[15.4,60.9]	[0.0,0.0]	
	Total	1.0	-11.5	-1.1	0.0	
		[0.5,4.4]	[-12.1,-4.2]	[-3.4,2.0]	[0.0,0.0]	
Number of "Full" Network Markets (out of 12)		6	1	4	12	
		[1,7]	[0,2]	[0,4]	[12,12]	
# Sys Excluded		0.5	2.3	1.2	0.0	
		[0.4,1.3]	[1.8,2.6]	[1.2,1.8]	[0.0,0.0]	
# Sys Excluded (Cond'l on Non-Full)		1.0	2.5	1.8	0.0	
		[1.0,1.4]	[2.1,2.6]	[1.8,2.0]	[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 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 a full network; outcome levels for full network (where all five major hospital systems are included) are presented in right-most column. 95% confidence intervals, reported below all figures (except for whether individual hospital systems are included), 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). For the individual hospital systems, the fraction of bootstrap samples under which individual system members are included are reported beneath predictions.

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 full 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 full in half of our markets (six out of twelve), and when exclusion does occur, only one system is predicted to be excluded. Compared to the full network, Blue Shield’s hospital payments are lower by 6%, premiums fall by approximately \$17 per enrollee, and Blue Shield’s market share increases by 0.4%. Since Blue Shield has lower estimated marginal costs per-enrollee than its competitors (see Table A2), 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 full

network. Furthermore, consumers benefit when hospitals are excluded: consumer surplus increases by \$12 per capita as a result of lower premiums. 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.

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). The Blue Shield optimal network is weakly narrower than the social optimum in all but one HSA, and is full in only four markets.⁴¹ 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 2, 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 full 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) is also not significantly different. On the other hand, Blue Shield’s average hospital payments are significantly lower (by 11.9%, or \$44 per enrollee per year) under its optimal network relative to the full network. This strongly suggests that rate-setting incentives—the focus of our NNTR bargaining model—are the primary motive behind exclusion.⁴² Indeed, in markets where exclusion actually occurs, average rate reductions are larger—18% across markets, and up to 30% 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.⁴³ These insights also explain why

⁴¹The insurer-optimal network under NNTR bargaining is broader than the social-optimal network only in Santa Barbara/Ventura; see Section 7.2.

⁴²We also compute Blue Shield’s realized payments to hospitals under the optimal network reported in column three if its hospital rates were *fixed* to be the same as rates negotiated under the full 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%) on average from the full 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.

⁴³Although we report only reductions in Blue Shield’s premiums, the premiums for other insurers also fall when Blue Shield’s premiums fall.

the network maximizing consumer welfare (column 2 of Table 1) is the most narrow in all markets: it is full 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 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’s network is full and rates negotiated according to NNTR (or, equivalently when networks are full, under 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.⁴⁴ Understandably, 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 vs. 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 4, the Nash-in-Nash solution does not allow hospitals excluded from an insurer’s network to influence negotiated prices for those hospitals included on the network. As a result, an 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

⁴⁴There are potentially additional effects. For example, although we find that premiums fall after exclusion in our setting, if an insurer is able to differentiate itself from broad-network competitors with a narrower network, consumers can be harmed if premiums actually increase as a result of softer insurer competition. This depends on the nature of premium competition (see further discussion in Ho and Lee (2017)). Additionally, reduced hospital payments may affect hospital investment, quality, or entry and exit patterns.

Table 2: Simulation Results for Sacramento

Objective		Social	Consumer	Blue Shield	Full
Surplus (per capita)	BS Profits	0.0%	3.1%	3.1%	316.2
		[0.0%,10.3%]	[1.7%,10.3%]	[1.7%,10.3%]	[290.2,325.9]
	Hospital Profits	0.0%	-26.0%	-26.0%	115.5
		[-40.1%,0.0%]	[-40.1%,-21.3%]	[-40.1%,-21.3%]	[102.2,170.7]
	Total Hosp Costs	0.0%	1.6%	1.6%	98.5
	[0.0%,3.6%]	[1.2%,3.6%]	[1.2%,3.6%]	[96.1,99.4]	
	Total Ins Costs	0.0%	-0.1%	-0.1%	2049.8
		[-0.6%,0.0%]	[-0.6%,0.0%]	[-0.6%,0.0%]	[2032.6,2068.5]
Transfers (per enrollee)	BS Pmts	0.0%	-1.5%	-1.5%	2619.7
		[-3.5%,0.0%]	[-3.5%,-1.1%]	[-3.5%,-1.1%]	[2593.9,2688.7]
	BS Hosp Pmts	0.0%	-16.8%	-16.8%	333.8
		[-30.4%,0.0%]	[-30.4%,-12.9%]	[-30.4%,-12.9%]	[307.4,444.8]
	BS Hosp Costs	0.0%	1.2%	1.2%	165.5
		[0.0%,1.2%]	[1.1%,1.3%]	[1.1%,1.3%]	[165.4,165.7]
Δ Welfare (per capita)	Consumer	0.0	23.3	23.3	
		[0.0,60.1]	[15.7,60.1]	[15.7,60.1]	
	Total	0.0	-3.4	-3.4	
		[0.0,5.0]	[-5.0,5.0]	[-5.0,5.0]	
	BS Market Share	0.0%	0.2%	0.2%	0.53
		[0.0%,2.6%]	[-0.2%,2.6%]	[-0.2%,2.6%]	[0.52,0.54]
Network	# Sys Excluded	0	3	3	
		[0,3]	[3,3]	[3,3]	
	Sys 1 (Sutter)	1	1	1	
		[1,0]	[1,0]	[1,0]	
	Sys 2 (Mercy)	1	1	1	
		[1,0]	[1,0]	[1,0]	
	Sys 3 (UCD)	1	0	0	
		[0,9]	[0,0]	[0,0]	
Sys 4 (Freemont)	1	0	0		
	[0,9]	[0,0]	[0,0]		
Sys 5 (Marshall)	1	0	0		
	[0,9]	[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 a full network; outcome levels for full network (where all five major hospital systems are included) are presented in right-most column. 95% confidence intervals are reported below all figures; see Table 1 for additional details.

exclusion on the part of Blue Shield, we report results under NNTR bargaining for the remainder of our results.

7.2 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 full network, whereas the 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% (\$57 per enrollee per year);

this results in 1.5% (\$39 per enrollee) lower premiums.

However, there are distributional consequences of selective contracting and narrow hospital networks. We illustrate these by providing a map of the Sacramento market in the top panel 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. 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 premiums. However, consumers in zip codes close to excluded hospitals experience welfare reductions. The magnitudes of consumer welfare changes can be fairly large: there are zip codes experiencing reductions as high as \$70 per person per month, and others with increases of up to \$40 per person per month. 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.

