

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

Kate Ho[†]

Robin S. Lee[‡]

December 2016

PRELIMINARY; COMMENTS WELCOME

(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 in order to evaluate the effects of *narrow networks* and restrictions on their use—a form of quality regulation. An insurer may wish to exclude hospitals in order to steer patients to less expensive providers, cream-skin enrollees, and negotiate lower reimbursement rates. 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 network exclusion. Pairing our framework with hospital and insurance demand estimates from Ho and Lee (forthcoming^a), we compare social, consumer, and insurer optimal networks for the largest non-integrated HMO carrier in California across 14 geographic markets. Both the insurer and consumers prefer narrower networks than the social planner in most markets: the insurer benefits from lower reimbursement rates (almost 50% in some markets), and passes a portion of the savings along in the form of lower premiums. However, the social planner may prefer a fuller network if it encourages the utilization of more efficient insurers and providers. We argue that the appropriateness of network regulation is context specific, and will depend on premium setting constraints and the generosity and efficiency of alternative insurance products.

Keywords: health insurance, narrow networks, selective contracting, hospital prices, bargaining, bilateral oligopoly

JEL: I11, L13, L40

*We thank Liran Einav, Glenn Ellison, Paul Grieco, Phil Haile, Barry Nalebuff, Mike Riordan, Bill Rogerson, Mike Whinston, Alexander Wolitzky, Ali Yurukoglu, and numerous conference and seminar participants for helpful discussion. 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 hospitals. Selective contracting by insurers 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 providers 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 hospital networks may adversely affect consumers by restricting 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 can be viewed as 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 diverges from the social optimum or the regulated level, and on any indirect effects of the regulation (such as the impact on prices). 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—thereby endogenizing the “cost of quality” for the insurer. An insurer thus may sacrifice social or productive efficiency in its choice of network in order strengthen its bargaining leverage with respect to providers.

In this paper we consider the potential effects of network adequacy regulations in the U.S. commercial (employer-sponsored) health insurance market. We first use a simple framework to highlight the fundamental economic trade-offs facing a profit maximizing insurer and a social planner when deciding whether or not to exclude a hospital. We then extend the estimated model of the U.S. commercial health care market developed in Ho and Lee (forthcoming*a*)—which incorporates insurer premium setting and consumer demand for hospitals and health insurers—to include strategic network formation by an insurer and a model of bargaining with endogenous outside options. Finally, we simulate the equilibrium hospital networks that would be chosen by a profit maximizing

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

insurer, and by a social planner maximizing either consumer or social welfare. By comparing the predicted outcomes to those generated if an insurer’s network is required to include all hospitals in a market, we are able to evaluate the effectiveness of certain forms of network regulation.

Our analysis focuses on the setting provided by the California Public Employees’ Retirement System (CalPERS), a large health benefits manager that offers access to three large insurance plans for over a million individuals across multiple geographic markets in 2004. Our simulations adjust the hospital network of one of the plans—an HMO product offered by Blue Shield—across 14 geographic markets in California. We hold fixed the networks of the other two plans: Kaiser Permanente, a vertically integrated insurer, and Blue Cross, a broad-network PPO plan that contracts with essentially every hospital in the markets that it covers. We allow these three plans to engage in premium competition with one another for enrollees. There are additional institutional constraints governing premium setting and cost-sharing, common in the health insurance industry, that we condition upon in our analysis.³ A key empirical component of our exercise is an estimated model of insurer and hospital demand from Ho and Lee (forthcoming a) which leverages detailed admissions, claims and enrollment data from CalPERS. Understanding the private and public benefits of exclusion requires us to predict how consumers’ insurance enrollment and hospital utilization decisions—which are critical inputs into insurers’ revenues and costs—are affected by counterfactual changes in insurer networks. The 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.

One of the methodological contributions of this paper is the extension of a bargaining concept that has been used in previous empirical work on insurer-hospital negotiations (e.g., Gowrisankaran, Nevo and Town (2015) and Ho and Lee (forthcoming a)) to allow for network formation and endogenous outside options. Commonly referred to as *Nash-in-Nash* bargaining (Collard-Wexler, Gowrisankaran and Lee, 2016), this bargaining concept predicts that each hospital is paid a fraction of its marginal contribution to an insurer’s network; an insurer thereby has an incentive to add hospitals to the network in order to reduce each hospital’s marginal contribution and, hence, reimbursement rate. However, *Nash-in-Nash* bargaining also has the feature that hospitals outside of an insurer’s network do not influence negotiated payments; this may be unreasonable if an insurer is able to replace an included hospital upon a bargaining disagreement with another hospital outside the current network. We develop a new bargaining concept, *Nash-in-Nash with Threat of Replacement* (NNTR), that relaxes this restriction. It both endogenizes the choice of an insurer’s network, and allows the insurer to replace an included hospital with an excluded alternative. We provide a non-cooperative extensive form and conditions under which a single network and set of

³ 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 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.

negotiated prices, governed by this solution concept, emerges as the unique equilibrium outcome. Furthermore, we show that while the Nash-in-Nash concept has difficulty rationalizing exclusion introduced by insurers in the year after our data, our NNTR concept does not.⁴

Our simulations indicate that the socially optimal network is often broad: in 10 out of 14 geographic markets, the social welfare maximizing hospital network includes every major hospital system. However, we predict that the network that Blue Shield would choose if reimbursement prices were determined via Nash-in-Nash bargaining is even broader: all major hospital systems are included in 13 out of 14 markets, consistent with the intuition that the Nash-in-Nash model tends to favor inclusion. In contrast, with NNTR bargaining, Blue Shield’s preferred network is full in only 6 out of 14 markets, and is strictly narrower than the social optimum in four markets. In these areas, Blue Shield achieves additional rate savings of between 8-48% through exclusion. Some of these savings are passed through to consumers in reduced premiums, hence consumer surplus is higher under Blue Shield’s preferred choice than under the broader socially optimal outcome. When we investigate this distinction we find that the social benefit from broader networks is due to consumer selection across plans. When Blue Shield drops hospitals, consumers benefit from the option to purchase a relatively low-premium, narrow-network product. However, some consumers respond by switching out of Blue Shield and into the broader-network Blue Cross PPO plan. That plan has a higher underlying cost of non-hospital care, an inefficiency which leads to a reduction in social surplus.

