Markov-Perfect Network Formation

An Applied Framework for Bilateral Oligopoly and Bargaining in Buyer-Seller Networks *

Robin S. Lee Kyna Fong

September 2013

Abstract

We develop a tractable and applicable dynamic model of network formation with transfers in the presence of externalities. Our primary application is the estimation of primitives and prediction of counterfactual outcomes in bilateral oligopoly and buyer-seller networks. The framework takes as primitives each agent's static profits, and provides the equilibrium recurrent class of networks and negotiated transfers. Importantly, agents anticipate future changes to the network, and links may be costly to form, maintain, or break. We detail the computation and estimation of a Markov-Perfect equilibrium in this environment, and highlight the approach using simulated data motivated by insurer-hospital negotiations. We explore the impact of hospital mergers on negotiated payments, insurer premiums, and consumer welfare, and demonstrate how accounting for dynamics yields substantively different predictions than traditional static approaches.

Keywords: Bilateral oligopoly, buyer-seller networks, bargaining, network formation
JEL: L13, L14, I11, C78, D85

*Lee: NYU Stern School of Business, Department of Economics, 44 West 4th St Ste 7-78, New York, NY 10012, rslee@stern.nyu.edu. Fong: Elation EMR, NBER, and Stanford University, 550 15th St Ste 27, San Francisco CA 94110, kyna.fong@gmail.com. We thank Victor Aguirregabiria, John Asker, C. Lanier Benkard, Adam Brandenburger, Allan Collard-Wexler, Ignacio Esponda, Igal Hendel, Rachel Kranton, Sarit Markovich, Ariel Pakes, Andy Skryzpacz, Lars Stole, Michael Whinston and seminar participants at Chicago Booth, Columbia, Duke, MIT, Northwestern, NYU Stern, Rochester, UCLA, U. Toronto, U. Warwick, Yale, the CEPR/JIE Applied IO Conference, the Hanyang Summer Microeconomics Workshop, the SITE Summer Workshop, the UBC Summer IO Conference, and the Utah Winter Business Economics Conference for helpful comments. All errors are our own.

1 Introduction

Network formation—the creation of relationships, trading partnerships, or *links*—is a general phenomenon underlying many areas of economic exchange and interaction, including buyer-seller networks and bilateral oligopoly. In these settings, externalities are pervasive as the value of a relationship between two parties is often influenced by relationships involving others. Though networks have attracted a great deal of interest and research, few models have yielded sharp predictions about equilibrium network structures and the division of surplus when there are externalities; additionally, none to our knowledge have analyzed dynamic settings with bargaining where relationships are constantly renegotiated and can be formed, broken, and reformed over time. As a result, our ability to estimate primitives or predict counterfactual outcomes in these environments has been limited.

This paper provides a tractable applied framework for the analysis of bilateral oligopoly that can be taken to data, which allows for the computation and estimation of both equilibrium *surplus division* and *network formation* when agents engage in repeated interaction. Our primary application is in industrial organization settings with bilateral contracting between oligopolistic firms in vertical markets: e.g., negotiations between manufacturers and retailers, health insurers and medical providers, content providers and distributors, and software developers and hardware providers.[1] By predicting who contracts with whom, our model is also useful for counterfactual analysis in settings where trading or contracting relationships between firms might change as a result of a policy intervention, merger, or other industry shock. We also demonstrate how a dynamic model can identify both agents' "bargaining power" and contracted transfers—objects which are often unobserved in applied work—and detail an estimator that can recover them by using observed equilibrium actions and networks.

We consider an infinite horizon, dynamic, discrete time network formation game between multiple agents with endogenous transfers, and do not impose restrictions on the set of feasible networks that may arise. At the beginning of each period, agents simultaneously announce the set of partners with whom they would like to bargain; if two agents both choose each other, they are then able to negotiate with one another in that period and potentially form a relationship. This determination of a period's "negotiation network" is a variant of the strategic network formation game proposed by Myerson (1991). Next, all agents who can negotiate with one another engage in simultaneous bilateral Nash bargaining to determine *period contracts*, where disagreement at this stage results in termination of the relationship in that period.[2] Finally, agents receive period payoffs as a function of the realized *network structure* and negotiated period contracts. We assume these payoffs are primitives of the analysis that may derive from some underlying subgame played among agents; as they are a function of the entire network, the model admits any general form of externalities among agents.

[1] The model can also be extended to other general network settings with transfers: e.g., horizontal mergers or alliances, firm-worker negotiations, and so forth.

[2] Period contracts may represent wholesale prices, capitation rates, carriage fees, royalty payments, etc.

The key feature of our model is that agents' disagreement points from any bilateral bargain are endogenously determined and are internally consistent with respect to the complete infinite horizon game. These disagreement points contain the continuation values of moving to a network state without that link; this captures the idea that agents repeatedly renegotiate (albeit perhaps with some delay) while anticipating future changes to the network. Furthermore, dynamics are introduced through both repeated interaction and the possibility that links may be costly to form, maintain, or break—i.e., the cost of creating, keeping, or terminating a relationship with another agent may depend on whether or not that link existed in the previous period—and may lead to persistence of a given network over time.

In light of the complexity of the problem, the focus is on tractability and computability so that the model can be taken to data. We restrict attention to Markov-Perfect equilibria (MPE) (Maskin and Tirole, 1988) in which agents' strategies (i.e., with whom they wish to negotiate in a given period) are a function of the previous network structure or *state*. We provide assumptions that guarantee the existence of an MPE, and provide a means of computation.

