Bargaining and Restricted Networks
ECON 590/890

Luca Maini*

March 22, 2021

See most recent version here

Contents

*These notes are a work in progress. Please check back frequently for an updated version, and email lmaini@email.unc.edu for corrections.

1 INTRODUCTION TO NARROW NETWORKS

A good way to introduce the core tradeoff at the heart of narrow networks is to think about the role of insurance. Insurance arises naturally out of risk-aversion. Individuals who wish to protect themselves from risk will willingly enter into contracts with less risk-averse entities (generally firms that can aggregate risk across many individuals). These contracts are welfare-enhancing because they increase the individual's utility. In exchange, the firm charges a premium. In an ideal world, the insurance contract is as comprehensive as possible, and all individuals are insured. Narrow networks move away from a comprehensive insurance contract by limiting the set of providers that a patient can visit to receive care. This is suboptimal, so one could naively propose a policy that forces insurers to cover all providers.[1] What would happen then? Insurers would be forced to accept whatever price hospitals want to charge. In turn, they would probably have to raise cost-sharing, as a means to keep spending down. This may hurts the value of the insurance contract even more than restricted networks do. The possibility to restrict access to certain providers then, can serve as a cost-containment measure. Lower costs mean lower premiums, and lower cost-sharing, which in turn expand access.

Narrow networks exist in multiple health care markets. We have already mentioned hospital networks. Pharmacy networks also exist, and can be used to lower the retail price of drugs. Access to drugs is itself also often restricted, through what are called formularies. Formularies are tiered lists of drugs covered by insurance plans. Most formularies exclude at least a few drugs. Occasionally, specific services may also be excluded from insurance coverage. For example, most health insurance plans exclude many cosmetic surgical interventions. Dental insurance plans are even more restrictive, and generally only cover a short list of procedures.

In these notes, I will focus on hospital networks, which is where the bulk of economic research is focused. We will discuss formularies when we talk about the pharmaceutical market. Research on service restrictions is extremely limited (within economics).

2 THE WELFARE IMPACT OF NARROW NETWORKS

I start by talking about the downside of narrow networks. The question we are trying to ask is simple (though the answer is anything but): how much insurance value do we lose by restricting networks?

2.1 Reduced-form evidence

We start from Dafny et al. (2017), a relatively simple reduced-form paper. The paper uses data on premium and characteristics of health insurance plans offered on the health exchanges of eight US states in 2014 and 2015. The main question they are asking is whether narrow networks con-

[1] This policy is not entirely naive: Medicare Part D regulation mandates that all drugs in certain protected classes are covered by insurers.

Dafny et al. (2017): Paper at glance

Positioning statement: Marketplace premiums in 2014-2016 were much lower than what the Congressional Budget Office projected. One possible explanation is that the plans offered on the ACA marketplaces covered fewer providers than expected.
Research question: To what extent are narrow networks responsible for the lower-than-projected premiums in the Affordable Care Act's health exchanges?
Answer: Narrow hospital and physician networks are associated with 16% lower premiums.

tributed to the lower-than-expected premiums for plans in the Affordable Care Act health exchanges, but our main interest in reading this paper is that it provides a simple way to assess the welfare loss from narrow networks. Intuitively, the difference in the monthly premium between a full network plan and a narrow network plan can give us a rough guess of the value that the average consumer places on upgrading the network. With that in mind, let's look at what the paper does.

First, some background. When the Affordable Care Act set up the health exchanges, the regulation required that qualified health plans maintain provider networks with sufficient number and types of providers to ensure that all services were accessible to enrollees. However, specific network coverage requirements did not exist, so when the Congressional Budget Office made projections for premiums, they did not have a clear idea of what networks were going to look like (and probably expected them to be more generous than they turned out to be).

Narrow networks can lead to lower premiums for two reasons. First, they can lower costs by excluding expensive and inefficient providers. Lower costs usually translate to lower premiums. Second, they can provide insurers with a powerful tool to negotiate lower prices: the threat of exclusion. An insurer that commits to a broad network will have limited options to replace a provider that is holding out for a high price. Conversely, an insurer with a limited network can simply replace holdouts using the large pool of excluded physicians.

The paper uses data on plan premiums and characteristics to study the association between network coverage and premiums. They measure coverage of hospitals and physicians separately, using a measure that places more weight on hospitals and physicians that treat a large volume of patients. They also control for a number of other observable plan characteristics, as well as geography (because premiums in the ACA exchanges can vary with location).

Their regression design does not allow them to claim a causal relationship between premium and network coverage. That's because there are all sorts of omitted variables that could affect the relationship. The most important unobservable variable that we might worry about here is that plans tend to be tailored to the population they target, so in areas with relatively narrower networks we expect patients to have a preference for those types of plans. Since premiums also depend on the characteristics of the population we have a classic case of omitted variable bias, where patient characteristics could influence both the outcome variable and the explanatory variable of interest.

That being said, correlational studies are still informative. They find that a plan with a narrow physician and hospital network is associated with a 16% lower premium. Narrow physician networks seem to matter more than narrow hospital networks, perhaps because most healthcare expenditure occurs in offices and outpatient settings (as opposed to hospitals, which are classified as inpatient settings). Their final numbers suggest that an increase from a narrow (70% coverage) to a full (100% coverage) hospital network is associated to an increase of $191 a year. An increase in the physician network from small (10%) to large (40%) is associated to a premium increase of $316 a year. They also track the implication of their results for the size of ACA subsidies, though this is less relevant for our topic.