We consider the implications for the impact of network regulation in these markets. If Blue Shield was required to negotiate with all hospitals (which we refer to as “full network regulation”), this would actively constrain Blue Shield in 8 out of 14 markets under NNTR. In those areas, Blue Shield’s hospital payments would increase by 26% per enrollee; total welfare would increase by \$12 per capita; but consumer welfare would fall by \$36 per capita due to the resulting increase in premiums. This finding underscores the importance of accounting for premium adjustments in response to network adequacy requirements. It also raises the question of how to define the social planner’s objective function. If the regulator—represented in our setting by CMS or the state of California—seeks to maximize social surplus, it might choose to impose full network regulation in the markets that we consider. If, however, a greater weight is placed on consumer surplus, then such regulations may not be optimal. We note the caveat that our findings are context-specific: in particular they rely on the existence of a broad-network and a relatively inefficient outside option. In ongoing work we are repeating our simulations in the absence of the broad network Blue Cross PPO plan as an outside option; our results will be added to a future iteration of this paper.

Our current results provide partial answers to the broader question regarding the impact of minimum quality standards. The magnitude of the quality distortion in the absence of regulation,

⁴The predictions of the Nash-in-Nash and the NNTR models coincide in the year of our data, in the CalPERS setting where the Department of Managed Health Care constrains networks to be essentially complete. This implies that the Nash-in-Nash assumption in Ho and Lee (forthcoming^a) is reasonable in the context considered there. The distinction between the two models has implications for hospital rates when a hospital is excluded, and therefore for insurers’ exclusion incentives, but not for predicted prices when networks are full.

and the effect of regulation on prices, are both empirically substantive. We also show that hospital rate negotiations introduce a new, and economically meaningful, incentive to distort quality away from the industry optimum. In many instances an insurer is able to secure lower hospital prices by committing to negotiate with a narrow hospital network. The magnitude and direction of this “pecuniary” distortion will depend on the details of the bargaining protocol that is used, and in particular, the extent to which exclusion affects rate negotiations. Nash-in-Nash bargaining generates an incentive to inefficiently *include* hospitals, while our NNTR model implies *exclusionary* incentives. Our approach may prove useful for evaluating the effect of quality standards in other markets where upstream prices are negotiated and firms can commit ex ante to the number of agents with which they will negotiate (for example in pharmaceutical formulary design).

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 and potential welfare consequences of narrow networks by developing a model that allows us to predict the impact of counterfactual regulatory schemes on outcomes such as the set of excluded hospitals and the resulting negotiated prices.

Our theoretical motivation for the NNTR bargaining concept adapts 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). Prager (2016) shows that similar incentives exist when insurers offer tiered hospital networks in which some hospitals are available at lower co-insurance rates than others. Arie, Grieco and Rachmilevitch (2016) incorporate repeated interaction and limits on the number of simultaneous negotiations by the same insurer; Ghili (2016) and Liebman (2016) allow excluded hospitals to affect insurers’ negotiated rates with included hospitals, as in this paper.⁵ These latter 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 (forthcominga), with the primary objective of quantifying the impact of narrow networks on negotiated prices. Our focus is on the broader issue of health care network regulation; the impact

⁵ There 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.

on insurer-hospital negotiations is one input into the welfare and efficiency considerations that we analyze.

Finally, we note that our exercise can be seen as similar in spirit to Handel, Hendel and Whiston (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, evidence 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. Networks have also narrowed in the commercial health insurance market. The 2015 Employer Health Benefits Survey, released by the Kaiser Family Foundation, found that seventeen percent of employers offering health benefits had high performance or tiered networks in their largest health plan which provided financial or other incentives for enrollees to use 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.⁶

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. Two institutional features of the US private commercial health care market are important to have in mind before continuing. First, different providers often negotiate different reimbursement rates with a particular insurer. 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. Insurers thus bear the majority of the incremental price differences when consumers visit a high-priced versus low-priced hospital. Consequently, narrow networks may represent an important instrument for insurers to steer patients towards lower-priced (and potentially lower-cost) providers.

⁶This 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/>.




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.

Second, 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 a particular provider.

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.⁷ For now, assume that premiums for this insurer remain fixed, and that the insurer is able to reimburse hospitals at their marginal costs (with the potential addition of a fixed fee). 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.

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 costs, 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.⁸

Consider what might occur if the insurer drops a hospital from its network. As depicted in

⁷ 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.

⁸Note that the marginal cost curve may be downward sloping in the presence of adverse selection.

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 generating a change in the identity of the marginal consumer. At the same time as a 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 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 insurer’s costs 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 skimming.

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 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, the social planner would consider whether $A + E + F < B$, where F is the adverse 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 inefficient exclusion if $E + F$ is positive. The sign of this term is theoretically ambiguous. For example, it may be negative if the outside option is more efficient than the insurer, implying a social gain from the plan dropping the hospital and thereby inducing consumers to enroll elsewhere ($F \ll 0$). $E + F$ is particularly likely to be negative if the insurer’s remaining enrollees have a low valuation for the excluded hospital (E is relatively small). Conversely, $E + F$ may be positive if the insurer is more efficient than the outside option and E is large.

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.⁹ 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.

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

⁹There 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.

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 to steer patients to more efficient 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 framework also suggests that the network chosen by a profit-maximizing insurer may differ from that preferred by the social planner. A private firm choosing quality will optimize with respect to the marginal rather than the average consumer. In addition, hospital prices are negotiated and may be influenced by the network that is chosen: thus, 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 (or quality) 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 Minimum Quality Standards. The previous literature emphasizes that 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, the endogenous cost of quality (i.e., broader networks) induced by insurer-hospital rate negotiation 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 Overview of Our Analysis and Empirical Application

The remainder of our paper examines a particular setting in which we quantify the incentives explored in the previous section. Following Ho and Lee (forthcoming*a*), 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 largest employer-sponsored health benefits purchaser in the United States; 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.

For over a decade, CalPERS employees have been able to access plans from several large carriers, including a PPO plan from Anthem Blue Cross (BC), an HMO from California Blue Shield (BS), and an HMO plan offered by Kaiser Permanente. The Blue Cross PPO is a broad network plan that offers essentially every hospital in the markets it covers. Kaiser Permanente is a vertically integrated insurer that owns 27 hospitals in California and generally does not provide access to non-Kaiser hospitals. The BS HMO historically offered a fairly broad hospital network, although narrower than that of the BC PPO: in 2004 its network contained 189 hospitals compared to 223 for Blue Cross.¹⁰

In June 2004, Blue Shield filed a proposal with the California Department of Managed Health Care (DMHC) to exclude 38 high-cost providers including 13 hospitals from the Sutter system. The proposal was vetted by the DMHC with the objective of ensuring access and continuity of care for CalPERS enrollees. 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 28 hospitals were excluded from the network in 2005.¹¹ This anecdote suggests that the DMHC's policies on network oversight actively constrained Blue Shield's network choices, preventing the exclusion of hospitals that would otherwise have been dropped.

The setting provided by CalPERS is ideal for studying network design for several reasons. First, we have sufficiently detailed data 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 chose to exclude hospitals in a time period close to the year of our data, and that it was permitted to exclude fewer hospitals than requested. 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. The question of whether potential interventions by the social planner (here the DMHC) to ensure access are welfare-improving is also clearly empirically relevant in this case.