We highlight the importance of dynamics in several ways. In our model the division of surplus within a given network depends on potential payoffs in all potential networks and not only strict sub-networks (which is the often case in other models of bargaining and surplus division); failing to account for this, or relying only on static Nash conditions, can affect predictions regarding transfers and other underlying fundamentals of interest. Through a simple example, we show our model admits the possibility that the "short-side" of the market in a bilateral oligopoly setting can achieve full rent extraction; this prediction cannot easily be generated by standard static bargaining models without additional assumptions.

Additionally, dynamics introduce a source of identification that is not present in static settings: agents' "bargaining powers," represented in our model by Nash bargaining parameters, directly affect agents' outside options and value functions, and influences both which networks are sustainable in equilibrium and the steady state distribution over those networks. This allows for the estimation of both bargaining power and contracted transfers using only information on agents' static profits, discount factors, and observed networks. This is in contrast to previous empirical work in bilateral contracting, which, in order to estimate parameters of interest, has also required making assumptions on the allocation of bargaining power (e.g., one side makes take-it-or-leave-it offers) or knowledge of contracted transfers.[3]

We apply our framework to a stylized bilateral contracting game between hospitals and health maintenance organizations (HMOs) using simulated data; this environment and the determinants of negotiated transfers between medical providers and insurers has been the focus of recent research given its central role in healthcare policy (Town and Vistnes, 2001; Capps, Dranove and Satterthwaite, 2003; Sorensen, 2003; Ho, 2009; Gowrisankaran, Nevo and Town, 2013; Ho and Lee, 2013). In our example, period payoffs are generated from an underlying demand system for insur-

[3]E.g., Ho (2009), Crawford and Yurukoglu (2012), Grennan (2013) Using a model such as Stole and Zweibel (1996) where agents' outside options include the renegotiation of all other links also would allow the set of equilibrium networks to be a function of bargaining power.

ers and providers, and we compute MPE networks and transfers for a large number of simulated markets. We find the complete or efficient network need not be visited in equilibrium, and that the total number of equilibrium network structures remains small even as the number of agents and potential network states increase. Using the same market-level primitives, we compare the equilibrium predictions of our dynamic model with those of a static bargaining model; we show the specification of bargaining power between agents matters greatly, and that incorporating dynamics yields different predicted equilibrium transfers than a static model. Furthermore, we illustrate how agents' bargaining powers can be estimated from observing equilibrium play in a single market.

Finally, contributing to the study of horizontal merger impacts on input prices in vertical markets (Horn and Wolinsky, 1988; Inderst and Wey, 2003), we use our model to study the effect of hypothetical hospital mergers on negotiated transfers and profits while crucially controlling for post-merger changes in network structure. We find that although mergers generally lead to higher premiums and fewer patients served, their effect depends on the division of bargaining power between firms. Furthermore, in certain circumstances, gains to merged hospitals in the form of higher negotiated payments are occasionally offset by an inability to direct patients away from utilizing higher cost, less efficient hospitals. Static merger analysis fails to capture the incentives for hospitals to merge, predicting most hospital mergers would lead to lower hospital profits.

Related Literature. Our paper builds on and contributes to the literatures on strategic network formation, dynamic bargaining, and contracting with externalities. As these literatures are vast, we recognize that certain papers and areas of research may have been unintentionally omitted below. However, it is worth emphasizing that our paper's primary contribution is not theoretical; rather, it is to propose a tractable and feasible framework that can be taken to data and can adequately capture important features and details of real-world industries.

There is a large literature on non-cooperative models of network formation in static environments both with and without transfers (e.g.., Jackson and Wolinsky (1996); Kranton and Minehart (2001); Bloch and Jackson (2007); Elliot (2013); see Jackson (2004) for a survey). In these static papers, often the primary focus is on the existence of stable and efficient networks, and they vary to the degree in which variable networks, heterogeneous agents, and externalities are admitted.[4] Papers on dynamic network formation are less numerous, and include those that assume agents are myopic (e.g., Watts (2001); Jackson and Watts (2002); Galeotti, Goyal and Kamphorst (2006)), and those that assume agents are farsighted but cannot engage in side payments (e.g., Dutta, Ghosal and Ray (2005)). Other work on surplus division in dynamic networked environments include Manea (2011) and Abreu and Manea (2013*a*,*b*), which focus on the division of unit surplus between networked agents, and papers on dynamic coalition formation games (e.g., Gomes (2005)). These dynamic papers crucially restrict how the network can change in that the network is either fixed, or can only either shrink or grow over time. More similar to our model is Gomes and Jehiel (2005), which examines a dynamic setting in which a random proposer (as opposed to every agent) each

[4]See also Corominas-Bosch (2004); Polanski (2007); Melo (2013) for models of bargaining in fixed or exogenous networks.

period can make transfer offers to others in order to move between states. Our use of bilateral Nash bargaining in a changing environment also relates to stochastic bargaining models (Binmore (1987); Merlo and Wilson (1995); c.f. Muthoo (1999)), whereby in our setting the current network affects payoffs and surplus to be divided.

Most previous work on contracting with externalities focuses on settings in which there is a single agent on one side the market (Cremer and Riordan, 1987; Horn and Wolinsky, 1988; Hart and Tirole, 1990; McAfee and Schwartz, 1994; Segal, 1999; Segal and Whinston, 2003). Exceptions are primarily static models, and include Prat and Rustichini (2003), which limits the nature of contracting externalities; and Inderst and Wey (2003) and de Fontenay and Gans (2013), which address surplus division within fixed networks. These latter two papers also provide non-cooperative foundations for allocation rules based on Shapley (1953) or Myerson (1977) values extended to networked settings. Jackson (2005) discusses limitations of applying such solution concepts to environments where both links and contracts are bargained over; similarly, we do not rely on these cooperative based allocations of value as they do not provide guidance over which networks arise in equilibrium, nor allow agents to anticipate changes to the existing network (other than dissolving into subnetworks) when bargaining over both links and allocations. Furthermore, due to the applications we have in mind where the contract space is often limited and there are contracting externalities, we do not wish to presume efficiency (which often is closely tied to these concepts) in our analysis.