2.2 A structural approach: Ho (2006)

To get a more accurate estimate of welfare we need to do a bit more work. One alternative to reduced-form analysis is to construct a full model of hospital and insurance choice to estimate how hospital networks affect the welfare of patients. Ho (2006) is the first paper that attempted this. The paper follows a three-step approach:

1. Estimate demand for hospitals using data on hospital visits from various markets.
2. Use the demand estimates to calculate the value of a given network of hospitals.
3. Estimate a model of plan choice using data on plan market shares, and using the estimates from 2. as an explanatory variable.

Conceptually, this is a relatively simple paper, though it is technically quite complex. I provide a brief overview of the model here.

To start, we specify a utility function for the utility of a patient i from visiting hospital h given a diagnosis l:

$$u_{ihl} = \delta_h + x_h \nu_{il} \beta + \varepsilon_{ihl} \tag{1}$$

Here we control for a fixed hospital quality component, δ_h, ν_{il} are observed patient characteristics, such as location and diagnosis. ε_{ihl} is an error term, which follows a Type 1 extreme value distribution. The probability that a patient will pick a certain hospital h is equivalent to the probability that

$$u_{ihl} = \max\left\{u_{ih'l}, h' \in H\right\}$$

where H is the set of all hospitals that patient i could have access to. The advantage of the Type 1 extreme value distribution is that it allows us to write this probability as

$$\rho_{ihl} = \frac{\exp\left(\delta_h + x_h \nu_{il} \beta\right)}{\sum\limits_{h' \in H} \exp\left(\delta_{h'} + x_{h'} \nu_{il} \beta\right)} \tag{2}$$

We can aggregate this probability across diagnoses to get a prediction for the number of patients going to a specific hospital (i.e. the hospital's market share):

$$s_h = \frac{1}{N} \sum_l \sum_{i=1}^{N_l} \rho_{ihl}$$

where N is the total number of patients, and N_l is the number of patients with diagnosis l. The assumption she makes is that all the variables are observed, so this can simply be solved using MLE (e.g. maximize the probability of observing the choices you see in the data, or some simple method of moments estimator (e.g. minimize the distance between observed and predicted market shares). Also notice that there is no price component in this utility function. This is an approximation that works as long as the cost of being admitted to any of the hospitals in the network is approximately the same.

The next step is to calculate the value of a network of hospitals. Intuitively, a patient will value a broader network if they worry about getting sick with a variety of diagnoses, and some hospitals are better at treating certain conditions. That's because for each diagnosis, the patient may have a different preferred hospitals that they want to go to. This is an expected value calculation, because the patient does not know in advance if they will need to go to a hospital. Once again, the properties of the Type 1 extreme value distribution allow us to write this expression analytically. Patient i's expected utility from the network of hospitals H_{jm} offered by plan j in market m is given by

$$EU_{ijm} = \sum_l p_{il} \mathbb{E}\left[\max_{h \in H_{jm}} u_{ihl}\right] = \sum_l p_{il} \log\left(\sum_{h \in H_{jm}} \exp\left(\delta_h + x_h \nu_{il} \beta\right)\right) \tag{3}$$

The final step is to take EU_{ijm} and use it as an explanatory variable in a model of health plan choice. This model is conceptually quite similar to the model of hospital choice. The utility of patient i from choosing plan j in market m is given by

$$w_{ijm} = \xi_{jm} + z_{jm}\theta + \gamma_1 EU_{ijm} + \gamma_2 \frac{P_{jm}}{y_i} + \omega_{ijm} \tag{4}$$

where z_{jm} are plan characteristics, P_{jm} is the premium, and ω_{ijm} is another Type 1 extreme value shock. In the utility the premium is specified as a fraction of a patient's income to capture some form of heterogeneity in price sensitivity. Exploiting the magical properties of Type 1 extreme value distributions once again we get an analytical function for market shares of each plan j in market m:

$$s_{ijm} = \sum_i \left(\frac{n_i}{n_m}\right) \frac{\exp\left(\xi_{jm} + z_{jm}\theta + \gamma_1 EU_{ijm} + \gamma_2 \frac{P_{jm}}{y_i}\right)}{\sum_{k \in J} \exp\left(\xi_{km} + z_{km}\theta + \gamma_1 EU_{ikm} + \gamma_2 \frac{P_{km}}{y_i}\right)} \tag{5}$$

Once again, this is something you can estimate, either using plan-market fixed effects to control for ξ_{jm}, or by using the BLP contraction mapping if the dataset is too small to back out a fixed

Ho (2006): Paper at glance

Positioning statement: Managed care health insurers in the US often restrict their enrollees' choice of hospitals to narrow networks. The welfare implications of this restriction are not well understood.
Research question: What is the welfare loss from restricted networks?
Answer: Across 43 markets in the US, restricted networks generate $1 billion in welfare losses (or $15.70 per person-year). The figure does not account for the potential benefit of narrow networks on prices.

effect.

The final calculations imply that the average patient loses about $15-16 a year because of network restrictions. This is not a very large number.