3.1 Overview of Model

In order to move beyond the simple discussion of costs and benefits provided in the previous section and examine the welfare impact of equilibrium health care markets, we rely on a stylized model of how insurers, hospitals, employers, and consumers interact in the U.S. commercial health care market. Our proposed timing is:

¹⁰A hospital is counted as in-network if it contains at least 10 admissions in our data.

¹¹The hospitals that were dropped included four major Sutter hospitals in the Greater Sacramento area and eight hospitals in the Greater Bay area. See Zaretsky and pmpm Consulting Group Inc. (2005) for details.

- 1a. Insurers or managed care organizations (MCOs) bargain with hospitals over whether they are included in their network, and the reimbursement rates that are paid.
- 1b. Simultaneously with the determination of hospital networks and negotiated rates, the employer and the set of MCOs bargain over the per-household premium charged by each MCO.
2. Given hospital networks and premiums, households choose to enroll in an MCO, determining household demand for each MCO.
3. After enrolling in a plan, each individual becomes sick with some probability; those that are sick visit some hospital in their network.

These assumptions approximate the timing of decisions in the commercial health insurance market, in which insurers negotiate networks and choose premiums in advance of each year’s open enrollment period. During that period, households observe insurance plan characteristics and choose a plan in which to enroll for the following year. Individual enrollees’ sickness episodes then arise stochastically throughout the year. Ho and Lee (forthcoming a) adopts a variant of this model to examine the welfare effects of insurer competition. That paper conditions on the set of hospital networks that are observed in the data when examining the determination of rates in Stage 1a, and holds networks fixed in its counterfactual simulations.

In this paper we follow Ho and Lee (forthcoming a) in the specification and implementation of Stages 1b, 2, and 3, which we outline in Section 5. We introduce a new specification for Stage 1a which allows an insurer’s hospital network to be endogenously determined. As we detail in the next section, this extension also motivates the adoption of a new bargaining solution when modeling the determination of negotiated hospital reimbursement rates. Consistent with our empirical setting, we assume that the networks offered by Blue Cross and Kaiser Permanente are determined *ex ante*. We model the network choice and hospital rates of Blue Shield’s HMO plan conditional on the other networks. We allow all three plans’ premiums to adjust in response to changes in BS networks and prices, through a Nash-in-Nash bargaining game with the employer.

4 Equilibrium Hospital Networks and Reimbursement Rates

This section extends the framework of the U.S. commercial health care market developed in Ho and Lee (forthcoming a) by endogenizing an insurer’s hospital network. We also introduce a bargaining solution that incorporates the ability of an insurer to “play off” hospitals with those excluded from its network in order to negotiate more advantageous rates.

We first define what we refer to as the *Nash-in-Nash with Threat of Replacement* (NNTR) bargaining solution for any particular hospital network that can be formed. We interpret this solution as one in which each bilateral hospital-insurer pair engages in simultaneous Nash bargaining over its combined gains-from-trade, and where the insurer can threaten to replace its bargaining partner with an alternative hospital that is not on the insurer’s network. This bargaining solution

naturally extends a commonly used surplus division rule in applied work on bilateral oligopoly—referred to as the *Nash-in-Nash* bargaining solution (c.f. Collard-Wexler, Gowrisankaran and Lee, 2016)—to environments where firms have endogenously determined outside options. Such a solution will only be defined for networks that we refer to as *stable*, where that implies no party has a unilateral incentive to terminate a relationship based on the negotiated price.

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 extensive form provides one potential justification for our approach, and an interpretation for how such an outcome might arise in practice.

4.1 Setup and Primitives

Consider a set of MCOs \mathcal{M} that are offered by an employer, and hospitals \mathcal{H} . We are concerned with the determination of equilibrium hospital networks and reimbursement rates for a single MCO (which we index by j); our analysis thus should be interpreted as conditioning on the networks and reimbursement rates that are set by other MCOs $-j$ in the market (which, for this section, we hold fixed). Furthermore, we initially hold fixed the set of premiums for all MCOs, ϕ , for exposition; we thus omit them from notation here, and re-incorporate them in the next section.

Let \mathcal{G}_j denote the set of all potential hospital networks that MCO j can form. For a given $g \in \mathcal{G}_j$, we say that hospital i is included on MCO j 's network if $i \in g$, and excluded if $i \notin g$. We denote by $\pi_j^M(g, \mathbf{p}) = \tilde{\pi}_j^M(g) - \sum_{i \in g} D_{ij}^H(g) p_{ij}$ and $\pi_i^H(g, \mathbf{p}) = \tilde{\pi}_i^H(g) + \sum_{n \neq j} D_{in}^H(g) p_{in}$ 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 the stylized model described in the previous section). The key assumptions that we maintain for now are that negotiated payments enter linearly into profits, and other aspects of firm profits (including objects represented by $\tilde{\pi}_j^M$ and $\tilde{\pi}_i^H$) depend only on the realized hospital network. We take profit and demand functions as primitives for this section's analysis, and provide explicit parameterizations for them later.

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 $[\Delta_{ij}\pi_i^H(g, \cdot)] \equiv \pi_i^H(g, \cdot) - \pi_i^H(g \setminus i, \cdot)$ denote the *gains-from-trade* to MCO j and hospital i when they come to an agreement, given all other agreements in g are formed at prices \mathbf{p}_{-ij} . Additionally, let $\Delta_{ij}\Pi_{ij}(g, \mathbf{p}) \equiv [\Delta_{ij}\pi_j^M(g, \mathbf{p})] + [\Delta_{ij}\pi_i^H(g, \cdot)]$ denote the (total) bilateral gains-from-trade (or surplus) created by an agreement between MCO j and hospital i over their respective disagreement points, represented by $\pi_j^M(g \setminus i, \mathbf{p}_{-ij})$ and $\pi_i^H(g \setminus i, \cdot)$. 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 they cancel out).

Negotiated prices will be a function of an MCO’s hospital network as well as the premiums that are chosen after bargaining concludes; furthermore, insurers typically negotiate with hospital systems as opposed to individual hospitals. We abstract from these complications in our initial discussion for expositional clarity, but incorporate them later in our application and full analysis (see Section 4.4).

4.2 Nash-in-Nash with Threat of Replacement

Fix the network of hospitals $g \in \mathcal{G}_j$ with which MCO j contracts. We define the *Nash-in-Nash with Threat of Replacement* (NNTR) prices induced by 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.¹²

Formally, each NNTR price satisfies

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

where

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

is the solution to the bilateral Nash bargaining problem between MCO j and hospital i with Nash bargaining parameter $\tau \in [0, 1]$; and

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

(referred to as the “outside option” price), 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.¹³ Hospital k ’s reservation price, in turn, is defined to be the

¹²As noted above, we hold fixed the hospital networks and reimbursement prices $\{p_{in}\}$ for other MCOs $n \neq j$.