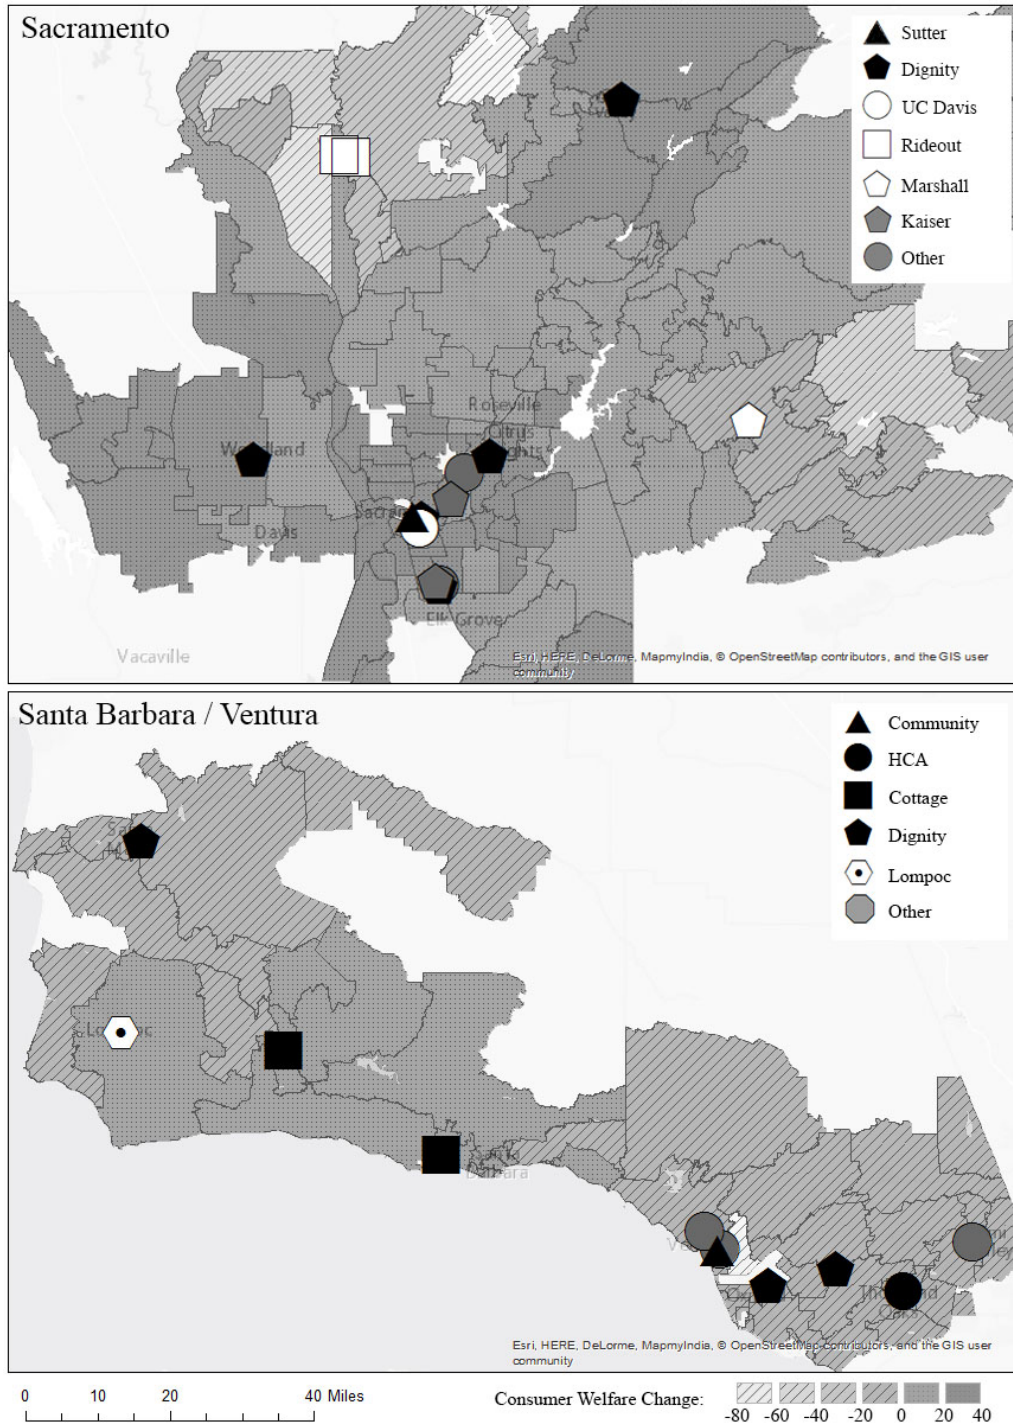
Finally, as discussed earlier, if we instead assumed that hospital rates were determined via Nash-in-Nash bargaining, we would predict that Blue Shield would not wish to exclude any hospital system from its market. This is inconsistent with the observation that a major hospital system in Sacramento was 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.⁴⁵

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 bargaining concept, the Blue Shield optimal network is predicted to be full (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% relative to the full network. However, since such an adjustment reduces Blue Shield's profits, the insurer does not choose to engage in exclusion. Again, as in Sacramento, we predict a very narrow consumer-optimal network.

The second panel 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

⁴⁵The 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 7.3 below for a discussion.

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. Grey shapes indicate other hospitals (Kaiser hospitals and non-Kaiser providers other than the top 5 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.

Table 3: Simulation Results for Santa Barbara / Ventura

Objective		Social	Consumer	Blue Shield	Full
Surplus (per capita)	BS Profits	-0.3%	-5.0%	0.0%	397.7
		[-0.3%,0.1%]	[-5.2%,-0.3%]	[0.0%,0.1%]	[382.9,403.3]
	Hospital Profits	0.0%	-1.5%	0.0%	240.4
		[-1.5%,0.4%]	[-15.3%,0.4%]	[-1.5%,0.0%]	[224.0,299.9]
	Total Hosp Costs	-1.0%	-3.5%	0.0%	115.8
	[-1.0%,-0.9%]	[-3.6%,-1.0%]	[-0.9%,0.0%]	[115.1,116.1]	
	Total Ins Costs	0.0%	0.5%	0.0%	1832.9
		[0.0%,0.0%]	[0.0%,0.6%]	[0.0%,0.0%]	[1815.1,1849.7]
Transfers (per enrollee)	BS Pmts	-0.1%	-0.5%	0.0%	2677.8
		[-0.3%,0.0%]	[-2.5%,0.0%]	[-0.3%,0.0%]	[2646.6,2751.6]
	BS Hosp Pmts	-0.5%	-3.1%	0.0%	363.9
		[-2.0%,-0.2%]	[-17.0%,-0.2%]	[-2.0%,0.0%]	[338.0,459.2]
	BS Hosp Costs	-1.4%	-4.6%	0.0%	126.0
		[-1.4%,-1.4%]	[-4.6%,-1.4%]	[-1.4%,0.0%]	[126.0,126.1]
Δ Welfare (per capita)	Consumer	1.6	7.0	0.0	
		[0.7,7.0]	[0.7,55.7]	[0.0,7.0]	
	Total	0.5	-15.2	0.0	
		[0.4,0.8]	[-15.7,0.5]	[0.0,0.8]	
	BS Market Share	-0.2%	-4.6%	0.0%	0.64
		[-0.2%,-0.1%]	[-4.7%,-0.2%]	[-0.1%,0.0%]	[0.63,0.64]
Network	# Sys Excluded	1	3	0	
		[1,1]	[1,3]	[0,1]	
	Sys 1 (Mercy)	1	1	1	
		[1.0]	[1.0]	[1.0]	
	Sys 2 (Community)	1	1	1	
		[1.0]	[1.0]	[1.0]	
	Sys 3 (Cottage)	1	0	1	
	[1.0]	[0.2]	[1.0]		
Sys 4 (Los Robles)	1	0	1		
	[1.0]	[0.2]	[1.0]		
Sys 5 (Lompoc MC)	0	0	1		
	[0.0]	[0.0]	[0.9]		

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

\$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.

7.3 Fit of the Model

Our information on the hospital systems that Blue Shield proposed to drop in 2005 provides an opportunity to investigate the model’s ability to predict observed exclusion.

We do not expect a perfect fit in terms of the identities of the hospitals included in the proposal because our simulations differ from the true data generating process for Blue Shield’s proposed hospital network in several important ways. 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; furthermore, decisions to drop particular hospital systems that are active in multiple markets may have been made at the state level (e.g., for Sutter). Second, in reality CalPERS and the California Department of Managed Health Care (DMHC) impose external constraints on networks. We know from Blue Shield’s experience in 2005—proposing to drop 38 hospitals from its network the following year, but being permitted to drop only 24—that the DMHC’s approval process can represent a binding

constraint. Such constraints are not explicitly modeled in our analysis. Given these differences, we do not use a direct comparison of our predicted 2004 networks to those observed for 2004, or proposed by Blue Shield for 2005, in order to test the validity of our model. 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.⁴⁶

The NNTR model perfectly predicts the number of excluded systems proposed by Blue Shield in three of the twelve 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 5 systems. Nash-in-Nash predicts a broader network than that proposed in all eleven 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 component 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 DRG-adjusted costs per admission than included hospitals (\$1,973 compared to \$1,840 on average). Our simulations predict average DRG-adjusted costs of \$2,038 for excluded hospitals compared to \$1,794 for other hospitals. Teaching hospitals were also more likely to be excluded, both in the observed 2005 data (18% of excluded hospitals compared to 15% of those included) and in the simulations (20% compared to 17%).

Finally we find evidence that political considerations may have constrained Blue Shield’s proposal. In Los Angeles, for example, the NNTR model predicts that Blue Shield should optimally exclude the University of California, Los Angeles (UCLA) medical center. Given the close links between CalPERS enrollees—i.e., California state and local employees—and the University of California system, this exclusion might well have been politically difficult for CalPERS to approve. In its actual proposal, Blue Shield did not propose to drop the UCLA medical center. It did, however, propose (and was allowed) to exclude the hospital belonging to University of Southern

⁴⁶Appendix Table A1 lists the hospitals 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.

California—the private school “rival” of UCLA.

7.4 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 Blue Cross, 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 Blue Cross PPO plan is not included in the CalPERS choice set. Table A5 reports average results across markets from this specification.

Unsurprisingly, when Blue Cross 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 Table A4, we report average results across markets under the assumption that premiums are fixed at the levels determined when Blue Shield’s hospital network is full. In this case, we predict that the social-optimal network is full in all but one market (Santa Barbara/Ventura, for steering reasons discussed in the previous subsection). 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% lower compared to 11.9% on average), but since it no longer passes along savings to consumers in the form of lower premiums, its profits increase by much more than before (8.8% as opposed to 2.6%).