Final Notes:

- The paper only has hospital choice data on a subset of markets. That's problematic when we expand to other markets because we don't know the hospital fixed effects. So we need to do an adjustment using observables. Ideally you would be able to use hospital choice data from each market to do this well, but you work with what you have.
- The paper also does not include diagnosis-specific hospital quality. Later papers have shown that this could be important, at least in some contexts (see e.g. Ho and Pakes, 2014). I don't think it is possible to introduce these sort of unobservables in the estimation and still retain structural identification, so this is not really a criticism of the paper. However, unobserved diagnosis-specific quality could be an important component of horizontal differentiation between hospitals. In turn, horizontal differentiation is one of the things that makes network breadth so important (e.g. compare the value of a network that covers one of two hospitals when the two hospitals are identical, vs. when one hospital specializes in oncology but does not offer cardiology services while the other specializes in cardiology but does not offer oncology services). If you think the demand system is not doing a good job of capturing this differentiation you worry that the welfare cost may be underestimated.
- Ericson and Starc (2015), who run a similar exercise on a narrower market, but with more complete data, find results that seem much larger in magnitude (I say seem because the results refer to observed differences in networks, which need not be the same across the two settings), on the order of about one thousand dollars per year.

3 Network Formation and Bargaining

The previous section established that narrow networks are bad for consumer welfare. In this section we seek to understand the benefit of narrow networks. The intuition is simple. By threatening hospitals with exclusions, insurers can charge lower prices. Lower prices, in turn, promote access.[2]

In order to fully model how narrow networks may affect prices, one has to take into account a relatively complicated series of steps. A complete model of narrow networks and their impact on welfare and prices would have four stages:

1. Hospitals (indexed by h) and insurers, indexed by j,
   (a) determine networks of hospitals H_j for each insurance plan, and
   (b) negotiate prices $p_j = \{p_{jh}\}_{h \in H_j}$
2. Insurers (sometimes also called MCOs) choose an insurance contract (i.e. set premiums P_j and cost-sharing $c_j = \{c_{jh}\}_{h \in H_j}$)
3. Patients (indexed by i) choose an insurance plan
4. Patients get sick and choose hospitals (within their network's plan)

Each of this four stages involves a maximization problem, sometimes over multiple variables. Stage 1 is also usually split in 1(a) and 1(b), and not all papers model both aspects.

I first present a general model of these four stages, and then briefly describe a few papers in the field. Since primitives from the latter stages are used as inputs in the earlier ones, the most intuitive way to present the stages is in reverse order, starting with hospital demand.

Stages 3 & 4 (hospital and insurance demand) The hospital and insurance demand model is covered in my description of Ho (2006) in the previous section, so I will only cover a few small differences here. To the utility function of patient i in plan j for hospital h we add the cost-sharing component c_j:

$$u_{ijh} = \delta_h + x_h \nu_i \beta - \alpha c_j + \varepsilon_{ijh}$$

(we also simplify by removing indices related to diagnoses). where H is the set of all hospitals that patient i could have access to. Plan-specific demand for a given hospital can be written as

$$D_{jh}\left(H_j, c_j\right) = \sum_{i \in I_j} f_i \times \rho_{ijh}\left(H_j, c_j\right) \tag{6}$$

where I_j is the set of patients of plan j, and f_i is the probability that patient i gets sick. You can obtain overall demand for a hospital by summing across plans.

Utility of patient i from choosing plan j is

$$w_{ij} = \xi_j + \gamma_1 EU_{ij}\left(H_j, c_j, f_i\right) + \gamma_2 P_j + \omega_{ij}$$

[2]There are other potential reasons for exclusion that we don't touch upon. Insurers may want to exclude hospitals that are very expensive, or simply bad, or they may want to exclude hospitals that are very good and expensive in order to nudge very costly patients away from their health plan (see e.g. Shepard, 2020).

where we have removed market indices and the income normalization, and we have summarized plan characteristics with a fixed effect. $EU_{ij}\left(H_j, c_j\right)$ is the utility from the network, which now also depends on the cost-sharing parameters. To calculate $U_{ij}\left(H_j, c_j\right)$ you would extend the formula from Ho (2006) to incorporate cost-sharing:

$$EU_{ij}\left(H_j, c_j, f_i\right) = f_i \times \mathbb{E}\left[\max_{h \in H_j} u_{ijh}\right] = f_i \times \log\left(\sum_{h \in H_j} \exp\left(\delta_h + x_h \nu_i \beta - \alpha c_j\right)\right) \tag{7}$$

Finally, demand for plans is given by

$$D_j\left(H_j, p_j, c_j, f_i\right) = N \times \frac{\exp\left(\xi_j + \gamma_1 EU_{ij}\left(H_j, c_j, f_i\right) + \gamma_2 P_j + \omega_{ij}\right)}{\sum_k \exp\left(\xi_k + \gamma_1 EU_{ij}\left(H_k, c_k, f_i\right) + \gamma_2 P_k + \omega_{ik}\right)} \tag{8}$$

Notice that f_i here introduces the possibility of adverse selection.