¹³The 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).

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.¹⁴

Assume that for any network g , $\mathbf{p}^*(g)$ exists and is unique.

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 to be stable if all agreements $i \in g$ are stable. Stability of an agreement $i \in g$ 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 4.1. *Network g is stable at NNTR prices \mathbf{p}^* iff, for all $i \in g$,*

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

(All proofs in the appendix). 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.¹⁵ In our extensive form representation provided in the next subsection,

¹⁴ Hospital k may earn profits even if excluded from MCO j 's network as it may contract with other MCOs $-j$.

¹⁵ 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

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}^S be the set of all networks g 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., $g^* = \arg \max_{g \in \mathcal{G}^S} \pi_j^M(g, \mathbf{p}^*(g))$.

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 reservation 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; i.e., in many circumstances, outside options only affect bargaining outcomes if they are credible and 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 $p_{kj}^{res}(\cdot)$.

To deal with the first complication, we build on what has been referred to the Nash-in-Nash (NN) bargaining solution (c.f. Collard-Wexler, Gowrisankaran and Lee, 2016) in bilateral oligopoly, which—following its use in Horn and Wolinsky (1988)—has been used by several recent applied papers to model bargaining between firms with market power (e.g., Draganska, Klapper and Villas-Boas, 2010; Crawford and Yurukoglu, 2012; Grennan, 2013). Defined for a particular network g , the Nash-in-Nash solution in our health care context specifies that reimbursement prices negotiated between hospital i and MCO j , $i \in g$, satisfy the solution to that pair’s Nash bargaining problem, assuming that all other bilateral pairs in $g \setminus i$ come to agreement; i.e., Nash-in-Nash prices satisfy

added to the network.

(2) if $p_{ij}^*(\cdot) = p_{ij}^{Nash}(\cdot) \forall i \in g$.

However, as noted in Lee and Fong (2013), the Nash-in-Nash bargaining solution does not allow the outside option from disagreement between any pair to involve firms *forming new agreements* with others.¹⁶ The NNTR concept addresses this limitation by extending Nash-in-Nash to environments where some firms can exercise an outside option and replace their current bargaining partner with another that they are not negotiating with. Viewed in this light, it is intuitive that the NNTR and the Nash-in-Nash bargaining solutions *coincide* when the network g in question is “complete”—i.e., all hospitals are in all insurers’ networks. Only when networks are incomplete, as is the case with narrow networks, will predictions differ, as the Nash-in-Nash concept may understate the extent to which an insurer can “play” hospitals off one another by forming selective networks. We return to this point when presenting the results from our application.

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; i.e., 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 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 NNTR

We now provide and analyze a non-cooperative extensive form game, in the spirit of the Nash program (c.f., 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 $g \subseteq \mathcal{G}^S$, and sends separate representatives r_i to negotiate with each hospital i s.t. $i \in g$;¹⁷
- 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

¹⁶ Assume that $i \in g$ but $k \notin g$: i.e., hospital k does not have a contract with insurer j and is not on j ’s network. Under this bargaining solution, there is no mechanism through which a hospital k serves as a constraint on the price that insurer j negotiates with hospital i .

¹⁷ The 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 may be no probability that network g actually 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.

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.¹⁸ 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 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 0.¹⁹

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 can be interpreted as one in which the seller engages in Nash bargaining with the highest valuation buyer,

¹⁸ In 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 $\delta \rightarrow 1$, and assuming an alternative specification (e.g., allowing representatives coordinate in their negotiations with hospital k for that period only, and if an agreement is reached with k , then all representatives negotiating with k are removed) will not affect our main result.

¹⁹ Using 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.

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.²⁰

Our extensive form 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.²¹

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 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). 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 4.2. *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 4.2 is straightforward. First, consider any subgame where MCO j announces a network g that contains a single hospital. By Proposition 4.1, 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 a τ fraction 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

²⁰ In 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.

²¹ The 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.

second-highest-surplus creating hospital—the MCO obtains the equivalent of $[\Delta_{kj}\Pi_{kj}(\{kj\})]$ by paying only the lump-sum equivalent of $p_{ij}^{OO}(g)$ 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 4.3. *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^* is announced in period 0;*
- *the probability that g^* is formed is greater than $1 - \varepsilon_1$;*
- *all agreements $i \in g^*$ are formed with ε_2 of NNTR prices.*

Such an equilibrium exists.

Thus, if δ is sufficiently high and we restrict attention to MPE in which the announced period-0 network is formed with sufficiently high probability, the unique MPE outcome will be the insurer optimal stable network g^* being announced and formed with probability (arbitrarily) close to 1 at prices close to NNTR prices.

An Alternative Interpretation: Sequential Negotiation. An alternative interpretation for why g^* 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 an MCO did not anticipate future renegotiation with other hospitals following the conclusion of any bargain. For example, assume that MCO j ’s network was g^* and negotiated prices were \mathbf{p}^* . Assume some hospital contract $i \in g^*$ 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 $k \notin g^*$ 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 4.2 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) to be 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 3.1, we assume in the timing of our model that premiums are determined 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). In the next section, we describe the process by which insurers bargain with employers over premiums. An equilibrium of our complete model will thus be 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 subsequent analysis, we adopt methods developed in Ho and Lee (forthcoming a) 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.²²

Synthesizing Nash-in-Nash and NNTR. Finally, our NNTR concept can 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 $\alpha \in [0, 1]$. 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 Premium Negotiations and Consumer Demand

Having developed a model of Stage 1a of the overall framework described in Section 3.1, we now move on to Stages 1b, 2 and 3, which generate the profit equations that are the primary inputs into the price bargaining model. We follow Ho and Lee (forthcoming a) in the specification of these components; we outline them briefly below.

5.1 Profit Equations and Premium Negotiations

The previous section conditioned on the networks and prices of a single MCO j and conditioned on the premiums charged by all insurers. Here, we condition on the hospital networks of all insurers, denoted G , and their premiums, $\phi \equiv \{\phi_j\}_{j \in \mathcal{M}}$.

²² We assume that the MCO and each hospital system negotiate over a single linear price; each hospital within the system receives a multiplier of this price, where the multiplier is chosen so that the ratio of prices of any two hospitals within the same system is equal to that observed in the data. See Ho and Lee (forthcoming a) for further details.