Close to the spirit of our paper is Dranove, Satterthwaite and Sfekas (2008), who also model insurer-medical provider negotiations, and compares predictions from a "naive" static bargaining model and a more sophisticated bargaining model motivated by Stole and Zweibel (1996) in which agents anticipate the future reactions of other agents. Dranove, Satterthwaite and Sfekas (2008) allow for agents to have different "levels of rationality" in bargaining, and assume agents are limited in their ability to anticipate future adjustments subject to a network change. Our work is complementary in that it does not require all agents to immediately renegotiate upon a network change, and solves for an exact MPE in a fully dynamic game.

Finally, although this paper aims to provide a tractable dynamic foundation for the empirical study of network formation, we recognize that it is a highly stylized model and provide the explicit caveat that certain applications may be ruled out given our assumptions. Our reliance on simulation is a result of the complexity and intractability of general network formation games in a dynamic context. In this light, our approach can be seen as in the spirit of the literature on industry dynamics (Pakes and McGuire (1994); Ericson and Pakes (1995); Doraszelski and Pakes (2007)). Furthermore, we build off several results in the literature on the estimation of dynamic games— particularly Hotz and Miller (1993), Hotz et al. (1994), Aguirregabiria and Mira (2007), and Benkard, Bajari and Levin (2007)—for the computation and estimation of our model.

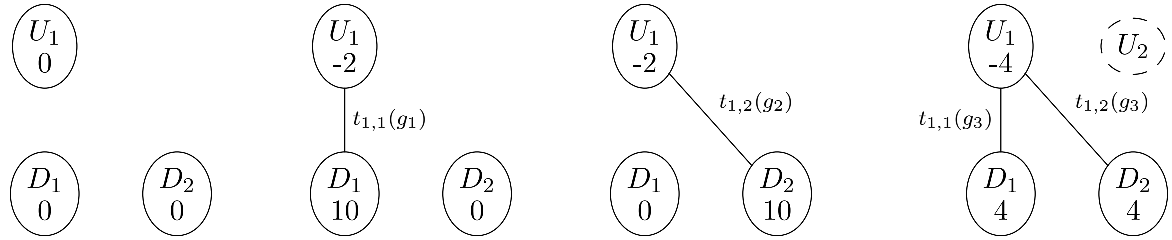

(a) g_0 (b) g_1 (c) g_2 (d) g_3

Figure 1: Potential Networks g_0, g_1, g_2, g_3 between firms U_1, D_1, D_2. Period payoffs contained within circles; $t_{ij}(g_k)$ represents payment between U_i and D_j under network g_k.

2 A Stylized Example

We first provide a simple example of bilateral contracting among firms which highlights the main innovations and objectives of this paper while clarifying the differences between a standard static analysis and our dynamic analysis. In particular, we emphasize why the division of surplus between agents for a given network fundamentally requires an understanding of surplus division in all other potential networks.

Consider a repeated setting with 1 upstream firm U_1 and two downstream firms D_1 and D_2. D_1 and D_2 compete for consumers, but must obtain inputs from U_1 to do so. Links between firms represent trading relationships between firms: i.e., if a link exists between U_1 and D_1, U_1 agrees to supply D_1 for a given period. Disregarding any payments between firms, the set of period revenues that would accrue to each firm for any set of links—i.e., a network—is given in Figure 1, and are taken to be a primitive of the analysis. Note that industry profits are maximized if U_1 supplies exclusively to D_1 or D_2; if U_1 supplies to both downstream firms, total industry profits are reduced (perhaps as a result of downstream competition). Hence, there are contracting externalities.

We assume that in each period, given a set of existing links defining a network g, U_1 negotiates with any downstream firm j with which it is linked a lump-sum payment $t_{1,j}(g)$ that is paid to U_1. Hence, if the current network was g_3, total period profits to each firm would be $\pi_{U_1}(\mathbf{t}(g_3)) = -4 + t_{1,1}(g_3) + t_{1,2}(g_3)$, $\pi_{D_1}(\mathbf{t}(g_3)) = 4 - t_{1,1}(g_3)$ and $\pi_{D_2}(\mathbf{t}(g_3)) = 4 - t_{1,2}(g_3)$, where $\mathbf{t}(\cdot) \equiv \{t_{1,1}(\cdot), t_{1,2}(\cdot)\}$.

Given these primitives, consider the following two questions: (i) for any given network, what is the division of surplus between agents (i.e., what are the values of $t_{i,j}(\cdot)$), and (ii) which network(s) do we expect to arise in equilibrium?

With regards to the first question, as there is only one upstream firm in the example, one might take the stance that U can offer take-it-or-leave-it offers (Hart and Tirole, 1990; Segal, 1999); alternatively, one might assume downstream buyers could make competing offers as in Bernheim and Whinston (1998). Choosing either extreme—an *offer game* or *bidding game*, using the terminology of Segal and Whinston (2003)—may be difficult to motivate in certain applications; and this approach does not easily generalize when there are multiple agents on both sides (e.g., if an additional upstream firm U_2 enters the market as depicted in Figure 1(d)). This paper will nest both possibilities by assuming firms engage in *simultaneous* bilateral Nash bargaining in each period: e.g., in network g_3, the negotiated transfer $t_{1,2}(g_3)$ will maximize the asymmetric Nash product of U_1 and D_2's gains from trade *given* U_1 *and* D_1*'s trade* $t_{1,1}(g_3)$:

$$t_{1,2}(g_3) \in \arg\max[\pi_{U_1}(\mathbf{t}(g_3); g_3) - \widetilde{\pi_{U_1}}(D_2; g_3)]^{b_{U_1}} \times [\pi_{D_2}(\mathbf{t}(g_3); g_3) - \widetilde{\pi_{D_2}}(U_1; g_3)]^{b_{D_2}} \quad (1)$$

where $\widetilde{\pi_k}(j; g)$ represents agent k's outside option or disagreement point from failing to come to an agreement with agent j in network g, and b_k represents agents k's Nash bargaining parameter. This setup embeds the possibility that agents make take-it-or-leave-it offers (by setting $b_k = 0$ for one agent), or any intermediate surplus division.[5] Nonetheless, regardless of the particular bargaining protocol used, there is still a need to specify each agent's outside option ($\widetilde{\pi_k}$).