Stage 2 (optimal premium setting) In this setting, health plans hold fixed their network and vector of hospital prices, and maximize profits by choosing the premium and the cost-sharing parameters. The problem takes the form of

$$\max_{P_j, c_j} D_j\left(H_j, p_j, c_j, f_i\right) \times P_j - TC_j\left(P_j, c_j\right) \tag{9}$$

where

$$TC_j\left(P_j, c_j\right) = \sum_{i \in I_j} f_i \sum_{h \in H_j} \rho_{ijh}\left(H_j, c_j\right) p_{jh} \tag{10}$$

In practice, no paper has ever allowed for endogenous cost-sharing. The most common assumption is that $c_{jh} = 0$ for all h. The only exception, as far as I know, is Gowrisankaran et al. (2015), which allows for an exogenously determined cost-sharing parameter c that is measured directly from claims data and not allowed to vary in counterfactuals. It is also possible to have premium-setting as a bargaining game, if patients are being represented by a larger entity, like an employer or a managed care organization.

Stage 1 (network formation and bargaining) This is the most complicated stage. Most papers split it in two by assuming that first the network is selected, and then prices are negotiated. Generally, the insurer is the one selecting the network. We follow this convention, and start by first writing down the bargaining problem holding H_j fixed.

The value of a given vector of prices p_j to the insurer (conditional on a network) can be expressed as

$$V\left(p_j; H_j\right) = D_j\left(H_j, p_j, c_j^{\star}\left(p_j; H_j\right), f_i\right) \times P_j^{\star}\left(p_j; H_j\right) - TC_j\left(P_j^{\star}\left(p_j; H_j\right), c_j^{\star}\left(p_j; H_j\right)\right) \tag{11}$$

where starred functions represent the optimal solutions to the Stage 2 problem.

Using this notation, we can write the Nash bargaining problem between each plan j and hospital h as

$$\max_{p_{jh}} \Bigg[V\left(p_{jh}, p_{j,-h}; H_j\right) - \underbrace{V\left(\infty, p'_{j,-h}; H_j \setminus \{h\}\right)}_{\text{Threat point of the plan}} \Bigg]^{b_{jh}} \times$$
$$\Bigg[D_{jh}\left(H_j, c_j^\star\left(p_j; H_j\right)\right) \times \left(p_{jh} - mc_j\right) + \sum_{k \in \mathcal{M}_h \setminus \{j\}} D_{kh}\left(H_k, c_k^\star\left(p_k; H_k\right)\right) \times \left(p_{kh} - mc_k\right)$$
$$- \underbrace{\sum_{k \in \mathcal{M}_h \setminus \{j\}} D_{kh}\left(H_k, c_k^\star\left(p'_k; H_k\right)\right) \times \left(p'_{kh} - mc_k\right)}_{\text{Threat point of the hospital}} \Bigg]^{1-b_{jh}}$$

What does all this notation mean?

- The two expressions contained in the underbraces represent the threat points of each agent.
- The prime (′) subscript allows for the possibility that plans and hospitals may re-negotiate prices in case of a disagreement. In practice, most papers do not allow for this possibility. Doing so substantially simplifies the math.
- $\mathcal{M}_h$ is the set of plans/MCOs that have struck a deal with h.

What about network formation? Intuitively, one could see how the optimal network would be the one that satisfies some notion of joint surplus maximization between hospitals and insurers (the joint surplus is what's split using prices). In theory one could solve for the optimal equilibrium conditional on all possible networks and pick the "best" one to be the equilibrium one. In practice that's very hard given the very large number of possible network configurations, and very few papers have attempted this. In the next section I discuss one of them.

3.1 A (nearly) complete model of network formation: Ho and Lee (2019)

The goal of Ho and Lee (2019) is to build a complete model to evaluate the impact of narrow networks. This is sort of the conclusion of the agenda that was started with Ho (2006), which only looked at the welfare effect. This paper looks at the complete picture, which includes the following effects of dropping a hospital from a network (see Figure 1 for some graphical intuition):

- Patient welfare loss, which leads to an inward shift of the demand function. All else equal, patients will leave the insurance plan to purchase a different plan, or stay uninsured.
- A reduction in the cost function, driven by (i) efficiency gains if the hospital was particularly expensive/inefficient, and (ii) selection of lower-cost enrollees (generally because the marginal enrollees tend to be higher-cost)