The consumer optimal network when premiums are fixed no longer tends to be narrow; rather, it is predicted to always be full. 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.

8 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 U.S. health care market developed in Ho and Lee (2017) by endogenizing an insurer’s hospital network and incorporating a new bargaining concept 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 towards 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 Grieco, and Shiran Rachmilevitch.** 2016. “Generalized Insurer Bargaining.” Unpublished.
- Binmore, Ken.** 1987. “Perfect Equilibria in Bargaining Models.” In *The Economics of Bargaining*, ed. Ken Binmore and Partha Dasgupta. Basil Blackwell, Oxford.
- Binmore, Ken, Avner Shaken, and John Sutton.** 1989. “An Outside Option Experiment.” *Quarterly Journal of Economics*, 104(4): 753–770.
- Bolton, Patrick, and Michael D. Whinston.** 1993. “Incomplete Contracts, Vertical Integration, and Supply Assurance.” *Review of Economic Studies*, 60(1): 121–148.
- Capps, Cory, David Dranove, and Mark Satterthwaite.** 2003. “Competition and Market Power in Option Demand Markets.” *RAND Journal of Economics*, 34(4): 737–763.
- Collard-Wexler, Allan, Gautam Gowrisankaran, and Robin S. Lee.** 2016. “Nash-in-Nash” Bargaining: A Microfoundation for Applied Work.” NBER Working Paper 20641.

- Cooper, Zack, Stuart V Craig, Martin Gaynor, and John Van Reenen.** 2015. “The price aint right? Hospital prices and health spending on the privately insured.” National Bureau of Economic Research.
- Crawford, Gregory S., and Ali Yurukoglu.** 2012. “The Welfare Effects of Bundling in Multi-channel Television Markets.” *American Economic Review*, 102(2): 643–685.
- Crawford, Gregory S., Robin S. Lee, Michael D. Whinston, and Ali Yurukoglu.** 2015. “The Welfare Effects of Vertical Integration in Multichannel Television Markets.” Unpublished.
- Cutler, David M., and Sarah J. Reber.** 1998. “Paying for Health Insurance: The Tradeoff between Competition and Adverse Selection.” *Quarterly Journal of Economics*, 113(2): 433–466.
- Dafny, L. S., I. Hendel, and N. Wilson.** 2016. “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: Papers and Proceedings*.
- Draganska, Michaela, Daniel Klapper, and 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.
- Ghili, Soheil.** 2016. “Network Formation and Bargaining in Vertical Markets: The Case of Narrow Networks in Health Insurance.” Unpublished.
- Giovanelli, J., K. Lucia, and S. Corlette.** 2016. “Health Policy Brief: Network Adequacy.” *Health Affairs*.
- Glied, S. A.** 2000. “Health Care Costs: On the Rise Again.” *Journal of Economic Perspectives*, 17(2): 125–148.
- 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): 147–177.
- Gruber, J., and R. McKnight.** 2014. “Controlling Health Care Costs Through Limited Network Insurance Plans: Evidence from Massachusetts State Employees.” *American Economic Journal: Economic Policy*.
- Handel, Benjamin, Igal Hendel, and Michael D. Whinston.** 2015. “Equilibria in Health Exchanges: Adverse Selection Versus Reclassification Risk.” *Econometrica*, 83(4): 1261–1313.
- Ho, Kate.** 2006. “The Welfare Effects of Restricted Hospital Choice in the US Medical Care Market.” *Journal of Applied Econometrics*, 21(7): 1039–1079.
- Ho, Kate, and Robin S. Lee.** 2017. “Insurer Competition in Health Care Markets.” *Econometrica*, 85(2): 379–417.
- Ho, Kate, and Robin S. Lee.** forthcoming. “Narrow Medical Provider Networks: Welfare Implications and Approaches to Market Design.” In *More Fair by Design: Economic Design Responses to Inequality.* , ed. Scott Duke Kominers and Alex Teytelboym. New York:Oxford University Press.

- Horn, Henrick, and Asher Wolinsky.** 1988. “Bilateral Monopolies and Incentives for Merger.” *RAND Journal of Economics*, 19(3): 408–419.
- 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, H.E.** 1979. “Quacks, Lemons, and Licensing: A Theory of Minimum Quality Standards.” *Journal of Political Economy*, 87: 1328–1346.
- Liebman, Eli.** 2016. “Bargaining in Markets with Exclusion: An Analysis of Health Insurance Networks.” Unpublished.
- Manea, Mihai.** forthcoming. “Intermediation and Resale in Networks.” *Journal of Political Economy*.
- Muthoo, Abhinay.** 1999. *Bargaining Theory with Applications*. Cambridge.
- Prager, E.** 2016. “Tiered Hospital Networks, Health Care Demand, and Prices.” Unpublished.
- Ronnen, Uri.** 1991. “Minimum quality standards, fixed costs, and competition.” *RAND Journal of Economics*, 22(4): 490–504.
- Royalty, A., and N. Solomon.** 1998. “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*, 97–109.
- Serrano, Roberto.** 2005. “Fifty Years of the Nash Program, 1953–2003.” *Investigaciones Económicas*, 24(2): 219–258.
- Shapiro, C.** 1983. “Premiums for High-Quality Products as Returns to Reputations.” *Quarterly Journal of Economics*, 98(4): 659–680.
- Shepard, Mark.** 2015. “Hospital Network Competition and Adverse Selection: Evidence from the Massachusetts Health Insurance Exchange.” Unpublished.
- Spence, Michael A.** 1975. “Monopoly, quality, and regulation.” *Bell Journal of Economics*, 6(2): 417–429.
- Stole, Lars A., and Jeffrey Zweibel.** 1996. “Intra-Firm Bargaining under Non-Binding Contracts.” *Review of Economic Studies*, 63(3): 375–410.
- Town, Robert J., and Gregory Vistnes.** 2001. “Hospital Competition in HMO Networks.” *Journal of Health Economics*, 20: 733–753.
- Zaretsky, H., and pmpm Consulting Group Inc.** 2005. “Report of the Analysis of the CalPERS/Blue Shield Narrow Network.” California Department of Managed Health Care.

A Additional Tables

Table A1: Hospitals Proposed to Be Removed from Blue Shield in 2005

Market Name	Hospital Name	System Name	Decision
Central California	Selma Community Hospital		Approved
	Sierra View District Hospital		Denied
	Delano Regional Medical Center		Withdrawn
	Madera Community Hospital		Withdrawn
East Bay	Eden Hospital Medical Center	Sutter	Approved
	Sutter Delta Medical Center	Sutter	Approved
	Washington Hospital		Approved
Inland Counties	Desert Regional Medical Center	Tenet	Approved
Los Angeles	Cedars Sinai Medical Center		Approved
	St. Mary Medical Center	Dignity	Approved
	USC University Hospital	Tenet	Approved
	West Hills Hospital Medical Center		Approved
	Presbyterian Intercommunity Hospital		Denied
	City of Hope National Medical Center		Withdrawn
	St. Francis Memorial Hospital	Verity	Withdrawn
St. Vincent Medical Center	Verity	Withdrawn	
North Bay	Sutter Medical Center of Santa Rosa	Sutter	Approved
	Sutter Warrack Hospital	Sutter	Approved
North San Joaquin	Memorial Hospital Medical Center - Modesto	Sutter	Approved
	Memorial Hospital of Los Banos	Sutter	Approved
	St. Dominics Hospital	Dignity	Approved
	Sutter Tracy Community Hospital	Sutter	Approved
Orange	Hoag Memorial Hospital Presbyterian		Approved
Sacramento	Sutter Davis Hospital	Sutter	Approved
	Sutter General Hospital	Sutter	Approved
	Sutter Memorial Hospital	Sutter	Approved
	Sutter Roseville Medical Center	Sutter	Approved
San Diego	Sharp Chula Vista Medical Center	Sharp	Withdrawn
	Sharp Coronado Hospital and Healthcare Center	Sharp	Withdrawn
	Sharp Grossmont Hospital	Sharp	Withdrawn
	Sharp Mary Birch Hospital for Women	Sharp	Withdrawn
	Sharp Memorial Hospital	Sharp	Withdrawn
Santa Barbara/Ventura	St John's Pleasant Valley Hosp	Dignity	Denied
	St John's Regional Med Center	Dignity	Denied
Santa Clara	OConnor Hospital	Verity	Approved
West Bay	California Pacific Medical Center Campus Hospital	Sutter	Approved
	Seton Medical Center	Verity	Approved
	St. Lukes Hospital	Sutter	Approved

Notes: List of hospitals that Blue Shield proposed to exclude in its filing to the California Department of Managed Health Care (DMHC) for the 2005 year. Source: DMHC “Report on the Analysis of the CalPERS/Blue Shield Narrow Network” (Zaretsky and pmpm Consulting Group Inc. (2005)). “Market name” denotes the Health Service Area of the relevant hospital; the two HSAs in California that are not listed here did not contain hospitals that Blue Shield proposed to exclude. “Decision” is the eventual outcome of the proposal for the relevant hospital.