In static bargaining environments, two approaches have been widely used. The first approach assumes agreements between other agents are binding given disagreement and do not change (e.g., as in Cremer and Riordan (1987); Horn and Wolinsky (1988)).[6] Under this assumption in the current example, the outside option for D_2 if it failed to come to an agreement with U_1 in g_3 would be D_2's profits under g_1, or $\widetilde{\pi_{D_2}}(U_1; g_3) = 0$, and U_1's outside option would be what it would get under g_1 *plus its negotiated transfer from* D_1 *in* g_3: i.e., $\widetilde{\pi_{U_1}}(D_2; g_3) = -2 + t_{1,1}(g_3)$. This is similar to Nash equilibrium reasoning in that agents do not anticipate changes to other bargains. Under this assumption, if all agents had equal Nash bargaining parameters, one could solve (1) (and the parallel equation for $t_{1,1}(g_3)$) to obtain $t_{1,1}(g_3) = t_{1,2}(g_3) = 3$.

Another approach, as utilized in Stole and Zweibel (1996), relaxes the assumption that transfers do not change upon disagreement; they assume contracts are non-binding, and thus all other agents can immediately renegotiate upon disagreement (although any disagreeing pair is prevented from recontracting again).[7] Under this assumption $\widetilde{\pi_{U_1}}(D_2; g_3) = -2 + t_{1,1}(g_1)$, where $t_{1,1}(g_1)$ would be what U_1 and D_1 would negotiate under g_1 if U_1 and D_2 never could recontract:

$$t_{1,1}(g_1) \in \arg\max[\pi_{U_1}(t_{1,1}(g_1); g_1) - \widetilde{\pi_{U_1}}(D_1; g_1)]^{b_{U_1}} \times [\pi_{D_1}(t_{1,1}(g_1); g_1) - \widetilde{\pi_{D_1}}(U_1; g_1)]^{b_{D_1}}. \quad (2)$$

Here, since if U_1 and D_1 disagree the network will be g_0 (as now U_1 could no longer recontract with either D_1 or D_2), the outside option for each agent would be 0; hence, if $b_{U_1} = b_{D_1}$, the equilibrium value of $t_{1,1}(g_1) = 6$, which solves (2). Similarly, one can obtain $t_{1,2}(g_2) = 6$ by similar reasoning in g_2. Using these values to construct U_1's outside options when negotiating in g_3, one can then solve (1) under the new renegotiating assumption, and obtain $t_{1,1}(g_3) = t_{1,2}(g_3) = 4$. Note that these payments are strictly higher than the static reasoning used before: e.g., since U_1 obtains a strictly higher payment when renegotiating with D_1 under g_1 than it would have under g_3 (i.e., $t_{1,1}(g_1) > t_{1,1}(g_3)$), it has a better outside option and hence can extract greater rents from D_2 when

[5] Such a setup is not without its own limitations; concerns will be discussed in the next section.

[6] Iozzi and Valletti (2013) highlight differences that may occur when other agents' actions—in addition to contracts—are not allowed to change under disagreement.

[7] de Fontenay and Gans (2013) show that contingent contracts (as used in Inderst and Wey (2003)), where contracts are negotiated at the beginning but payments depend on the realized network structure, generates a similar division of surplus given by the Myerson-Shapley value.

negotiating $t_{1,1}(g_3)$.

This second approach captures the flavor of dynamics in that U anticipates renegotiating with D_1 under g_1 upon disagreement with D_2 under g_3 whereas the first approach does not. Nonetheless, both approaches share the limitation that disagreements do not anticipate the potential for new links to be formed, and neither allows for payments in networks that are not strict subnetworks to influence negotiated rents. This is most stark in networks g_1 and g_2: although U_1 can presumably contract with either D_1 or D_2 exclusively, the presence of a competing downstream firm does not influence the determination of $t_{1,1}(g_1)$ or $t_{1,2}(g_2)$ under either approach; both yield $t_{1,1}(g_1) = t_{1,2}(g_2) = 6$ under equal Nash bargaining parameters. One might presume that U would have a higher outside option than 0 since he could recontract with either D_1 and D_2, and thus U might be able to capture a higher transfers *even with the same Nash bargaining parameter*.

Our paper's model attempts to address this point directly and build upon both static approaches by incorporating the following features: (i) agents can not only break, but also form new links in the future (including with agents with whom they previously disagreed), (ii) adjustments to the network may not be immediate and there may be some delay (including after disagreement), and (iii) agents anticipate these changes and account for them in bargaining when computing expected surpluses from contracting and outside options. We feel that these are particularly important aspects of any model that attempts to capture network formation and bargaining in real world, repeated settings.