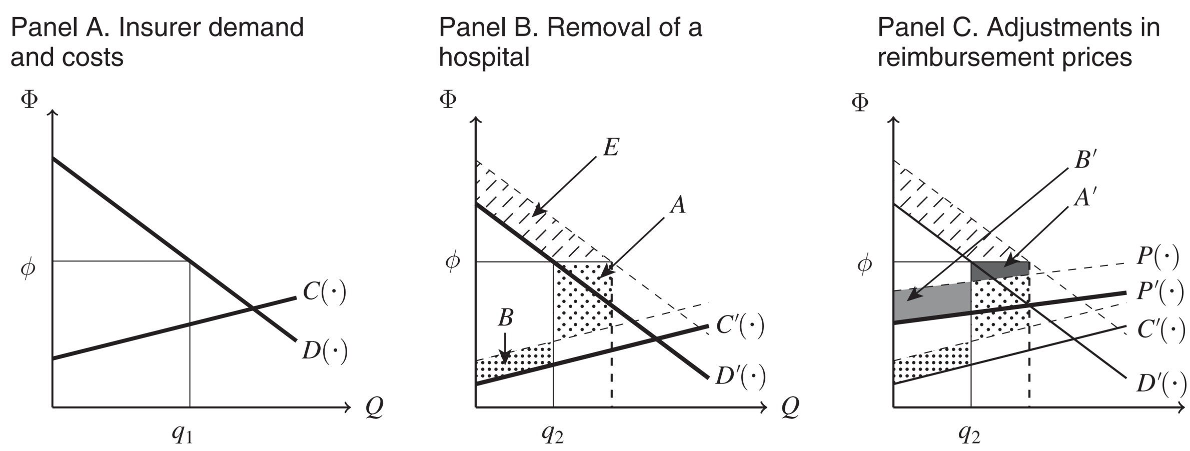

Figure 1. Removing a Hospital from an Insurer's Network

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

Figure 1: Removing a Hospital from an Insurer's Network (Ho and Lee, 2019, Figure 1)

- A change in the negotiated rates with other hospitals.

The paper's main contribution is in the way we think about this last effect. In the traditional Nash-in-Nash setup, the assumption is that failure to come to an agreement leads to the removal of a hospital from the network. This has relatively unintuitive implications:

- In equilibrium, exclusions only occur if there are no gains from trade (or if there is asymmetric information)
- If hospitals are substitutes (which they generally are, to some degree), broader networks imply lower prices, because the marginal contribution of each hospital is lower if there are more hospitals in the network.
- Increased competition leads to lower prices, but only through the channel above. In other words, if a new hospital enters the market, but is not included in a network, there is no impact on price.

This paper proposes a new equilibrium concept, which they call Nash-in-Nash with threat of replacement (NNTR). Under this equilibrium concept, the insurer can replace each hospital by threatening to replace them with any excluded hospital. The insurer can do so by making a take-it-or-leave-it offer to any one of the excluded hospitals. Under this new solution, excluding at least one hospital makes sense. Excluding more than one hospital can be beneficial if hospitals are sufficiently differentiated.

Theoretical setup

I modify the notation in the paper to be more consistent with the notation I have used so far. I also focus on the network-formation and price setting stage, and take all other primitives as given. This is consistent also with the expositional approach of the paper, which relies on primitives estimated by a previous paper (Ho and Lee, 2017). They take the network of the hospital as given and also condition on the networks and reimbursement rates of other MCOs.[3]

Let $\mathcal{H}_j$ denote the set of all possible networks of hospitals that the insurer could form. Each possible network $H_j \in \mathcal{H}_j$ is a set of hospitals. Each hospital h is included in the network of insurer j if and only if $h \in H_j$. Profits of hospitals and insurers are written respectively as

$$\pi_j^M\left(H_j, p_j\right) = \tilde{\pi}_j^M\left(H_j\right) - \sum_{h \in H_j}\left(D_{hj}^H\left(H_j\right) \times p_{hj}\right)$$
$$\pi_h^H\left(H_j, p_j\right) = \tilde{\pi}_h^H\left(H_j\right) + D_{hj}^H\left(H_j\right) \times p_{hj} + \sum_{n \neq j}\left(D_{hn}^H\left(H_j\right) \times p_{hn}\right)$$

Notice two key assumptions here:

1. The presence of a network-specific lump-sum component $\tilde{\pi}\left(H_j\right)$, and

2. Demand that only depends on the network, and not on the reimbursement prices p_{hj}

Denote the gains from trade between an insurer j and a hospital h as

$$\Delta_{hj}\pi_j^M\left(H_j, p_j\right) = \pi_j^M\left(H_j, p_j\right) - \pi_j^M\left(H_j \backslash h, p_{-hj}\right)$$
$$\Delta_{hj}\pi_h^H\left(H_j, p_j\right) = \pi_h^H\left(H_j, p_j\right) - \pi_h^H\left(H_j \backslash h, p_{-hj}\right)$$

where p_{-hj} is a vector of all prices other than p_{hj}, and the disagreement points $\pi_j^M\left(H_j \backslash h, p_{-hj}\right)$ and $\pi_h^H\left(H_j \backslash h, p_{-hj}\right)$ are computed assuming that all other agreements are formed at prices p_{-hj}. The total trade surplus is given by

$$\Delta_{hj}\Pi_{hj}\left(H_j, p_j\right) = \Delta_{hj}\pi_j^M\left(H_j, p_j\right) + \Delta_{hj}\pi_h^H\left(H_j, p_j\right)$$

Here the key implication of the assumption that demand for hospital is price inelastic is that the trade surplus does not depend on p_{hj}.