Table A2: Summary Statistics and Parameter Estimates

		Blue Shield	Blue Cross	Kaiser
Premiums (per year)	Single	3782.64	4192.92	3665.04
	2 party	7565.28	8385.84	7330.08
	Family	9834.84	10901.64	9529.08
Hospital Network	# Hospitals in network	189	223	27
	# Hospital systems in network	119	149	-
	Avg. hospital price per admission	6624.08 (3801.24)	5869.26 (2321.57)	-
	Avg. hospital cost per admission	1693.47 (552.17)	1731.44 (621.33)	-
Household Enrollment	Single	19313	8254	20319
	2 party	16376	7199	15903
	Family	35058	11170	29127
	Avg # individuals per family	3.97	3.99	3.94
Parameter Estimates	η (Non-inpatient cost per enrollee)	1691.50 (10.41)	1948.61 (8.14)	2535.14 (0.62)
	τ^H (Hospital bargaining weight)	0.31 (0.05)	0.38 (0.03)	-
(Ho and Lee, 2017)	τ^ϕ (Premium bargaining weight)		0.47 (0.00)	

Notes: The first three panels report summary statistics by insurer. The number of hospitals and hospital systems for Blue Shield and Blue Cross are determined by the number of in-network hospitals or systems with at least 10 admissions observed in the data. Hospital prices and costs per admission are averages of unit-DRG amounts, unweighted across hospitals (with standard deviations reported in parentheses). The fourth panel reports estimates from Ho and Lee (2017) of marginal costs for each insurer (which do not include hospital payments for Blue Shield and Blue Cross), and (insurer-specific) hospital price and (non-insurer specific) premium Nash bargaining weights; standard errors are reported in parentheses.

Table A3: Admission Probabilities and DRG Weights

Age/Sex	Admission Probabilities		DRG Weights		
	BS	BC	BS	BC	All
0-19 Male	1.78%	2.08%	1.78	1.49	1.70
20-34 Male	1.66%	2.07%	1.99	1.77	1.92
35-44 Male	2.79%	3.21%	1.95	1.89	1.93
45-54 Male	5.29%	5.32%	2.07	2.05	2.07
55-64 Male	10.13%	9.70%	2.25	2.25	2.25
0-19 Female	1.95%	2.04%	1.31	1.39	1.32
20-34 Female	11.75%	10.22%	0.84	0.87	0.85
35-44 Female	7.31%	7.73%	1.32	1.33	1.32
45-54 Female	6.16%	6.82%	1.90	1.83	1.87
55-64 Female	9.01%	9.26%	2.03	2.02	2.03

Notes: Average admission probabilities and DRG weights per admission by age-sex category.

Table A4: Simulation Results for All Markets (Averages), Fixed Premiums

Objective		Social	Consumer	Blue Shield		Full
		(NNTR)	(NNTR)	(NNTR)	(NN)	(NNTR/NN)
Surplus (\$ per capita)	BS Profits	0.0%	0.0%	8.8%	0.0%	304.7
		[0.0%,0.1%]	[0.0%,0.0%]	[5.0%,18.2%]	[0.0%,0.0%]	[287.5,312.1]
	Hospital Profits	0.0%	0.0%	-23.1%	0.0%	170.0
		[-0.1%,0.0%]	[0.0%,0.0%]	[-34.0%,-16.5%]	[0.0%,0.0%]	[159.4,209.4]
	Total Hosp Costs	-0.1%	0.0%	-1.1%	0.0%	95.6
	[-0.1%,0.1%]	[0.0%,0.0%]	[-1.7%,-1.0%]	[0.0%,0.0%]	[94.1,96.3]	
	Total Ins Costs	0.0%	0.0%	0.6%	0.0%	2008.5
		[0.0%,0.0%]	[0.0%,0.0%]	[0.5%,0.8%]	[0.0%,0.0%]	[1990.4,2025.7]
Transfer / Cost (\$ per enrollee)	BS Premiums	0.0%	0.0%	0.3%	0.0%	2640.1
		[0.0%,0.0%]	[0.0%,0.0%]	[0.3%,0.4%]	[0.0%,0.0%]	[2615.8,2695.1]
	BS Hosp Pmts	0.0%	0.0%	-18.8%	0.0%	369.3
		[-0.1%,0.0%]	[0.0%,0.0%]	[-29.4%,-13.2%]	[0.0%,0.0%]	[347.5,449.3]
	BS Hosp Costs	-0.1%	0.0%	1.3%	0.0%	146.2
		[-0.1%,0.1%]	[0.0%,0.0%]	[1.1%,1.4%]	[0.0%,0.0%]	[146.1,146.3]
BS Market Share		0.0%	0.0%	-3.7%	0.0%	0.52
		[0.0%,0.0%]	[0.0%,0.0%]	[-5.4%,-3.7%]	[0.0%,0.0%]	[0.51,0.53]
Welfare Δ (\$ per capita)	Consumer	-0.1	0.0	-8.7	0.0	
		[-0.1,-0.1]	[0.0,0.0]	[-12.1,-8.6]	[0.0,0.0]	
	Total	0.0	0.0	-16.2	0.0	
		[0.0,0.0]	[0.0,0.0]	[-23.4,-16.0]	[0.0,0.0]	
Number of "Full" Network Markets (out of 12)		11	12	0	12	
		[11,11]	[12,12]	[0,2]	[12,12]	
# Sys Excluded		0.1	0.0	1.9	0.0	
		[0.1,0.1]	[0.0,0.0]	[1.8,2.5]	[0.0,0.0]	
# Sys Excluded (Cond'l on Non-Full)		1.0	0.0	1.9	0.0	
		[1.0,1.0]	[0.0,0.0]	[1.9,2.5]	[0.0,0.0]	

Notes: Unweighted averages across markets when premiums are fixed to be the same as when Blue Shield's network is full. See Table 1 for details.

Table A5: Simulation Results for All Markets (Averages), No Blue Cross

Objective		Social	Consumer	Blue Shield		Full
		(NNTR)	(NNTR)	(NNTR)	(NN)	(NNTR/NN)
Surplus (\$ per capita)	BS Profits	1.1%	3.2%	3.5%	0.0%	365.8
		[0.4%,3.0%]	[1.9%,8.5%]	[2.3%,8.9%]	[0.0%,0.0%]	[344.9,375.9]
	Hospital Profits	-6.7%	-32.6%	-27.2%	0.1%	118.6
		[-13.6%,-1.6%]	[-50.9%,-27.1%]	[-48.4%,-20.0%]	[-0.3%,0.1%]	[107.9,158.5]
	Total Hosp Costs	-0.2%	-0.1%	0.0%	-0.1%	89.3
	[-0.4%,0.4%]	[-0.6%,0.3%]	[-0.3%,0.4%]	[-0.1%,0.0%]	[88.0,89.9]	
	Total Ins Costs	0.0%	0.0%	0.0%	0.0%	2005.6
		[-0.2%,0.0%]	[0.0%,0.2%]	[-0.1%,0.1%]	[0.0%,0.0%]	[1988.8,2023.9]
Transfer / Cost (\$ per enrollee)	BS Premiums	-0.3%	-1.5%	-1.1%	0.0%	2603.1
		[-0.8%,0.1%]	[-3.1%,-1.1%]	[-2.9%,-0.8%]	[0.0%,0.0%]	[2584.3,2643.2]
	BS Hosp Pmts	-3.6%	-17.4%	-15.0%	0.0%	336.4
		[-8.4%,-1.3%]	[-31.8%,-13.8%]	[-30.3%,-10.9%]	[-0.2%,0.0%]	[318.1,404.4]
	BS Hosp Costs	-0.4%	0.3%	-0.1%	-0.1%	146.5
		[-0.5%,0.3%]	[0.2%,0.4%]	[-0.2%,0.1%]	[-0.1%,0.0%]	[146.5,146.6]
BS Market Share		0.2%	-0.3%	0.0%	0.0%	0.63
		[0.1%,0.7%]	[-0.8%,0.1%]	[-0.4%,0.5%]	[0.0%,0.0%]	[0.62,0.63]
Welfare Δ (\$ per capita)	Consumer	5.4	20.7	15.8	0.0	
		[1.9,14.7]	[14.1,47.8]	[10.3,45.3]	[-0.3,0.0]	
	Total	0.6	-7.5	-6.8	0.1	
		[0.4,1.9]	[-14.7,-7.3]	[-9.9,-5.1]	[-0.6,0.1]	
Number of "Full" Network Markets (out of 12)		7	2	2	11	
		[6,9]	[0,3]	[0,4]	[11,12]	
# Sys Excluded		0.4	1.7	1.7	0.1	
		[0.3,0.7]	[1.6,2.4]	[1.3,2.3]	[0.0,0.1]	
# Sys Excluded (Cond'l on Non-Full)		1.0	2.0	2.0	1.0	
		[1.0,1.3]	[2.0,2.4]	[1.9,2.3]	[0.0,1.0]	

Notes: Unweighted averages across markets when Blue Cross is unavailable. See Table 1 for details.