We thus specify a game in which each period, links and relationships can be adjusted; however, upon disagreement, contracts among other parties are not able to immediately change. Solving for an equilibrium in this game endogenously determines the transition probabilities across networks, which in turn can be used to construct an internally consistent measure of outside options. To illustrate, we rewrite the negotiated transfer $t_{1,1}(g_1)$ between U_1 and D_1 in (2) as:

$$t_{1,1}(g_1) \in \arg\max \left[(\pi_{U_1}(t_{1,1}(g_1); g_1) + \beta_{U_1} V_{U_1}(g_1)) - \beta_{U_1} V_{U_1}(g_0)\right]^{b_{U_1}} \\ \times \left[(\pi_{D_1}(t_{1,1}(g_1); g_1) + \beta_{D_1} V_{D_1}(g_1)) - \beta_{D_1} V_{D_1}(g_0)\right]^{b_{D_1}} \quad (3)$$

where each agent k anticipates additional payoffs $\beta_k V_k(g_1)$ or $\beta_k V_k(g_0)$ upon reaching an agreement or disagreeing. These represent the expected (discounted) *continuation values* to agent k of being in network g_l at the end of the period, and β_k represents agent k's discount factor. Implicitly, the continuation values account for future renegotiation of payments as well as the fact that, e.g., although the network may be g_0 in the next period upon disagreement, it is unlikely to remain there forever.

Specifying this dynamic game also answers the second question of which networks we expect to arise in equilibrium. The advantage of our dynamic model is that, as opposed to potentially admitting several "stable" or equilibrium networks, it specifies a distribution over networks that are reached in each period. This allows for a constantly evolving network—e.g., firms breaking and/or recontracting over time—within equilibrium play. In the next section, we describe and specify our model; in Section 3.7, we return to this example and illustrate how results are substantively different once we account for dynamics. In particular, we will show how our model will substantively change predictions for $t_{1,1}(g_1)$, for example, and the formulation will allow for U to capture even greater share of the surplus in g_1 as the number of downstream firms increases.

3 Model

We study an infinite horizon, discrete time, dynamic network formation game with transfers between a set of n agents denoted by $N = \{1, 2, \ldots, n\}$. At any point in time, agents may be *linked* to one another thereby defining the current *network structure* or *state*. Let $\mathbf{G} \subset \{0,1\}^{n\times(n-1)}$ represent the set of all *feasible* networks: it may include all possible networks, or solely bi-partite networks that admit only links between two subsets of agents (e.g., as in buyer-seller networks). We focus only on networks which are undirected graphs, so that if i is linked to j in network $g \in \mathbf{G}$, j is linked to i in g as well; this is denoted by $ij \in g$. Let g_i denote the set of links in network g that contain agent i, and g_{-i} denote the network that remains if all links containing agent i are removed.[8] Let $\mathbf{G}_i$ represent the set of feasible links which contain i, and $N_i(g)$ the set of agents i is connected to in state g.

Associated with any state g are *per-period payoffs* $\pi(g, \mathbf{t}_g) \equiv \{\pi_i(g, \mathbf{t}_g)\}_{i\in N}$, where $\mathbf{t}_g \equiv \{t_{ij;g}\}_{ij\in g}$ are *per-period contracts*; $t_{ij;g}$ represents negotiated payments between agents i and j in network structure g.[9] We denote the space of feasible per-period contracts as $\mathbf{T}$, which is determined by the particular application being modeled: e.g., $\mathbf{T} \equiv \mathcal{R}$ if contracts are lump-sum payments or simple linear fees, or $\mathbf{T} \equiv \mathcal{R}^n$ for n-part tariffs. We assume that the payoffs $\pi_i(g, \cdot)$ are continuous in per-period contracts for any g and payments can only be made between linked agents.

For our analysis, we will assume that per-period payoffs are primitives which arise from some exogenously specified subgame, and that payoffs to each agent are uniquely determined for any network structure and set of per-period contracts. E.g., in a buyer-seller setting, period payoffs arise from a price competition game among downstream retailers for consumers given wholesale prices $\mathbf{t}_g$ paid to upstream manufacturers, where links in a given network g determine trading relationships.[10] Finally, we assume every agent i discounts future payoffs at constant rate $\beta \in (0, 1)$.

3.1 Timing and Actions

Consider period τ, when the last period state was $g^{\tau-1}$. Each period of the game can be divided into two stages: a stage in which the set of links open for negotiation is determined, and a stage in which

[8] In this regards, we abuse notation slightly for expositional clarity and interpret $g \in \mathbf{G}$ as both a network as well as a set of links.

[9] Payoffs also may contain any per-period costs of maintaining links to potential trading partners.

[10] In applied work, these payoff functions $\pi(\cdot)$ (even for network structures never observed) can be recovered from separate analysis (c.f. Ackerberg et al. (2007)). For instance, in the example given in Section 2, a structural model of consumer demand for products and a model of retailer pricing can be utilized to obtain period profits for agents given any network structure.

agents bargain over per-period payments for each link open for negotiation. A link then exists for a given period if it is first negotiated and then bargained over without coming to a disagreement. We first provide a brief overview of the timing before detailing specifics.

1. Network Formation: Given last period state $g^{\tau-1}$, the set of links open for negotiation $\tilde{g}$ is determined.

   (a) Every agent i simultaneously announces the set of links $a_i \in A_i \equiv \mathcal{P}(\mathbf{G_i})$ that he wishes to negotiate over, where $\mathcal{P}(\cdot)$ denotes the powerset over all feasible links. By announcing a_i, each agent receives a period payoff shock $\epsilon_{a_i,i}$ (privately observed by each agent before the choice is made), where $\epsilon_i \equiv \{\epsilon_{1,i}, \ldots, \epsilon_{|A_i|,i}\}$ is drawn independently each period from a continuous density $f_i^\epsilon(\epsilon_i)$ with finite first moments, and $|\cdot|$ represents the dimensionality of the set.