They define the *Nash-in-Nash with Threat of Replacement* (NTTR) prices as

$$p_j^\star = \left\{p_{hj}^\star\left(H_j, p_{-hj}^\star\right)\right\} = \min\left\{p_{ij}^{Nash}\left(H_j, p_{-hj}^\star\right), p_{hj}^{OO}\left(H_j, p_{-hj}^\star\right)\right\} \tag{12}$$

where $p_{ij}^{Nash}\left(H_j, p_{-hj}^\star\right)$ is the solution to the standard Nash bargaining problem, and $p_{hj}^{OO}\left(H_j, p_{-hj}^\star\right)$

[3] This is probably more reasonable in their setting than in others, because there are only three competitors. One has a complete network (which you can interpret as a fixed choice), and the other is Kaiser Permanente, which is a vertically integrated insurer-provider system. You could quibble with the assumption that rates are fixed, but the bottom line with these complicated structural papers is that multiple-agent games are essentially intractable at this time.

is defined as the solution to

$$\pi_j^M\left(H_j, \left\{p_{hj}^{OO}\left(\cdot\right), p_{-hj}^{\star}\right\}\right) = \max_{k \notin H_j} \tilde{\pi}_j^M\left(\left(H_j \backslash h\right) \cup k, \left\{p_{kj}^{res}\left(H_j \backslash h, p_{-hj}^{\star}\right), p_{-hj}^{\star}\right\}\right) \tag{13}$$

and $p_{kj}^{res}\left(H_j \backslash h, p_{-hj}^{\star}\right)$ is hospital k's reservation price of being added to the network of insurer j, and defined as the solution to

$$\pi_k^H\left(\left(H_j \backslash h\right) \cup k, \left\{p_{kj}^{res}\left(H_j \backslash h, p_{-hj}^{\star}\right), p_{-hj}^{\star}\right\}\right) = \pi_k^H\left(H_j \backslash h, p_{-hj}^{\star}\right) \tag{14}$$

Some intuition for these iterative equations. The assumption is that the insurer can go out of network and make TIOLI offers to excluded hospitals. This essentially provides a "floor" price for the insurer: they will not pay a price above what would make them switch to one of these excluded hospitals.

This setup leads to a two-part pricing equation:

$$p_{hj}^{\star}\left(\cdot\right) = \begin{cases} p_{hj}^{OO}\left(H_j, p_{-hj}^{\star}\right) & \text{if } \tau \times \Delta_{hj}\Pi_{hj}\left(H_j, p_j\right) < \max_{k \notin H_j} \Delta_{hj}\Pi_{hj}\left(\left(H_j \backslash h\right) \cup k, p_{-hj}^{\star}\right) \\ p_{ij}^{Nash}\left(H_j, p_{-hj}^{\star}\right) & \text{otherwise} \end{cases} \tag{15}$$

In other words, the exclusion threat binds whenever there is another hospital $k \notin H_j$ whose inclusion (as a replacement for j) generates an additional surplus that exceeds τ-share of the bilateral gains from trade generated by hospital j (τ here is the bargaining power of the insurer). Another way of restating this intuition is that as long as a hospital does not have very close competitors, the equilibrium will still be Nash-in-Nash. If the hospitals has close competitors, the price will decline below the Nash-in-Nash equilibrium. In an extreme case, where, say, two identical hospitals are sitting on opposite sides of the street, the insurer will essentially be able to extract marginal cost pricing from either one of them.

One possible concern with this new equilibrium concept is that a hospital may want to unilaterally reject this outcome. This is not something you have to worry about in the usual Nash-in-Nash model: it will only happen if the price is below the Nash outcome (i.e. if the price is $p_{hj}^{OO}\left(H_j, p_{-hj}^{\star}\right)$). So to complete the definition of an equilibrium, the authors include a requirement that the network be "stable". In other words:

$$\Delta_{hj}\pi_j^M\left(H_j, p_j^{\star}\right) \geq 0 \,\forall\, h$$

There is also a similar requirement for j, though (in my reading of the paper) that appears to be less crucial.

To conclude the derivation, they define the *insurer-optimal stable network* as the stable network that maximizes the insurer's profits at NNTR prices $p_j^{\star}$:

$$H_j^{\star} = \arg\max_{H_j \in \mathcal{H}_j} \pi_j^M\left(H_j, p_j^{\star}\right) \tag{16}$$

Estimation

The data covers the 2004 enrollees of the CalPers system, and includes: (i) health plan enrollment, (ii) claims, and (iii) hospital admissions. They focus on the strategic choices of Blue Shield, which is one of three large insurers in the market (the others are Blue Cross and Kaiser Permanente, which they disregard because their networks appear largely fixed, for a variety of reasons). The estimation relies largely on the work of a previous paper (Ho and Lee, 2017), which we will cover at a later date. I cover the basics here:

Hospital demand: This is based on Ho (2006). Groups are age-sex categories (ten in total), and there are six possible diagnoses. Demand does not respond to negotiated rates. Another necessary assumption is that there is no selection across insurance plans on unobservable consumer preferences for hospitals.

Demand for insurance plans: This is also very similar to Ho (2006). They calculate the willingness to pay for a network and use it in the utility function of a plan.

Premium setting: Premiums here are set through negotiations between employers and insurers. Employers maximize patient welfare (from the demand stage above) minus premium payments, while insurers maximize premium payments minus expected cost (which come from the hospital demand stage). There is no specific mention of adverse selection here, though some adverse selection concerns should be implicitly channeled through the different admission probabilities of various patient categories (see e.g. Table III of Ho and Lee, 2017). Occasionally, adverse selection can imply non-existence, or multiplicity of equilibria. This concern is not mentioned here.

The impact of narrow networks

The empirical contribution of the paper is not in the estimation, but rather in a simulation exercise. The simulation exercise uses the NNTR model to identify equilibria (including the optimal network) satisfying various criteria (highest insurer surplus, higher overall surplus, etc.). The simulations accomplish various goals. Here, we are most interested in what they can tell us about the tradeoff between full networks and limited networks. In this regard, we want to pay close attention to the tradeoff between consumer surplus and reduced rates.