B Proofs

For the following proofs, we introduce the following notation:

- $\mathbf{p}_{(ij=0)} \equiv \{0, \mathbf{p}_{-ij}\}$ (i.e., $\mathbf{p}_{(ij=0)}$ replaces $p_{ij} = 0$ in the vector of prices \mathbf{p}). Note that $[\Delta_{ij}\Pi_{ij}(G, \mathbf{p})] = [\Delta_{ij}\Pi_{ij}(G, \mathbf{p}_{(ij=0)})] \forall i \in G, \forall \mathbf{p}$: i.e., the payment made between MCO j and hospital i cancels out when considering changes in the joint surplus $[\Delta_{ij}\Pi_{ij}(G, \mathbf{p})]$.
- For each $G \in \mathcal{G}_j$ and hospital $i \in G$, $N_i(G) \equiv (\mathcal{H} \setminus G) \cup i$ (i.e., i and all hospitals not in G).
- $v_h^i(G, \mathbf{p}) \equiv [\Delta_{hj}\Pi_{hj}((G \setminus i) \cup h, \mathbf{p}_{-ij})]$ denotes the bilateral gains from trade created by MCO j and hospital h if i is replaced by h in network G .
- Let $v_{(1)}^i(\cdot)$ and $v_{(2)}^i(\cdot)$ represent the first and second-highest values in the set $\mathbf{v}^i(\cdot) \equiv \{v_h^i(\cdot)\}_{h \in N_i(G)}$, and $k_{(1)}^i(\cdot)$ and $k_{(2)}^i(\cdot)$ their respective indices.⁴⁷ In the case of lump-sum transfers (for Proposition 4.4 and Proposition 4.5), we omit the argument \mathbf{p} as linear prices are assumed to be 0 (and lump-sum transfers do not otherwise affect bilateral gains from trade).

Lemma B.1. *Fix G and \mathbf{p} . For any $i \in G$, if $k = \arg \max_{h \in N_i \setminus i} [\Delta_{hj}\Pi_{hj}((G \setminus i) \cup h, \mathbf{p}_{-ij})]$, then:*

$$\tau[\Delta_{ij}\Pi_{ij}(G, \mathbf{p}_{(ij=0)})] \geq [\Delta_{kj}\Pi_{kj}((G \setminus i) \cup k, \mathbf{p}_{-ij})] \iff p_{ij}^{Nash}(G, \mathbf{p}_{-ij}) \leq p_{ij}^{OO}(G, \mathbf{p}_{-ij})$$

Proof. Nash-in-Nash prices are given by the solution to (2):

$$p_{ij}^{Nash}(G, \mathbf{p}_{-ij}) D_{ij}^H(G) = (1 - \tau)[\Delta_{ij}\pi_j^M(G, \mathbf{p}_{(ij=0)})] - \tau[\Delta_{ij}\pi_i^H(G, \mathbf{p}_{(ij=0)})] \forall i \in G. \quad (11)$$

Reservation prices for any $h \notin G$ replacing i are given by the solution to (4):

$$p_{hj}^{res}(G \setminus i, \mathbf{p}_{-ij}) D_{hj}^H(G \setminus i) = -[\Delta_{hj}\pi_h^H((G \setminus i) \cup h, \mathbf{p}_{-ij})].$$

Substituting this expression for p_{hj}^{res} into the right hand side of (3) yields: $\pi_j^M((G \setminus i) \cup h, \{p_{hj}^{res}(G \setminus i, \mathbf{p}_{-ij}), \mathbf{p}_{-ij}\}) = \pi_j^M((G \setminus i) \cup h, \mathbf{p}_{-ij}) + [\Delta_{hj}\pi_h^H((G \setminus i) \cup h, \mathbf{p}_{-ij})]$; hence, k also maximizes the right-hand-side of (3) over $h \notin G$.

Re-arranging (3) yields:

$$\begin{aligned} p_{ij}^{OO}(G) D_{ij}^H(G) &= \pi_j^M(G, \mathbf{p}_{(ij=0)}) - \pi_j^M((G \setminus i) \cup k, \mathbf{p}_{-ij}) - [\Delta_{kj}\pi_k^H((G \setminus i) \cup k, \mathbf{p}_{-ij})] \\ &= [\Delta_{ij}\pi_j^M(G, \mathbf{p}_{(ij=0)})] - [\Delta_{kj}\pi_j^M((G \setminus i) \cup k, \mathbf{p}_{-ij})] - [\Delta_{kj}\pi_k^H((G \setminus i) \cup k, \mathbf{p}_{-ij})] \\ &= [\Delta_{ij}\pi_j^M(G, \mathbf{p}_{(ij=0)})] - [\Delta_{kj}\Pi_{kj}((G \setminus i) \cup k, \mathbf{p}_{-ij})]. \end{aligned} \quad (12)$$

Finally, note that:

$$\begin{aligned} &\tau[\Delta_{ij}\Pi_{ij}(G, \mathbf{p}_{(ij=0)})] \geq [\Delta_{kj}\Pi_{kj}((G \setminus i) \cup k, \mathbf{p}_{-ij})] \\ (\iff) &\quad -[\Delta_{kj}\Pi_{kj}((G \setminus i) \cup k, \mathbf{p}_{-ij})] \geq -\tau[\Delta_{ij}\Pi_{ij}(G, \mathbf{p}_{(ij=0)})] \\ (\iff) &\Delta_{ij}\pi_j^M(G, \mathbf{p}_{(ij=0)}) - [\Delta_{kj}\Pi_{kj}((G \setminus i) \cup k, \mathbf{p}_{-ij})] \geq (1 - \tau)[\Delta_{ij}\pi_j^M(G, \mathbf{p}_{(ij=0)})] - \tau[\Delta_{ij}\pi_i^H(G, \mathbf{p}_{(ij=0)})] \\ (\iff) &\quad p_{ij}^{OO}(G, \mathbf{p}_{-ij}) \geq p_{ij}^{Nash}(G, \mathbf{p}_{-ij}) \end{aligned}$$

where the last line follows from substituting in the expressions from (11) and (12) and dividing through by $D_{ij}^H(G)$. \square

B.1 Proof of Proposition 4.1

Fix G and prices for other MCOs $-j$, \mathbf{p}_{-j} , and omit them as arguments in subsequent notation. Let H denote the number of hospitals that MCO j contracts with in G . Define the mapping $\rho : [-\bar{p}, \bar{p}]^H \rightarrow [-\bar{p}, \bar{p}]^H$

⁴⁷ For our analysis, we assume that all values $\{v_h^i(\cdot)\}$ are distinct so that there are unique values of i and k . The analysis in Manea (forthcoming) allows for equal values.

where, for each $i \in G$:

$$\rho_i(\{p_{hj}\}_{h \in G}) = \max \left\{ -\bar{p}, \min\{\rho_i^{Nash}(\{p_{hj}\}_{h \neq i, h \in G}), \rho_i^{OO}(\{p_{hj}\}_{h \neq i, h \in G}), \bar{p}\} \right\} \quad (13)$$

$$\rho_i^{Nash}(\{p_{hj}\}_{h \neq i, h \in G}) = \left((1 - \tau)[\Delta_{ij}\pi_j^M(\mathbf{p}_{(ij=0)})] - \tau[\Delta_{ij}\pi_j^H(\mathbf{p}_{(ij=0)})] \right) / D_{ij}^H \quad (14)$$

$$\rho_i^{OO}(\{p_{hj}\}_{h \neq i, h \in G}) = \left([\Delta_{ij}\pi_j^M(\mathbf{p}_{(ij=0)})] - \max_{k \in \mathcal{H} \setminus G} [\Delta_{kj}\Pi_{kj}((G \setminus i) \cup k, \{p_{hj}\}_{h \neq i, h \in G})] \right) / D_{ij}^H \quad (15)$$

Given our assumptions on firm profit functions (which are linear in prices), (14) and (15) are continuous in $\{p_{hj}\}_{h \neq i, h \in G}$ for all $i \in G$, and thus $\rho_i(\cdot)$ is a continuous mapping from a compact convex set into itself. By Brouwer's fixed-point theorem, there exists a fixed point of $\rho(\cdot)$. It is straightforward to show that any fixed point of $\rho(\cdot)$ satisfies (2)-(5) (as (14) follows from (11) and (15) from (12)), and thus represents a vector of NNTR prices for network G and other MCO prices \mathbf{p}_{-j} .