   (b) Given the set of announcements $\mathbf{a} \equiv \{a_i\}$, the *negotiation network* $\tilde{g}(\mathbf{a})$ is formed where $ij \in \tilde{g}(\mathbf{a})$ if and only if $ij \in a_i$ and $ij \in a_j$; i.e., a link ij is open for negotiation if and only if both i and j announce it. Each agent i then incurs a cost $c_i(\tilde{g}(\mathbf{a})|g^{\tau-1})$, which represents the cost of negotiating existing or new links, or not negotiating (breaking) links, from the previous network $g^{\tau-1}$.[11,12]

2. Bargaining: Given the negotiation network $\tilde{g}$, agents engage in bilateral Nash bargaining to determine period contracts, the realized network g^τ, and period payoffs.

   (a) First, each pair $ij \in \tilde{g}$ observes a publicly observable pair-specific shock η_{ij}, where $\eta \equiv \{\eta_{ij}\}_{ij\in\tilde{g}}$ is drawn from continuous density f^η, and η_{ij} continuously affects period profits for both i and j if they reach an agreement in the current period. We assume payoff shocks enter additively into per-period payoffs.

   (b) Any pair $ij \in \tilde{g}$ for which there exists no Nash bargaining solution (i.e., no gains from trade *conditional on their expectations of whether other pairs will maintain their agreements*) is *unstable* and immediately dissolves, generating a new network $\tilde{g}'$. If $\tilde{g}'$ also has links that are unstable, those dissolve as well. This repeats until a stable network $g^\tau \subseteq \tilde{g}' \subseteq \tilde{g}$ (i.e., every pair $ij \in g$ has gains from trade) is reached.

   (c) Per-period contracts $\mathbf{t}_{g^\tau}$ are determined via bilateral Nash bargaining, and each agent i obtains *total per-period payoffs* $\overline{\pi}_i(g^\tau, \eta, \mathbf{t}_g^\tau) \equiv \pi_i(g^\tau, \mathbf{t}_g^\tau) + \sum_{ij\in g^\tau} \eta_{ij}$.

One interpretation of the timing is that at the beginning of each period, each agent decides which existing links from the previous period to drop or try to maintain, and which links that were

[11] The model can easily be extended to allow costs to depend on each agent's own actions as well as the realized negotiation network $\tilde{g}$: i.e., $c_i(a_i, \tilde{g}(\mathbf{g})|g^{\tau-1})$.

[12] An alternative modeling assumption would be to allow costs to depend on the previous period's *negotiation* network $\tilde{g}(\boldsymbol{a}^{\tau-1})$ as opposed to the previous realized network $g^{\tau-1}$: this would ensure costs would only be incurred if two agents did not *attempt* to contract in the previous period, as opposed to not having a link between them (which could have been caused by a dissolution of unstable links). The state space would not increase and computation would not be affected.

not in existence to try to form. To maintain an existing link or form a new link, an agent must send a representative to negotiate with the other party; a link is open for negotiation only if both agents on each side send a representative. There may be differential costs of negotiating a link depending on whether or not it was previously in existence; e.g., it may be more costly for two firms who did not previously have a trading relationship to form a new link than for two firms who have interacted in the past. Similarly, if a link was in existence but one party does not wish to continue the relationship by sending a representative, a termination or breakup cost may be incurred.

All agents and representatives then observe all links that are potentially open for negotiation. Each pair of representatives that meet subsequently engages in simultaneous Nash bargains (using either the Rubinstein (1982) or Binmore, Rubinstein and Wolinsky (1986) non-cooperative implementations between each pair), with beliefs over the contracts and outcomes that will be reached in other negotiations.[13] We discuss additional issues related to and provide justification for this particular bargaining framework in Section 3.3.

3.2 Strategies and Value Functions

In the network formation stage of each period, agents announce which set of potential links they would like to open for negotiation. We restrict attention to Markov strategies denoted by $\sigma = \{\sigma_i(g, \epsilon_i)\}$, where $\sigma_i : \mathbf{G} \times \mathbb{R}^{|\mathbf{G}_i|} \to A_i$. Thus each agent conditions only on the last period network state g and their draw of payoff shocks when deciding which links to announce.

Following Hotz and Miller (1993) and Hotz et al. (1994), we define $P_i^\sigma(a_i|g)$ for a given vector of strategies σ to be the *conditional choice probabilities* of action a_i being chosen given last period state g:

$$P_i^\sigma(a_i|g) = Pr(\sigma_i(g, \epsilon_i) = a_i) = \int I\{\sigma_i(g, \epsilon_i) = a_i\} f_i(\epsilon_i) d\epsilon_i, \tag{4}$$

where $I\{\cdot\}$ represents the indicator function. $P_i^\sigma(a_i|g)$ can be interpreted as the probability an opponent that does not observe ϵ_i assigns to i announcing a_i in state g. Given agents' period payoff shocks ϵ are independent, the probability i assigns to the negotiation network being g' given he announces a_i, the last period state is g, and other agents' strategies are (perceived to be) σ can be expressed as:

$$q_i^\sigma(g'|a_i, g) = \sum_{a_{-i} \in \Pi_{j \neq i} A_j} \left(\Pi_{j \neq i} P_j^\sigma(a_{-i}[j]|g)\right) I\{\tilde{g}(a_i, a_{-i}) = g'\}, \tag{5}$$

where $a_{-i}[j]$ is the jth action in the vector of actions of a_{-i} of all other agents excluding i.