To run the simulation, they make some simplifying assumptions. As usual, they hold everything but the Blue Shield network and rates as fixed. They run the simulation by Hospital Service Area (HSA, a measure of a hospital market). They have a total of 12 HSAs. I believe this is largely to reduce the size of the set of all possible networks (choosing the network is a discrete maximization problem, so they have to check the surplus of every possible network configuration, which is computationally burdensome). For the same reason (probably), they focus only on negotiations with the five larges systems in each market.

Table 1: Simulation results (Table 1 from Ho and Lee, 2017)

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

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

Table 1 shows the results of the simulation. We focus here on two columns. The first one (Social, NNTR), and the third one (Blue Shield, NNTR). All changes listed in the Table are relative to the numbers on the right-most column, which represents the base case (where Blue Shield's network is exogenously set to be complete).

First, consider the first column, which shows the socially optimal network under the NNTR model. We start from the bottom. Under the model, the optimal network is *not necessarily* complete. Instead, half the market have an incomplete network. This allows Blue Shield to extract lower rates from hospitals (payments decline by 5.6%), which translate to lower premiums (0.6%, or about $16). The fall in premiums more than compensates enrollees for the decline in the quality of the network: patient welfare increases by $12 on average (all numbers are yearly). Overall welfare also increases. These results are encouraging, but not necessarily surprising, since what we are doing is essentially adding a degree of freedom to the problem: instead of forcing Blue Shield to have a complete network, we are allowing it to vary.

A more interesting comparison is with the third column. This is the network that we would expect to arise if the market was perfectly competitive.[4] In this situation, Blue Shield is more aggressive in excluding hospitals. This is likely because the insurer benefits from a reduction in rates (these rates are neutral in the problem of the social planner since they represent a transfer between firms). This turns out to benefit consumers even more than in the social planner case. Welfare increases by $20 per year, thanks to a reduction in premiums by over $30. Overall surplus however, falls, due to the reduction in hospital profits. Indeed, exclusions do not seem to harm consumers all that much in these simulations. Column 2, which shows results from the network that maximizes consumer surplus, results in the most restricted networks, and has the largest reduction in premiums. This result is not necessarily intuitive (one would think that if Blue Cross could generate higher consumer surplus by removing hospitals and lowering premiums, it could also increase premiums just enough to capture some of that extra surplus). The authors do not present this explanation directly, but my interpretation is that the jump in consumer surplus is largely driven by customers of other insurers (Blue Cross and Kaiser), whose premiums also go down because of competition with Blue Shield. Those consumer welfare gains are quite prominent in the calculations of column two because they are "free" (the networks of Blue Cross and Kaiser have not gotten worse). However, they do not factor very much in the first column (premiums are also transfers to a first-order approximation), and not at all in the third.

Some notes

- One of the first reactions you might have to the structure of the outside offers is that being able to make a TIOLI offer is very convenient from the point of view of the insurer. Perhaps even too convenient, and resulting in artificially lower prices. It's a reasonable reaction, and it's helpful to think through possible alternatives.
  - One possibility is the reverse. Should we let hospitals make TIOLI offers? This may seem appealing, but if you think about it, does not appear like a feasible alternative. Why? If excluded hospitals can make TIOLI offers, they will make the insurer indifferent between the Nash-in-Nash price and selecting another hospital. This leaves us

[4] As a side note, this is *not* necessarily the network that arises in the data. You could find this problematic: does the model lack some feature that would otherwise allow it to fit the data perfectly? I believe the authors argue that Blue Shield has to satisfy certain network requirements, so the market is not at a competitive equilibrium (Blue Shield, who has a restricted network, had to ask for permission to exclude a list of hospitals, and was only allowed to exclude a subset of the submitted list).

with a fairly unsatisfying final equilibrium: the hospital will choose the Nash-in-Nash option, leaving us with exactly the same model as before.

– Another possibility is something in between the two TIOLI options. This is likely more realistic than what the authors do, but can create very problematic estimation issues. To see why, imagine that we set up the outside option as an additional bargaining problem, where the insurer negotiates with the originally excluded hospital using some form of Nash bargaining. The issue with this new maximization problem is the outside option. What is it? If this negotiation fails, where can the insurer do? The most natural option might be to go to a potential second excluded hospital as an alternative. This sets up a very complicated series of iterated maximization problems, which will almost certainly be computationally unfeasible (not to mention incredibly convoluted). One could try to achieve the same result by simply imposing an exogenous split of surplus between the insurer and the new hospital. However, there is probably no real reason to prefer an ad-hoc assumption about surplus split to the TIOLI assumption that the authors make.

- The requirement of stability does not rule out that additional links would not be profitable. In other words, after setting $\left(H_j, p_j^\star\right)$, it is possible that the insurer may find it profitable to strike an agreement with an additional hospital. Ghili (2020) criticizes this aspect of the model and addresses it in his own (contemporaneous) paper on network formation. He points out that this allows the insurer to commit to non-credible exclusions to extract lower prices. The chief criticism focuses on an example where two non-competing hospitals are pitted against one another (see the discussion in Section 3.3 of that paper). Ho and Lee (2019) do acknowledge the limitation (footnote 32), but argue that it is unreasonable to assume that two-agent deviations (both the insurer and the hospital have to agree to the addition of an extra hospital) would leave the rest of the agreement unchanged. Given that these are relatively recent papers (Ghili in particular is still unpublished) it may be a while before a clearer picture of the theoretical implications emerges.