B.2 Proof of Proposition 4.2

Assume first G is stable. Proceed by contradiction, and assume that $[\Delta_{hj}\Pi_{hj}((G \setminus i) \cup h, \mathbf{p}_{-ij}^*)] > [\Delta_{ij}\Pi_{ij}(G, \mathbf{p}^*)]$ for some agreement $i \in G$ and $h \in (N_i \setminus i) \cup \emptyset$. If this holds for $h = \emptyset$, then agreement i is unstable since there are negative bilateral gains from trade; contradiction. Thus, it must be that if $k = \arg \max_{h \in N_i \setminus i} [\Delta_{hj}\Pi_{hj}((G \setminus i) \cup h, \mathbf{p}_{-ij}^*)]$, $[\Delta_{kj}\Pi_{kj}((G \setminus i) \cup k, \mathbf{p}_{-ij}^*)] > \tau[\Delta_{ij}\Pi_{ij}(G, \mathbf{p}^*)]$ (since $\tau < 1$). By Lemma B.1, $p_{ij}^{OO}(G) \leq p_{ij}^{Nash}(G)$, and $p_{ij}^*(G) = p_{ij}^{OO}(G)$. However, at this payment, hospital i receives:

$$\begin{aligned} \pi_i^H(G, \mathbf{p}_{(ij=0)}^*) + p_{ij}(G)D_{ij}^H(G) &= \pi_i^H(G, \mathbf{p}_{(ij=0)}^*) + \underbrace{[\Delta_{ij}\pi_j^M(G, \mathbf{p}_{(ij=0)}^*)] - [\Delta_{kj}\Pi_{kj}((G \setminus i) \cup k, \mathbf{p}_{(ij=kj=0)}^*)]}_{\text{From (12)}} \\ &= \pi_i^H(G \setminus i, \mathbf{p}_{(ij=0)}^*) + \underbrace{[\Delta_{ij}\Pi_{ij}(G, \mathbf{p}_{(ij=0)}^*)] - [\Delta_{kj}\Pi_{kj}((G \setminus i) \cup k, \mathbf{p}_{-ij}^*)]}_{< 0 \text{ by assumption (since } [\Delta_{ij}\Pi_{ij}(G, \mathbf{p}^*)] = [\Delta_{ij}\Pi_{ij}(G, \mathbf{p}_{(ij=0)}^*)])} \\ &\quad (16) \end{aligned}$$

and hospital i would prefer rejecting the payment $p_{ij}^*(G)$. Thus G is not stable, yielding a contradiction.

Next, assume that $[\Delta_{ij}\Pi_{ij}(G, \mathbf{p}^*)] \geq [\Delta_{kj}\Pi_{kj}((G \setminus i) \cup k, \mathbf{p}^*)] \forall k \in (\mathcal{H} \setminus G) \cup \emptyset \forall i \in G$. We now prove that this implies G is stable. Assume by contradiction that some agreement $i \in G$ is not stable at \mathbf{p}^* . If $p_{ij}^* = p_{ij}^{Nash}$, then agreement i is unstable only if $[\Delta_{ij}\Pi_{ij}(G, \mathbf{p}^*)] < 0$: contradiction. If $p_{ij}^* = p_{ij}^{OO}$, then by Lemma B.1, $\tau[\Delta_{ij}\Pi_{ij}(G, \mathbf{p}_{(ij=0)}^*)] < [\Delta_{kj}\Pi_{kj}((G \setminus i) \cup k, \mathbf{p}_{-ij}^*)]$ for $k = \arg \max_{h \in N_i \setminus i} [\Delta_{hj}\Pi_{hj}(G \setminus i \cup h, \mathbf{p}_{-ij}^*)]$. By (16), such an agreement will be unstable at such prices only if $[\Delta_{ij}\Pi_{ij}(G, \mathbf{p}_{(ij=0)}^*)] < [\Delta_{kj}\Pi_{kj}((G \setminus i) \cup k, \mathbf{p}_{-ij}^*)]$; contradiction. Thus, G is stable at \mathbf{p}^* .

B.3 Proof of Propositions 4.3-4.4

In the proofs for these propositions and for Proposition 4.5, we restrict attention to lump-sum transfers negotiated between MCO j and each hospital; i.e., total payments are made when an agreement is formed.^{48,49}

We derive the equivalent lump-sum NNTR prices $P_{ij}^*(G) \equiv \min\{P_{ij}^{Nash}(\cdot), P_{ij}^{OO}\}$, where (using (11) and (12)):

$$P_{ij}^{Nash}(G) = (1 - \tau)[\Delta_{ij}\pi_j^M(G)] - \tau[\Delta_{ij}\pi_i^H(G)], \quad (17)$$

$$P_{ij}^{OO}(G) = [\Delta_{ij}\pi_j^M(G)] - [\Delta_{kj}\Pi_{kj}((G \setminus i) \cup k)]. \quad (18)$$

⁴⁸Transfers between hospitals and other MCOs $-j$ are still allowed to be linear; however, we omit these prices \mathbf{p}_{-j} from notation as they remain fixed in our analysis and changes they induce in payments are otherwise subsumed into hospital profits as a function of MCO j 's network.

⁴⁹We restrict attention to lump-sum transfers for analytic tractability. Using linear fees may imply that flow payoffs that accrue to each firm will depend on the set of prices that have previously been agreed upon, which significantly complicates analysis.

Note that these values for each pair $ij \in G$ depend only on profit terms, which are assumed to be primitives; thus, $P_{ij}^*(G) \forall ij \in G$ exists and is unique and Proposition 4.3 is proven.

In this setting, Lemma B.1 implies that for any G and $i \in G$, if $k = \arg \max_{h \in N_i \setminus i} [\Delta_{hj} \Pi_{hj}((G \setminus i) \cup h, \mathbf{p}_{-ij})]$ then:

$$\tau[\Delta_{ij} \Pi_{ij}(G)] \geq [\Delta_{kj} \Pi_{kj}((G \setminus i) \cup k)] \iff P_{ij}^{Nash}(G) \leq P_{ij}^{OO}(G).$$

In turn, Proposition 4.2 implies that if G is stable, then:

$$[\Delta_{ij} \Pi_{ij}(G)] \geq [\Delta_{kj} \Pi_{kj}((G \setminus i) \cup k)] \forall k \in (\mathcal{H} \setminus G) \cup \emptyset,$$

where all bilateral surpluses can be expressed as a function of the network only (as lump sum transfers cancel out).

Single hospital announced at period-0. We first prove the conditions of Proposition 4.4 hold for subgames where the period-0 network contains a single hospital; we defer establishing existence until the more general multiple hospital case.

Consider any subgame where stable network G is announced in period 0 by MCO j , $G \equiv \{i\}$ (i.e., G contains a single hospital i), and no agreement has yet been formed by MCO j . Relative to no agreement—where each period MCO j receives $(1-\delta)\pi_j^M(\{\emptyset\})$ and each hospital i receives $(1-\delta)\pi_i^H(\{\emptyset\})$ —an agreement with hospital i results in an increase in total discounted profits of $(1-\delta)[\Delta_{ij}\pi_j^M(G) + \Delta_{ij}\pi_i^H(G)]/(1-\delta) = [\Delta_{ij}\Pi_{ij}(G)]$ for MCO j and hospital i .

By Proposition 4.2, it must be that $i = k_{(1)}^i(G)$, else G is not stable. Let $k = k_{(2)}^i(G)$ so that $[\Delta_{ij}\Pi_{ij}(G)] > [\Delta_{kj}\Pi_{kj}((G \setminus i) \cup k)]$. This subgame corresponds exactly to the single seller and multiple buyer case analyzed in Manea (forthcoming), where the MCO j can transact with any hospital $h \in \mathcal{H}$ and generate surplus $v_h^i(\{i\}) = [\Delta_{hj}\Pi_{hj}(\{h\})]$. A direct application of Proposition 1 of Manea (forthcoming) implies that all MPE of this subgame are outcome equivalent, and for any family of MPE (i.e., a collection of MPE for different values of δ), expected payoffs for MCO j (above its disagreement point) converge as $\delta \rightarrow 1$ to $\max(\tau[\Delta_{ij}\Pi_{ij}(G)], [\Delta_{kj}\Pi_{kj}((G \setminus i) \cup k)])$. Furthermore, there exists $\underline{\delta}$ such that for $\delta > \underline{\delta}$, if $\tau[\Delta_{ij}\Pi_{ij}(G)] > [\Delta_{kj}\Pi_{kj}((G \setminus i) \cup k)]$, trade occurs only with hospital i ; otherwise, the MCO engages with positive probability with either i or k , but the probability that the MCO comes to agreement with hospital i converges to 1 as $\delta \rightarrow 1$.