Let $v^\sigma(a_i, g)$ represent i's expected current and future profits (net of period payoff shocks ϵ) if he chooses action a_i in state g, and behaves optimally in future periods:

$$v_i^\sigma(a_i, g) = \sum_{g'} q_i^\sigma(g'|a_i, g) \left(c_i(g'|g) + E_\eta \left[\overline{\pi}_i(g'', \eta, \mathbf{t}^\sigma_{g''}) + \beta V_i^\sigma(g'') : g'' = \Gamma(g'; \eta, V^\sigma)\right]\right), \tag{6}$$

[13] Collard-Wexler, Gowrisankaran and Lee (2013) provide a non-cooperative foundation for this bargaining outcome in bilateral oligopoly when firms are allowed to make and receive multiple offers in each period.

These *choice-specific value functions* takes expectations over the per-period contracting shocks η and the resultant network g'' that arises after bargaining occurs: i.e., $g'' = \Gamma(g'; \cdot)$ represents the *stable* network which arises from g' as *unstable* links dissolve. Γ and our notion of stable networks and unstable links will be defined later. Finally, $V^\sigma \equiv \{V_i^\sigma(\cdot)\}$ are the set of value functions for every agent i, where:

$$V_i^\sigma(g) = \int \left[\max_{a_i \in A_i} \left(\epsilon_{a_i,i} + v_i^\sigma(a_i, g) \right) \right] f_i(\epsilon_i) d\epsilon_i \tag{7}$$

and represents i's expected current and future profits at the beginning of any period (before the realization of payoff shocks) from behaving optimally, given the last period network was g and all other agents act according to strategies σ_{-i}. Note that for a fixed set of payoffs $\{\pi(\cdot)\}$, the right hand side of (7) is a contraction mapping (c.f. Aguirregabiria and Mira (2002)), and thus there is a unique V_i^σ which solves (7) for any given σ.

3.3 Per-period Contracts and Nash Bargaining

To close the model, we define the process by which per-period contracts are determined. In turn, this allows us to define the function $\Gamma(g; \cdot)$, which finds the network that arises from any negotiation network g as *unstable* links are broken.

We assume the set of all period contracts $\mathbf{t}^\sigma(\eta) \equiv \{\mathbf{t}_g^\sigma(\eta)\}$ are determined endogenously via (asymmetric) Nash bargaining. Assume g is *stable* (which we will define shortly); in this case, every contract $t_{ij;g}(\eta) \in \mathbf{t}_g^\sigma(\eta)$ is assumed to satisfy the following:

$$\begin{aligned} t_{ij;g}(\eta) \in \arg\max_{\tilde{t}} & \Bigg[\underbrace{[\overline{\pi}_i(g, \eta, \{\tilde{t}, \mathbf{t}^\sigma_{-ij;g}\}) + V_i^\sigma(g)]}_{S^\sigma_{i,j}(g;\eta,\mathbf{t}^\sigma_g)} - \underbrace{[\overline{\pi}_i(g', \eta, \mathbf{t}^\sigma_{-ij;g}) + V_i^\sigma(g')]}_{S^\sigma_{i,j}(g-ij;\eta,\mathbf{t}^\sigma_{-ij;g})} \Bigg]^{b_{ij}} \\ & \times \Bigg[\underbrace{[\overline{\pi}_j(g, \eta, \{\tilde{t}, \mathbf{t}^\sigma_{-ij;g}\}) + V_j^\sigma(g)]}_{S^\sigma_{j,i}(g;\eta,\mathbf{t}^\sigma_g)} - \underbrace{[\overline{\pi}_j(g', \eta, \mathbf{t}^\sigma_{-ij;g}) + V_j^\sigma(g')]}_{S^\sigma_{j,i}(g-ij;\eta,\mathbf{t}^\sigma_{-ij;g})} \Bigg]^{b_{ji}} \end{aligned} \tag{8}$$

where $g' = g - ij$, $\mathbf{t}^\sigma_{-ij;g} \equiv \{\mathbf{t}^\sigma_g \setminus t^\sigma_{ij;g}\}$, and b_{ij}, b_{ji} represent agent i and j's Nash relative bargaining parameters, which are primitives of the analysis. Thus, each $t_{ij;g}(\eta)$ maximizes the (weighted) Nash product of i and j's gains from trade (represented by $\Delta S^\sigma_{i,j}(g;\eta,\mathbf{t}^\sigma_g) \equiv S^\sigma_{i,j}(g;\eta,\mathbf{t}^\sigma_g) - S^\sigma_{i,j}(g - ij;\eta,\mathbf{t}^\sigma_{ij;g})$ and $\Delta S^\sigma_{j,i}(g;\eta,\mathbf{t}^\sigma_g) \equiv S^\sigma_{j,i}(g;\eta,\mathbf{t}^\sigma_g) - S^\sigma_{j,i}(g - ij;\eta,\mathbf{t}^\sigma_{-ij;g})$), given the contracts of all other linked pairs of agents and the strategies of agents employed in the larger network formation game.

The bargaining outcome in (8) is a version of the static bilateral bargaining equilibria between an upstream supplier and downstream firms used in Cremer and Riordan (1987) and Horn and Wolinsky (1988), which has been adapted in applied work to model negotiations between upstream content providers and downstream multichannel video distributors (Crawford and Yurukoglu, 2012), between medical device suppliers and hospitals (Grennan, 2013), and between hospitals and insurers (Gowrisankaran, Nevo and Town, 2013; Ho and Lee, 2013). We extend the existing literature by building on this framework, known for its tractability and ability to capture firm heterogeneity, and admitting endogenously determined dynamic outside options. In addition, though originally motivated a cooperative solution concept, the asymmetric Nash bargaining outcome among multiple firms that we leverage can be motivated via different noncooperative extensive forms.[14,15] We thus assume that agents engage in one of these extensive forms once the negotiation network is determined in each period, thereby generating contracts given by (8).