- Shepard (2020) finds evidence that selection of patients across insurance plans on unobservable preference for hospitals does exist, though in a different setting. The real question here is not whether the selection exists (it almost definitely does), but whether it is large enough to really be problematic for the estimation.

3.2 Brief history/overview of other papers on the topic

Ho and Lee (2019) is, in a way, the culmination of a strand of literature started by Ho (2006) (other papers had asked similar questions about narrow networks, but none had attempted the structural route before). Over these fifteen or so years, several papers made contributions and pushed the boundary of what was possible. Also, parallel to Ho and Lee (2019), two other papers provided solutions to the question of network formation. Here I provide a brief history of these other papers, and what they contributed.

- Ho (2009): estimates a full model of hospital network formation, but models stage 1 as a game where hospitals make TIOLI offers to insurers, who then decide whether to include them in networks or not. Narrow networks may arise because some hospitals make very high offers (these hospitals have limited capacity, or are so high quality that they can fill up demand even if they don't contract with all insurers).

- Gowrisankaran et al. (2015): estimates a bargaining model between hospitals and MCOs (basically insurers that provide services for specific employers). The main model has only stages 1(b) and 4. The assumptions that make this work are that stage 3 is trivial—patients are employees and take the employer plan as given—and premiums don't exist, so stage 2 is irrelevant—the employer/MCO simply maximizes employee welfare. They also report calibrated results from a model that adds stages 2 and 3. Finally, they investigate the impact of coinsurance rates but only in counterfactual (there is no section of the model where the MCO chooses the coinsurance optimally). The main value of this paper, relative to others in the literature is that they examine the role of coinsurance rates (though they hold them fixed instead of allowing firms to set them endogenously).

- Lewis and Pflum (2015) explores the possibility that joining a hospital system may also change a hospital's ability to bargain (i.e. its Nash Bargaining parameter). In general, bargaining parameters are essentially estimated as residuals. In this specific case, the paper uses hospital revenue and cost data and attempts a more economically rigorous estimation of this parameters. According to their findings changes in bargaining ability after joining a hospital system can explain most of the subsequent increase in prices.

- Ho and Lee (2017): they estimate stages 1(b) to 4, holding the network fixed (i.e. stage 1(a)). Their counterfactuals of interest are about insurer mergers.

- Ghili (2020); Liebman (2018): like Ho and Lee (2019), they estimate all four stages. They each have a different assumption about how to model the threat point in part 1(b). Ghili (2020) works his way up from theoretical requirements, which crucially include the assumption that the "threat" of replacement has to be incentive-compatible (e.g. if a hospital is replaced, the insurer does not then have an incentive to still contract with the now excluded hospital). This may be more realistic because it ensures that hospitals used as threats have to be substitutes. However, it requires some additional assumptions. In Liebman (2018), the threat is a random hospital among the excluded ones.

4 BIBLIOGRAPHY

Dafny, Leemore S., Igal Hendel, Victoria Marone, and Christopher Ody, "Narrow Networks on the Health Insurance Marketplaces: Prevalence, Pricing, and the Cost of Network Breadth," *Health Affairs*, 2017, *36* (9), 1606–1614.

Ericson, Keith M. Marzilli and Amanda Starc, "Measuring Consumer Valuation of Limited Provider Networks," *The American Economic Review: Papers and Proceedings*, 2015, *105* (5), 115–119.

Ghili, Soheil, "Network Formation and Bargaining in Vertical Markets: The Case of Narrow Networks in Health Insurance," 2020.

Gowrisankaran, Gautam, Aviv Nevo, and Robert Town, "Mergers when prices are negotiated: Evidence from the hospital industry," *American Economic Review*, 2015, *105* (1), 172–203.

Ho, Kate and Ariel Pakes, "Physician payment reform and hospital referrals," *American Economic Review*, 2014, *104* (5), 200–205.

_ and Robin S. Lee, "Insurer Competition in Health Care Markets," *Econometrica*, 2017, *85* (2), 379–417.

_ and _ , "Equilibrium provider networks: Bargaining and exclusion in health care markets," *American Economic Review*, 2019, *109* (2), 473–522.

Ho, Katherine, "The welfare effects of restricted hospital choice in the US medical care market," *Journal of Applied Econometrics*, nov 2006, *21* (7), 1039–1079.

_ , "Insurer-provider networks in the medical care market," *American Economic Review*, 2009, *99* (1), 393–430.

Lewis, Matthew S. and Kevin E. Pflum, "Diagnosing hospital system bargaining power in managed care networks," *American Economic Journal: Economic Policy*, 2015, *7* (1), 243–274.

Liebman, Eli, "Bargaining in Markets with Exclusion: An Analysis of Health Insurance Networks," 2018.

Shepard, Mark, "Hospital Network Competition and Adverse Selection: Evidence from the Massachusetts Health Insurance Exchange," 2020.