Insurer and hospital profit equations. We assume that MCO j 's profits are

$$\pi_j^M(G, \mathbf{p}, \phi) = D_j(G, \phi) \times (\phi_j - \eta_j) - \sum_{h \in G_j^M} D_{hj}^H(G, \phi) \times p_{hj}, \quad (5)$$

where $D_j(\cdot)$ is household demand for MCO j and $D_{hj}^H(\cdot)$ is the number of enrollees in MCO j who visit hospital i , both predicted from the demand system outlined below. The first term on the right-hand-side of the equation represents MCO j 's total premium revenues (given by premium payments per household, ϕ_j , net of non-inpatient hospital costs, η_j). The second 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 (denoted G_j^M), the price per admission (p_{hj}) negotiated with the hospital multiplied by the number of patients admitted. We account for variation in disease severity across admissions by scaling $D_{h,j,m}^H$ by the expected DRG weight (based on age and gender) for the relevant patient. In the empirical specification we also make the distinction that an MCO's non-inpatient hospital costs η_j are incurred on an individual and not a household basis; we use data on the number of individuals per household to scale up this component of the insurer's costs appropriately.

We assume that profits for a hospital i are

$$\pi_i^H(G, \mathbf{p}, \phi) = \sum_{n \in G_i^H} D_{in}^H(G, \phi) \times (p_{in} - c_i), \quad (6)$$

which sums, over all MCOs n with which hospital i contracts (denoted G_i^H), 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). We scale D_{in}^H by the expected DRG weight for each patient, implying that prices and costs vary across age-sex categories and across each insurer-hospital pair.

Consistent with the presence of limited cost sharing, we assume that utilization of a particular hospital h on MCO j —denoted $D_{hj}^H(\cdot)$ —does not depend on the set of negotiated reimbursement rates \mathbf{p} .

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

$$\phi_j = \arg \max_{\phi_j} \left[\underbrace{\pi_j^M(G, \mathbf{p}, \{\phi, \phi_{-j}\})}_{GFT_j^M} \right]^{\tau\phi} \times \left[\underbrace{W(\mathcal{M}, \{\phi, \phi_{-j}\}) - W(\mathcal{M} \setminus j, \phi_{-j})}_{GFT_j^E} \right]^{(1-\tau\phi)} \quad \forall j \in \mathcal{M}, \quad (7)$$

(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" (GFT) from coming to agreement, i.e., from MCO j being included on 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 $\tau^\phi \in [0, 1]$.

5.2 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; Farrell et al., 2011).

We assume that an individual of type κ 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 :

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

where δ_i are hospital fixed effects, 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.

We use the estimated hospital demand model to construct 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 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:

$$u_{f,j,m}^M = \delta_{j,m} + \alpha_f^\phi \phi_j + \sum_{\forall \kappa} \alpha_\kappa^W \sum_{k \in f, \kappa(k)=\kappa} WTP_{k,j,m} + \varepsilon_{f,j,m}^M \quad (9)$$

where $\delta_{j,m}$ is an 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, α_κ . Finally ε_{fjm}^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.

6 Empirical Analysis

In this section we summarize the data and estimation strategy used in Ho and Lee (forthcoming *a*), before moving on to the simulations that are the focus of this paper.

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. For enrollees in BS and 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.²³

²³We 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.

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 (MLRs) for each insurer by dividing total medical and hospital costs by total revenues.

In 2004, Blue Shield offered 189 hospitals across California, compared to 223 for Blue Cross and just 27 for the vertically integrated Kaiser HMO. Annual premiums for single households across BS, BC, and Kaiser were \$3782, \$4193, and \$3665; 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.

6.2 Estimation

Ho and Lee (forthcoming^a) estimate the parameters of the hospital demand equation via maximum likelihood using the admissions data. We assume that there is no outside option, and that the enrollee can choose any in-network hospital in his HSA that is within 100 miles of his zip code. The insurer demand model is also estimated via maximum likelihood, using our household-level data on plan choices, location and family composition. Kaiser is the outside option in markets where it is available.

Insurer non-inpatient hospital costs η_j and the bargaining weights (τ^ϕ, τ_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. These moments are particularly helpful in estimating the insurer’s non-inpatient hospital costs η_j .

The assumption of Nash-in-Nash bargaining over hospital prices is reasonable in the framework that holds networks fixed because CalPERS and the California Department of Managed Health Care constrain networks to be close to full in reality. In 11 out of 14 markets, all five of the largest systems are included in the BS network. This implies that, given the observed networks, the Nash-in-Nash model assumed in Ho and Lee (forthcoming^a) coincides with the NNTR model. Furthermore, omitting the three markets where one of the five largest hospital systems is excluded,

and rerunning the estimation, yields statistically similar estimates for all parameters including Nash bargaining parameters and estimated non-inpatient marginal costs. Thus the estimated $(\eta_j, \tau^\phi, \tau_j)$ can be utilized in our simulations under both bargaining models with no internal inconsistency. We hold the estimated values of $(\eta_j, \tau^\phi, \tau_j)$ constant in our simulations.

Several of the estimates can be compared to outside sources. For example the estimated values of η_j , 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 indicates average spending of \$1,644 per person per year on professional services for the three largest commercial insurers in the years 2010-12.²⁴ Similarly, the own-premium elasticities for each insurer can be compared to those estimated in previous papers.²⁵ 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 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 the estimates provide a reliable basis for our simulations. 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. All of these components of demand will affect insurer and hospital incentives during bargaining and network formation.

7 Simulations: The Impact of Narrow Networks

Our objective is to examine how Blue Shield, the only carrier offering a non-integrated HMO plan in our empirical setting, is predicted to adjust its network under different regulations when facing an outside option (the broad-network BC PPO plan and Kaiser Permanente) whose premiums are set strategically as part of the model but whose networks are not. Our simulations assess the empirical importance of the incentives to reallocate admissions to relatively inexpensive hospitals; to prompt

²⁴Kaiser 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.

²⁵We 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.

relatively unhealthy enrollees to select into other insurers; and to negotiate low reimbursement rates, all of which we have argued are likely to affect insurer incentives to exclude. By comparing these predictions to the BS network that would be chosen by the social planner maximizing either social or consumer welfare, we also assess the importance of premium adjustments and the non-inpatient costs of Blue Shield relative to other insurers in the choice set, both of which factor into our calculation of welfare.

We simulate market outcomes under different models of price bargaining and network formation across 14 different geographic markets in California. We use the model of Nash bargaining with the employer to allow all three plans' premiums to adjust when networks change, with one simplification: we conduct our simulations market-by-market, implying an assumption that negotiations occur at the market level, while in reality CalPERS constrains premiums to be set at the state level. We focus on BS's negotiations with the five largest systems in each market. We hold fixed its contracts with the remaining smaller systems and individual hospitals (which together make up less than 50% of admissions in every market, and less than 27% of admissions in 11 out of 14 markets).

We begin by finding the socially optimal stable network: i.e., the network that maximizes the sum of insurer and hospital profit and employer surplus (consumer surplus net of the employer's contribution towards premium payments). We compare this outcome to the stable network that maximizes consumer (employer) surplus alone, without accounting for firm profits. We then consider the predictions of the model if BS were unconstrained in its choice of profit maximizing network, subject to NNTR or Nash-in-Nash bargaining over reimbursement rates. Comparing the outcomes of the two models establishes the importance of the additional incentive to exclude in order to reduce rates in the NNTR model. Implementation details can be found in Appendix B.