To show that this result implies that negotiated prices converge to NNTR prices, consider the following two cases:

1. $\tau[\Delta_{ij}\Pi_{ij}(G, \mathbf{p})] > [\Delta_{kj}\Pi_{kj}((G \setminus i) \cup k)]$. MPE expected payoffs (above its disagreement point) for the MCO then converge to:

$$\begin{aligned} \tau[\Delta_{ij}\Pi_{ij}(G)] &= \tau[\Delta_{ij}\pi_j^M(G, \mathbf{p})] + \tau[\Delta_{ij}\pi_i^H(G)] \\ &= \Delta_{ij}\pi_j^M(G) - \left((1-\tau)[\Delta_{ij}\pi_j^M(G)] - \tau[\Delta_{ij}\pi_i^H(G)] \right) \\ &= [\Delta_{ij}\pi_j^M(G)] - P_{ij}^{Nash}(G) \end{aligned}$$

where the last line follows from (17). By Lemma B.1, $P_{ij}^*(\cdot) = P_{ij}^{Nash}(\cdot)$.

2. $\tau[\Delta_{ij}\Pi_{ij}(G)] \leq [\Delta_{kj}\Pi_{kj}((G \setminus i) \cup k)]$. MPE expected payoffs for the MCO then converge to:

$$[\Delta_{kj}\Pi_{kj}((G \setminus i) \cup k)] = [\Delta_{ij}\pi_j^M(G)] - P_{ij}^{OO}(G)$$

where the equality follows from (18). By Lemma B.1, $P_{ij}^*(\cdot) = P_{ij}^{OO}(\cdot)$.

Thus, for the payoffs for MCO j to converge to $\max(\tau[\Delta_{ij}\Pi_{ij}(G)], [\Delta_{kj}\Pi_{kj}((G \setminus i) \cup k)])$, equilibrium payments must converge to P_{ij}^* .

We now provide insight into MPE outcomes and strategies of the subgame where network $G = \{i\}$ is announced. Since G is stable, $v_{(1)}^i(G) > 0$. For sufficiently high δ , arguments used in Proposition 1 of Manea

(forthcoming) show that any MPE of this subgame results in immediate agreement with any hospital with which the MCO engages with probability greater than 0; and that all MPE are characterized by the following conditions:

$$\begin{aligned} u_0^i &= \tau(v_{(1)}^i(G) - \delta u_{(1)}^i) + (1 - \tau)\delta u_0^i \\ u_h^i &= \Lambda_h^i(\tau\delta u_h^i + (1 - \tau)(v_h^i(G) - \delta u_0^i)) \quad \forall h \in N_i \\ \Lambda_h^i &= \begin{cases} \frac{1 - \delta + \delta\tau}{\delta\tau} - \frac{(1 - \delta)(1 - \tau)v_{(h)}^i(G)}{\delta\tau(v_{(h)}^i(G) - u_0^i)} & \text{if } u_0^i < v_h^i(G) \frac{\tau}{1 - \delta + \delta\tau} \\ 0 & \text{otherwise} \end{cases} \quad \forall h \in N_i \end{aligned}$$

where u_0^i and u_h^i are the expected payoffs for MCO j and hospital h , and Λ_h^i is the probability that the MCO engages with h at the beginning of each period where agreement has not yet occurred. Furthermore, only the two highest surplus creating hospitals, $k_1 = k_{(1)}^i(G)$ and $k_2 = k_{(2)}^i(G)$, have positive probabilities of being engaged with in any period. In any MPE, Manea proves that there exists a unique value of u_0 such that $\Lambda_{k_1}^i + \Lambda_{k_2}^i = 1$; this pins down all equilibrium outcomes, and MPE strategies that generate these payoffs and probabilities are easily constructed. Furthermore, $\Lambda_{k_1}^i \rightarrow 1$ as $\delta \rightarrow 1$, and if $\tau v_{k_1}^i \geq v_{k_2}^i$, then $\Lambda_{k_1}^i = 1$ for sufficiently high δ .

Multiple hospitals announced at period-0. We next examine subgames where stable network G is announced in period 0 by MCO j , G contains more than one hospital, and no agreements have yet been formed by MCO j .

Consider the bargain being conducted by representative r_i , $i \in G$, holding fixed its beliefs over the outcomes of other negotiations. Let $\mathbf{\Lambda}^h \equiv \{\Lambda_k^h\}_{k \in (\mathcal{H} \setminus i) \cup \emptyset}$ represent the perceived probabilities held by r_i and all hospitals representatives contained in N_i over whether another MCO representatives r_h , $h \in G \setminus i$, forms agreements with other hospitals $k \neq i$. Denote by $\mathbf{\Lambda}^{-i} \equiv \{\mathbf{\Lambda}^h\}_{h \in G \setminus i}$; this induces a distribution $f(\tilde{G} | \mathbf{\Lambda}^{-i})$ over the set of all other networks not involving i that may form, $\tilde{G} \subseteq \mathcal{H} \setminus i$. Let $\tilde{v}_h^i(\mathbf{\Lambda}^{-i}) \equiv \sum_{\tilde{G} \subseteq \mathcal{H} \setminus i} [\Delta_{hj} \Pi_{hj}(\tilde{G} \cup h)] \times f(\tilde{G} | \mathbf{\Lambda}^{-i})$ represent the expected bilateral gains from trade created when r_i and h come to an agreement given beliefs $\mathbf{\Lambda}^{-i}$.⁵⁰ We establish the following result:

Lemma B.2. *For any $\varepsilon_1, \varepsilon_2 > 0$, there exists $\underline{\Lambda} < 1$ and $\underline{\delta} > 0$ such that if $\Lambda_h^i > \underline{\Lambda} \quad \forall h \in G \setminus i$ and $\delta > \underline{\delta}$, any MPE involves r_i coming to agreement with hospital i with probability greater than $1 - \varepsilon_1$ and payoffs are within ε_2 of $\max(\tau[\Delta_{ij} \Pi_{ij}(G)], [\Delta_{kj} \Pi_{kj}((G \setminus i) \cup k)]) = \max(\tau v_{(1)}^i(G), v_{(2)}^i(G))$, where $k = k_{(2)}^i(G)$.*

Proof. In this setting, any representative r_i is engaged in the same bargaining protocol with hospitals $h \in N_i$ as before, but now expects to generate surplus $\tilde{v}_h^i(\cdot)$ upon agreement with any hospital. For sufficiently high $\underline{\Lambda}$ (so that the probability $f(G \setminus i | \mathbf{\Lambda}^{-i}) > \underline{\Lambda}^{|g|-1}$ is close to 1, where $|g|$ represents the number of agreements in G), $\tilde{v}_h^i(\cdot)$ can be made to be arbitrarily close to $v_h^i(G)$, and the indices for the first and second-highest values over $\tilde{v}_h^i(\cdot)$ will coincide with the indices for the first and second-highest values over $v_h^i(G)$.⁵¹ As before, applying the results from Proposition 1 of Manea (forthcoming), shows that payoffs in any MPE must converge to $\max(\tau \tilde{v}_{(1)}^i(\cdot), \tilde{v}_{(2)}^i(\cdot))$ and the probability that r_i engages and comes to agreement with $k_{(1)}^i(G)$, given by $\Lambda_{(1)}^i$, converges to 1 as $\delta \rightarrow 1$. Furthermore, for large enough $\underline{\Lambda}$, payoffs will converge to be within ε_2 of $\max(\tau v_{(1)}^i(G), v_{(2)}^i(G))$; by the arguments of the single-hospital case, this also ensures that payments are within ε_2 of NNTR prices. Finally, since G is assumed to be stable, by Proposition 4.2, $i = k_{(1)}^i(G)$ and the result follows. \square

We now prove that there exists an MPE of our game for sufficiently high δ . We adapt the proof of Proposition 4 of Manea (forthcoming); following his arguments, MPE payoffs and probabilities of engagement

⁵⁰ Implicit in this construction is the possibility that r_i may negotiate with some hospital $k \in \tilde{G}$, $k \notin G$, and that the representative from k may have some expectation that an agreement may form between a different representative for k and another representative for MCO j (r_h , $h \neq i$). This can occur if, as discussed in footnote 27, both r_i and r_h negotiate with k that neither representative was initially assigned to engage with ($k \neq i, h$). This is consistent with our assumption that such a hospital k also employs separate representatives to engage with each separate MCO representative, and must act without knowledge of other representatives' actions.