Finally we consider the implications of the simulations for the effects of full network regulation—regulations that require networks to be complete—in this setting. 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.

A note on the observed networks. Before presenting the results of our simulations, we note that they differ from the data generating process for the networks observed in the data in at least two ways. First, as noted, we simplify the premium bargaining game by assuming market-level premiums and that negotiations are independent across markets; in the data, premiums are constrained to be constant across the entire state. 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 28—that the DMHC's approval process can represent a binding constraint.²⁶ Thus there is no reason to expect our simulated networks to “fit the data”

²⁶ The proposal was assessed against the standard that the remaining in-network hospitals (and physician groups) should have sufficient capacity to treat local BS enrollees in each market. Details are provided in a technical Report on the Analysis of the CalPERS/Blue Shield Narrow Network (Zaretsky and pmpm Consulting Group Inc. (2005)).

Table 1: Santa Barbara / Ventura

Objective		Social (NNTR)	Consumer (NNTR)	Blue Shield (NNTR) (NN)	
Surplus (\$ per capita)	BS Profits	400.02	382.11	400.02	398.42
	Hospital Profits	225.41	219.89	225.41	240.62
	Total Hosp + Insurer Costs	1,946.74	1,953.37	1,946.74	1,948.81
Transfers / Costs (\$ per enrollee)	BS Premiums	2,655.05	2,640.61	2,655.05	2679.47
	BS Hosp Pmts	338.24	324.56	338.24	364.36
	BS Hosp Costs	124.26	120.27	124.26	126.04
Δ Welfare from Full (\$ per capita)	Consumer Welfare	21.75	29.42	21.75	-
	Total Welfare	1.64	-13.90	1.64	-
	BS Market Share	0.64	0.61	0.64	0.64
Network	# Major Systems Excluded	1	3	1	0
	Sys 1 (Mercy)	1	1	1	1
	Sys 2 (Community)	1	1	1	1
	Sys 3 (Cottage)	1	0	1	1
	Sys 4 (Los Robles)	1	0	1	1
	Sys 5 (Lompoc MC)	0	0	0	1

Notes: Simulation results from Santa Barbara / Ventura HSA. Each column represents the network that maximizes social welfare, consumer welfare, or Blue Shield’s profits, under either NNTR or Nash-in-Nash (NN) bargaining over hospital reimbursement rates. All welfare figures are computed as differences from the “full network,” where all five major hospital systems are included.

in the sense of matching the observed networks, so we do not use a comparison of observed to predicted networks to test the validity of the model.

7.1 Simulation Results: Specific Markets

We begin by presenting results of our simulations for two specific markets: Santa Barbara / Ventura and Sacramento. These two markets illustrate our findings regarding the incentives of interest. We then summarize our results across all 14 markets in California.

Santa Barbara/Ventura. Table 1 presents simulation results for the Santa Barbara/Ventura HSA under the four different simulation scenarios. We focus on hospital and insurer profits, payments and costs, the change in consumer and total welfare from the full network, and whether or not each of the five largest hospital systems in the market are included in the BS network.

Our first finding is that the network that maximizes social surplus in this market also maximizes Blue Shield’s profits under NNTR bargaining. That network excludes the smallest of the five hospital systems. The consumer surplus-maximizing network is much narrower, excluding three out of five systems. A comparison of columns 1 and 2 of Table 1 shows that consumer surplus increases by approximately \$8 per capita per year when we move from the social optimal to the consumer optimal network. The reason is that, when BS drops three systems, this enables it to negotiate lower hospital prices and premiums also fall, generating a benefit to consumers that outweighs the loss from reduced hospital choice. The reduction in total welfare is due largely to a

The specific requirement was that enrollees should have access to primary medical services—both hospitals and physicians—within 30 minutes or 15 miles of their residence or workplace.

Table 2: Sacramento

Objective		Social (NNTR) BS (NN)	Consumer (NNTR) BS (NNTR)
Surplus (\$ per capita)	BS Profits	316.87	325.63
	Hospital Profits	115.86	39.24
	Total Hosp + Insurer Costs	2,148.86	2,156.73
Transfers / Costs (\$ per enrollee)	BS Premiums	2,622.69	2,508.68
	BS Hosp Pmts	334.74	173.95
	BS Hosp Costs	165.51	154.47
Δ Welfare from Full (\$ per capita)	Consumer Welfare	-	55.60
	Total Welfare	-	-28.82
	BS Market Share	0.53	0.51
Network	# Major Systems Excluded	0	4
	Sys 1 (Sutter)	1	0
	Sys 2 (Mercy)	1	1
	Sys 3 (UCD)	1	0
	Sys 4 (Freemont)	1	0
	Sys 5 (Marshall)	1	0

Notes: Simulation results from Sacramento HSA. See notes from Table 1.

\$7 per capita per year increase in insurer and non-hospital costs when enrollees switch out of BS and into the BC PPO, whose underlying non-hospital costs (η_j in the insurer profit equation) are relatively high.²⁷

A comparison of columns 3 and 4 of Table 1 illustrates the point that the Nash-in-Nash bargaining model understates the incentives for Blue Shield to exclude because it does not allow an insurer to play included and excluded hospitals off against one another. In this market the network predicted to maximize BS's profits under Nash-in-Nash bargaining is the full network. Table 1 shows that moving to NNTR, under which the insurer optimally chooses to exclude one system, implies a 7% reduction in payments made to hospitals. Consumers are better off despite the smaller network, again due to a premium reduction when a large hospital system is removed.

Finally we can interpret column 4 of Table 1 as a prediction of the outcome when networks are mandated to be complete but the insurer is permitted to negotiate prices under NNTR bargaining (or, equivalently when networks are full, under Nash-in-Nash bargaining). Such regulation reduces consumer surplus relative to the other three columns because premiums increase by approximately 2%. It also increases hospital profits by 7-9% on average, with little effect on overall welfare in this market.

Sacramento. Our results for Sacramento are set out in Table 2. In this case we predict that the socially optimal network, and that chosen by BS under Nash-in-Nash bargaining, are complete. Note again the intuition that the insurer is rarely predicted to exclude hospitals under Nash-in-Nash bargaining. In contrast, the network that maximizes consumer surplus also maximizes BS's profits under NNTR bargaining: this is much narrower, including only Mercy out of the five largest systems in the market.

²⁷There is also a reduction in consumer valuation for the BS network.

Table 3: All Markets

Objective		Social (NNTR)	Consumer (NNTR)	Blue Shield		Full (-)
				(NNTR)	(NN)	
Surplus (\$ per capita)	BS Profits	0.26%	-2.57%	1.21%	0.01%	295.17
	Hospital Profits	-1.41%	-14.42%	-11.26%	-0.03%	197.81
	Total Hosp + Insurer Costs	-0.02%	0.20%	-0.03%	0.00%	2,087.75
Transfer / Cost (\$ per enrollee)	BS Premiums	-0.17%	-1.98%	-1.05%	0.00%	2,657.12
	BS Hosp Pmts	-1.37%	-16.46%	-9.79%	-0.04%	382.18
	BS Hosp Costs	0.00%	0.28%	0.10%	0.00%	528.75
Welfare Δ from Full (\$ per capita)	Consumer Welfare	3.16	29.87	15.44	0.08	-
	Total Welfare	0.33	-12.66	-2.31	0.04	-
Number of "Full" Network Markets		10/14	1/14	6/14	13/14	-
Avg Systems Excluded (Cond'l on Narrow)		1.00	2.38	1.63	1.00	-

Notes: Unweighted averages across markets. Percent figures and welfare calculations represent changes relative to outcomes under a full network; outcome levels for full network are presented in right-most column.

Recall that Sutter, the largest system in Sacramento, was dropped by BS in 2005 after a lengthy process of negotiation with the Department of Managed Health Care. This together with the existence of a broad-network alternative option, the Blue Cross PPO, makes the narrow predicted network empirically reasonable. Furthermore, the Nash-in-Nash bargaining concept, when paired with the ability for BS to choose its profit maximizing network, has difficulty rationalizing the observed exclusion of Sutter hospitals.

Table 2 demonstrates that mandating a complete network in this market would double Blue Shield's payments to hospitals and triple hospital profits, while also generating an approximately 5% premium increase and a corresponding reduction in consumer welfare. Social surplus would increase by approximately \$29 per capita per year, again largely due to a movement of enrollees into Blue Shield, implying a reduction in insurer costs.

7.2 Summary Results Across Markets

Table 3 summarizes the simulation results across all markets. The socially optimal network is predicted to be complete in 10 out of 14 markets in California. Moving from the full network (fourth column) to the socially optimal network results in an average 1.4% reduction in hospital profits, a \$3.16 increase in consumer surplus and a small (\$0.33 per capita) increase in social surplus. The network that maximizes consumer surplus is complete in only one out of 14 markets. This narrower network generates much lower premiums: on average 2% lower than the full network and 1.7% lower than the socially optimal network.

In contrast to the network that would be chosen by the social planner, the insurer-optimal stable network under NNTR bargaining is complete in only six markets, and is strictly narrower than the social optimum in four markets. Blue Shield achieves rate savings of over 8% on average from exclusion, relative to the social optimum, and up to 48% savings in the four markets where the networks differ. This difference leads to premium reductions and resulting increases in consumer surplus. The social losses from exclusion are thus due largely to substitution of enrollees to less-

efficient outside options (particularly the Blue Cross PPO plan).

We note that these findings are context-specific: in particular, they rely on the presence of a broad-network alternative plan option which mitigates the consumer harm from narrow networks. Our welfare conclusions may be different in a setting without such an outside option. We consider this possibility in ongoing work by repeating our simulations under the assumption that Blue Cross is not offered to CalPERS enrollees.

7.3 Implications for Minimum Quality Standards

Our simulations imply that full network regulation would constrain Blue Shield, assuming NNTR bargaining, in eight out of 14 markets. BS payments to hospitals would increase by 10% on average, and by 26% in the markets where the constraints would be binding. Consumer surplus would fall by \$36 per capita in these eight markets; but total welfare would increase by \$12 per capita per year. The results underscore the importance of accounting for premium adjustments in response to network regulation: if premiums were held fixed, regulation would yield a \$17 per capita gain for consumers in the eight markets with binding constraints.

More generally, our simulations provide evidence regarding the impact of narrow networks on insurer payments, costs, and welfare. We predict greater incentives to exclude hospitals for Blue Shield (under the NNTR bargaining model) and for the consumer than for the social planner. Full network regulation may encourage utilization of the more efficient insurer—hence increasing social surplus—but it limits insurers’ ability to offer differentiated products, thereby limiting consumer choice. This raises a question regarding the objective function of the social planner. If the goal is to maximize consumer surplus, then full network regulation may not be optimal in our setting.

The results also have implications for a broader question regarding the impact of minimum quality standards. We show that rate negotiations introduce a new incentive for the insurer to distort quality away from the industry optimum. The direction of this distortion relative to the social optimum depends on the details of the bargaining model: a Nash-in-Nash formulation tends to promote inclusive networks, while our NNTR model introduces a new exclusionary incentive which generates relatively narrow networks and hence a substantial effect of minimum quality standards on realized networks (and, through them, on prices and premiums). These findings are likely to generalize to other settings where upstream prices are negotiated.

8 Conclusion and Future Work

Our model examines provider network design from the perspective of a profit-maximizing insurer and a social planner who may maximize social surplus or consumer surplus. We discuss several reasons why the insurer might choose to exclude a hospital: to cream-skin healthy consumers; to steer remaining enrollees to lower-priced hospitals; and to negotiate lower rates with those that remain. We show that insurer-hospital bargaining over prices generates an endogenous component of insurer costs which has implications for the effect of quality standards in this market. Rather

than intensifying price competition, as in markets with exogenous costs, such standards can lead to increases in both hospital rates and consumer premiums.

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.
- Crawford, Gregory S., and Ali Yurukoglu.** 2012. “The Welfare Effects of Bundling in Multi-channel Television Markets.” *American Economic Review*, 102(2): 643–685.
- 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.
- Farrell, Joseph, David J. Balan, Keith Brand, and Brett W. Wendling.** 2011. “Economics at the FTC: Hospital Mergers, Authorized Generic Drugs, and Consumer Credit Markets.” *Review of Industrial Organization*, 39: 271–296.
- 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. Whiston.** 2015. “Equilibria in Health Exchanges: Adverse Selection Versus Reclassification Risk.” 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.** forthcoming^a. “Insurer Competition in Health Care Markets.” *Econometrica*.
- Ho, Kate, and Robin S. Lee.** forthcoming^b. “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 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 $[\Delta_{ij}\Pi_{ij}(g, \mathbf{p})] = [\Delta_{ij}\Pi_{ij}(g, \mathbf{p}_{(ij=0)})] \forall i \in g, \forall \mathbf{p}$.
- For each $g \in \mathcal{G}$ 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.²⁸ In the case of lump-sum transfers (for Proposition 4.2 and Proposition 4.3), 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 A.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):

$$p_{ij}^{Nash}(g, \mathbf{p}_{-ij})D_{ij}^H(g) = (1 - \tau)[\Delta_{ij}\pi_j^M(g, \mathbf{p}_{(ij=0)})] - \tau[\Delta_{ij}\pi_i^H(g, \mathbf{p}_{(ij=0)})] \forall i \in g. \quad (10)$$