⁵¹This follows since profits are assumed to be finite for any potential network.

for each r_i , $i \in G$, and its bargaining partners must satisfy:

$$u_0^i = \sum_{h \in N_i} \Lambda_h^i (\tau(\tilde{v}_h^i(\mathbf{\Lambda}^{-i}) - \delta u_h^i) + (1 - \tau)\delta u_0^i) \quad (19)$$

$$u_h^i = \Lambda_h^i (\tau \delta u_h^i + (1 - \tau)(\tilde{v}_h^i(\mathbf{\Lambda}^{-i}) - \delta u_0^i)) \quad \forall h \in N_i \quad (20)$$

where u_0^i is the expected payoff created for the MCO by representative r_i , u_h^i is the expected payoff for the hospital $k_h^i(G)$, and Λ_h^i is the probability that the MCO engages with either hospital in the beginning of a period. Again, all expected payoffs are above what would occur if no agreement by r_i and any of the hospitals were reached.

For any arbitrary vector $\mathbf{\Lambda}^i$ describing a probability distribution over N_i (i.e., whom r_i engages with at the beginning of each period), Manea shows that the system of equations given by (19) and (20), given $\mathbf{\Lambda}^{-i}$, satisfies the conditions of the contracting mapping theorem and has a unique fixed point $\tilde{\mathbf{u}}^i(\mathbf{\Lambda}^i | \mathbf{\Lambda}^{-i})$; furthermore, he shows that this solution, expressible as the determinants of this system of linear equations using Cramer's rule, varies continuously in $\mathbf{\Lambda}^i$. Given the construction of \tilde{v}_h^i and similar arguments, it is straightforward to show that $\tilde{\mathbf{u}}(\cdot)$ also varies continuously in $\mathbf{\Lambda} \equiv \{\mathbf{\Lambda}^h\}_{h \in G}$.

By Lemma B.2, we can find $\underline{\Lambda}$ such that if $\Lambda_h^i > \underline{\Lambda} \forall h \in G \setminus i$ (and thus r_i expects that agreements $G \setminus i$ will form with sufficiently high probability), the indices for the first and second highest values over $\tilde{v}_h^i(\cdot)$ will coincide with those of the first and second-highest values over $v_h^i(G)$ (and that these values can be made arbitrarily close to one another). Choose $\underline{\Lambda} < 1$ such that this holds for all $i \in G$. Furthermore, by the previous claim, we can find $\underline{\delta}$ such that for all $\delta > \underline{\delta}$ and $i \in G$, any MPE where $\Lambda_h^i > \underline{\Lambda} \forall h \in G \setminus i$ implies that $\Lambda_i^i > \underline{\Lambda}$.

Let $\mathcal{L}(\underline{\Lambda})$ denote the set of probability distributions $\mathbf{\Lambda}$ over all representatives i such that $\Lambda_i^i \geq \underline{\Lambda} \forall i \in G$. For any vector of $\mathbf{u} \equiv \{u^i\}_{i \in G}$, let $\tilde{\mathbf{\Lambda}}(\mathbf{u}; \mathbf{\Lambda})$ denote the set of probabilities in $\mathcal{L}(\underline{\Lambda})$ consistent with optimization by representative i —i.e., $\tilde{\Lambda}_h^i(\cdot) > 0$ only if $h \in \arg \max \tilde{v}_h^i(\mathbf{\Lambda}^{-i}) - \delta u_h^i$, and $\tilde{\Lambda}_i^i \geq \underline{\Lambda}$. Consider the correspondence $\tilde{\mathbf{\Lambda}}(\tilde{\mathbf{u}}(\mathbf{\Lambda}); \mathbf{\Lambda}) \Rightarrow \mathbf{\Lambda}$ restricted to the domain $\mathcal{L}(\underline{\Lambda})$. By construction, the correspondence is non-empty valued: by the previous claim, a best response for each r_i given that $G \setminus i$ forms with sufficiently high probability (Guaranteed for values in $\mathcal{L}(\underline{\Lambda})$) is to engage with i with positive probability (since $i \in \arg \max \tilde{v}_i^i(\cdot) - \delta \tilde{u}_i^i(\cdot)$ for $\delta > \underline{\delta}$). Such a correspondence also has a closed graph and is convex valued, and since $\mathcal{L}(\underline{\Lambda})$ is compact and convex, an application of Kakutani's fixed point theorem ensures the existence of a fixed point $\mathbf{\Lambda}^*$. This fixed point ensures that expected payoffs $\tilde{\mathbf{u}}(\mathbf{\Lambda}^*)$ and expected bilateral gains from trade $\tilde{\mathbf{v}}(\mathbf{\Lambda}^*)$ are consistent with probabilities induced by $\mathbf{\Lambda}^*$, and probabilities $\mathbf{\Lambda}^*$ are consistent with expected payoffs from actions. Following the arguments of Manea, construction of strategies that yield the desired payoffs, verification that they comprise an MPE, and verification that payoffs to all agents are non-negative is straightforward. Furthermore, for sufficiently high δ , the constructed MPE will result in G forming at prices arbitrarily close to NNTR payments.

B.4 Proof of Proposition 4.5

The proof of Proposition 4.4 establishes that for $\underline{\Lambda}$ and $\underline{\delta}$ sufficiently high, if $\delta > \underline{\delta}$, any MPE outcome in any subgame with stable network G being announced has network G being formed with probability $\Lambda > \underline{\Lambda}$ at prices arbitrarily close to NNTR prices. Consequently, for sufficiently high δ , the unique best response for MCO j at period 0 is to announce the insurer optimal stable network G^* at period-0 in any MPE where the announced network forms with probability $\Lambda > \underline{\Lambda}$.

C Simulation Details

For every market, we proceed by examining *all* possible BS networks $G \in \mathcal{G}_{BS}$, and computing the set of NNTR prices $\mathbf{p}^*(G, \phi^*(\cdot))$ and premiums $\phi^*(G, \mathbf{p}^*(\cdot))$ such that (the hospital system equivalent) for equations (2)-(4) hold for all hospital systems negotiating with BS, and premiums for all MCOs satisfy (10). Given the set of premiums, prices, and implied insurance enrollment and hospital utilization patterns by consumers, we evaluate whether the network G is stable by testing whether $[\Delta_{ij} \pi_j^M(G, \mathbf{p}^*, \phi^*)] > 0$ and $[\Delta_{ij} \pi_i^H(G, \mathbf{p}^*, \phi^*)] > 0$ for all $i \in G$, $j \in \{BS\}$. Finally, once the set of stable networks \mathcal{G}_{BS}^S for BS has been

determined, we are able to select the stable network that maximizes the appropriate objective (i.e., social welfare, consumer surplus, or BS profits). A similar procedure is used when we solve for NN as opposed to NNTR prices.

To determine NNTR prices and premiums for a given G , we employ the following algorithm:

1. Initialize \mathbf{p}^0 and $\boldsymbol{\phi}^0$ to observed prices and premiums.
2. At each iteration t , for a given $\boldsymbol{\phi}^{t-1}$ and \mathbf{p}^{t-1} :
 - (a) Update premiums and demand terms so that $\boldsymbol{\phi}^t(\cdot)$ satisfy (10) for all MCOs (see Ho and Lee, 2017, for further details).
 - (b) Update prices \mathbf{p}^t by iterating on the following:
 - i. Update Nash-in-Nash prices $\mathbf{p}^{Nash}(\cdot)$ for all $i \in G$ via the hospital system equivalent of the matrix inversion of the first-order condition for (2) (see Ho and Lee, 2017).
 - ii. For all $i \in G$ in sequence, search over $k \notin G$, and compute $p_{ij}^{OO}(G)$ as the solution to:
$$\pi^M(G, \{p_{ij}^{OO}(G), \mathbf{p}_{-ij}, \boldsymbol{\phi}^t\}) = \max_{k \in \mathcal{H} \setminus G} \left[\pi^M(G \setminus ij \cup kj, \{p_{kj}^{res}(\cdot), \mathbf{p}_{-ij}, \boldsymbol{\phi}^t\}) \right].$$
 - iii. Update $p_{ij} = \min(p_{ij}^{Nash}, p_{ij}^{OO})$.
3. Repeat step 2 until premiums converge (sup-norm of \$1).