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:

$$\begin{aligned} p_{ij}^{OO}(g)D_{ij}^H(g) &= \pi_j^M(g, \mathbf{p}_{(ij=0)}) - \pi_j^M(g \setminus i \cup k, \mathbf{p}_{-ij}) - [\Delta_{kj}\pi_k^H(g \setminus i \cup k, \mathbf{p}_{-ij})] \\ &= [\Delta_{ij}\pi_j^M(g, \mathbf{p}_{(ij=0)})] - [\Delta_{kj}\pi_j^M(g \setminus i \cup k, \mathbf{p}_{-ij})] - [\Delta_{kj}\pi_k^H(g \setminus i \cup k, \mathbf{p}_{-ij})] \\ &= [\Delta_{ij}\pi_j^M(g, \mathbf{p}_{(ij=0)})] - [\Delta_{kj}\Pi_{kj}(g \setminus i \cup k, \mathbf{p}_{-ij})]. \end{aligned} \quad (11)$$

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 (10) and (11) and dividing through by $D_{ij}^H(g)$. \square

A.1 Proof of Proposition 4.1

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 $k = \arg \max_{h \in N_i \setminus i} [\Delta_{hj}\Pi_{hj}(g \setminus$

²⁸ 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.

$i \cup h, \mathbf{p}_{-ij}^*$], $[\Delta_{kj}\Pi_{kj}(g \setminus i \cup h, \mathbf{p}_{-ij}^*)] > \tau[\Delta_{ij}\Pi_{ij}(g, \mathbf{p}^*)]$ (since $\tau < 1$). By Lemma A.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:

$$\begin{aligned} \pi_i^H(g, \mathbf{p}_{(ij=0)}^*) + p_{ij}(g)D_{ij}^H(g) &= \pi_i^H(g, \mathbf{p}_{(ij=0)}^*) + \underbrace{[\Delta_{ij}\pi_j^M(g, \mathbf{p}_{(ij=0)}^*)] - [\Delta_{kj}\Pi_{kj}(g \setminus i \cup k, \mathbf{p}_{(ij=kj=0)}^*)]}_{\text{From (11)}} \\ &= \pi_i^H(g \setminus i, \mathbf{p}_{(ij=0)}^*) + \underbrace{[\Delta_{ij}\Pi_{ij}(g, \mathbf{p}_{(ij=0)}^*)] - [\Delta_{kj}\Pi_{kj}(g \setminus i \cup k, \mathbf{p}_{-ij}^*)]}_{< 0 \text{ by assumption (since } [\Delta_{ij}\Pi_{ij}(g, \mathbf{p}^*)] = [\Delta_{ij}\Pi_{ij}(g, \mathbf{p}_{(ij=0)}^*)])} \\ & \tag{12} \end{aligned}$$

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 A.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 (12), 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}^* .

A.2 Proof of Proposition 4.2

For this proof and for Proposition 4.3, 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.^{29,30}

We derive the equivalent lump-sum NNTR prices $P_{ij}^*(g) \equiv \min\{P_{ij}^{Nash}(\cdot), P_{ij}^{OO}\}$, where (using (10) and (11)):

$$P_{ij}^{Nash}(g) = (1 - \tau)[\Delta_{ij}\pi_j^M(g)] - \tau[\Delta_{ij}\pi_i^H(g)], \tag{13}$$

$$P_{ij}^{OO}(g) = [\Delta_{ij}\pi_j^M(g)] - [\Delta_{kj}\Pi_{kj}(g \setminus i \cup k)]. \tag{14}$$

Furthermore, Lemma A.1 in this setting 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 4.1 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 4.2 hold for subgames where the period-0 network contains a single hospital; we defer establishing existence until the more general multiple hospital case (which also applies here).

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 .

²⁹Transfers 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.

³⁰We 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.

By Proposition 4.1, 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 (13). By Lemma A.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 (14). By Lemma A.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:

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

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 $\mathbf{\Lambda}^{-i}$.³¹

Claim: For any $\varepsilon_1, \varepsilon_2 > 0$, there exists $\underline{\Lambda} < 1$ and $\underline{\delta} > 0$ such that if $\Lambda_h^i > \underline{\Lambda} \forall h \in g \setminus i$ 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 $\max(\tau[\Delta_{ij} \Pi_{ij}(g)], [\Delta_{kj} \Pi_{kj}(g \setminus i \cup k)]) = \max(\tau v_i^i(g), v_k^i(g))$, 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 $v_h^i(g)$.³² 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 4.1, $i = k_{(1)}^i(g)$ and the result follows.

We next establish 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 for each r_i , $i \in g$, and its bargaining partners must satisfy:

$$u_0^i = \sum_{h \in N_i} \Lambda_h^i (\tau(\tilde{v}_h^i(\mathbf{\Lambda}^{-i}) - \delta u_h^i) + (1 - \tau)\delta u_0^i) \quad (15)$$

$$u_h^i = \Lambda_h^i (\tau \delta u_h^i + (1 - \tau)(\tilde{v}_h^i(\mathbf{\Lambda}^{-i}) - \delta u_0^i)) \quad \forall h \in N_i \quad (16)$$

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 (15) and (16), 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 the previous claim, 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

³¹ 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 18, 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.

³²This follows since profits are assumed to be finite for any potential network.

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 $\Lambda_i^i > \underline{\Lambda} \forall i \in g$. 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 $\tilde{\Lambda}_i^i > \underline{\Lambda}$. 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.

A.3 Proof of Proposition 4.3

The proof of Proposition 4.2 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 $\Lambda > \underline{\Lambda}$.

B 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 (7). 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}^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 ϕ^0 to observed prices and premiums.
2. At each iteration t , for a given ϕ^{t-1} and \mathbf{p}^{t-1} :
 - (a) Update premiums and demand terms so that $\phi^t(\cdot)$ satisfy (7) for all MCOs (see Ho and Lee, forthcoming*a*, 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, forthcoming*a*).
 - ii. For all $i \in g$ in sequence, search over $k \notin g$, and compute $p_{ij}^{OO}(g)$ as the solution to:
$$\pi^M(g, \{p_{ij}^{OO}(g), \mathbf{p}_{-ij}, \phi^t\}) = \max_{k \in \mathcal{H} \setminus g} \left[\pi^M(g \setminus ij \cup kj, \{p_{kj}^{res}(\cdot), \mathbf{p}_{-ij}, \phi^t\}) \right].$$
 - iii. Update $p_{ij} = \min(p_{ij}^{Nash}, p_{ij}^{OO})$.
3. Repeat step 2 until premiums converge (sup-norm of \$1).