Note the gains from trade for each agent ($\Delta S^{\sigma}_{i,j}(g;\eta,\mathbf{t}^{\sigma})$ and $\Delta S^{\sigma}_{j,i}(g;\eta,\mathbf{t}^{\sigma})$) consists of what each agent would obtain if i and j link minus what each agent expects to obtain upon disagreement. Importantly, in a significant departure from previous literature, we do not assume the disagreement point if i and j fail to contract to be fixed nor a function of one particular network structure (which would be implied if links and contracts were never renegotiated), nor do we necessarily assume all other pairs immediately renegotiate (as in Stole and Zweibel (1996)). Rather, disagreement points are defined to be what i or j expect to obtain if they do not contract in the current period plus any impact on future payoffs given by continuation values $V^{\sigma}_i(g')$ and $V^{\sigma}_j(g')$. Our notion of a disagreement point is thus internally consistent with the larger dynamic dynamic game: when two agents fail to come to an agreement, period profits are derived from a network in which they are not linked, but agents anticipate subsequent changes to the network. Indeed, i and j may even anticipate contracting again in the future.

Nonetheless, there is some flexibility in specifying how agents perceive current period profits will be upon disagreement. As written in (8), if agents i and j disagree, all other period contracts that would have negotiated under network g (i.e., $\mathbf{t}^{\sigma}_{-ij;g}$) are binding, even if the new network is $g' = g - ij$. This is similar to the assumption used in Horn and Wolinsky (1988) on contracts (though we assume agents will respond to the new network structure $g - ij$ for any other period actions which affect the determination of $\pi_i(\cdot)$). We believe this is reasonable for the applied settings we have in mind between firms: since bargains occur simultaneously and disagreements in stable networks are off-equilibrium events, we do not assume that other firms can immediately renegotiate their contracts. However, if we assume agreements are non-binding or can be made contingent on the network as in Stole and Zweibel (1996), then our specification can be adjusted

[14] E.g., Collard-Wexler, Gowrisankaran and Lee (2013) show that a non-cooperative game with alternating offers between many upstream and downstream firms admits the solution to (8) as an equilibrium outcome as the time period between offers goes to 0. They extend the Rubinstein (1982) bargaining game to a bilateral oligopoly setting: downstream firms simultaneously make take-it-or-leave-it offers in "odd" periods to upstream firms in g with whom they do not have an agreement; upstream firms simultaneously accept or reject offers that are made. In the subsequent "even" periods, upstream firms make take-it-or-leave-it offers to downstream firms with whom they do not have an agreement, and downstream firms choose whether or not to accept offers. The game continues until all agreements have been reached. The authors assume agents' beliefs following off-equilibrium offers are *passive* (c.f. McAfee and Schwartz (1994)) in that beliefs over the status or progress of other negotiations are not updated, and each agent i discounts payoffs between periods by $\delta_i \equiv \exp(-r_i\Lambda)$, where Λ is the time between periods and r_i is i's discount rate. Alternatively, $(1-\delta_i)$ can represent the exogenous probability of breakdown after each offer is made (Binmore, Rubinstein and Wolinsky, 1986). As Λ approaches 0, there exists an equilibrium with payoffs approaching (8) with $b_{ij} = r_i/(\lambda_i + \lambda_j)$.

[15] Another motivation for this setup in the literature is assuming each firm sends different "representatives" to bargain simultaneously with each of its linked partners; bargains happen simultaneously according to a non-cooperative alternating offers game as in Rubinstein (1982), and agents from the same firm do not coordinate with one another across separate bargains. See also Crawford and Yurukoglu (2012).

so that the period profits under disagreement are a function of contracts that would have been negotiated if i and j did not contract (i.e., $\mathbf{t}^{\sigma}_{\Gamma(g';\cdot)}$); in addition, given the outside option would then depend on per-period contracts negotiated in other network states, agent gains from trade $\Delta S^{\sigma}_{i,j}$ would be a function of all transfers $\mathbf{t}^{\sigma}$ and not just those negotiated in one network state.

Unstable Networks There may be instances in which for certain pairs of linked agents $\{ij \in g\}$, draws of η, and perceived continuation values V, there is no value of $\mathbf{t}_g$ in which both $\Delta S^{\sigma}_{i,j}(g;\eta,\mathbf{t}^{\sigma}_g)$ and $\Delta S^{\sigma}_{j,i}(g;\eta,\mathbf{t}^{\sigma}_g)$ are positive; i.e., there are no potential gains from trade. In this case, the Nash bargaining solution is undefined. Without restrictions on π, this may occur in equilibrium due to the presence of general externalities: e.g., an agent i by forming a new link or dissolving an existing link may cause a link between agents j and k, which might have previously exhibited gains from trade, to not be profitable to maintain.[16]

Any network g for which there exists no set of period-contracts $\mathbf{t}^{\sigma}_g$ such that there are gains from trade between all pairs of agents $ij \in g$ is considered *unstable*. Conversely, if $\mathbf{t}^{\sigma}_g$ exists which satisfies (8) such that there are gains from trade between all connected agents, g is considered *stable*.

Given there is some flexibility in the order in which links dissolve, we adopt the following rule: if a network is unstable, then any link $ij \in g$ in which there exists some $\mathbf{t}_{-ij;g}$ such that $\Delta S_{i,j} < 0$ or $\Delta S_{j,i} < 0$ for all $t_{ij;g}$ is an *unstable link* and is immediately and simultaneously broken. This will yield a new network, which will either be stable or unstable; if unstable, the process by which unstable links dissolve continues. Eventually, a stable network is reached, which is given by the function $\Gamma(g;\eta,V)$ and is defined recursively as:

$$\Gamma(g;\eta,V) = \begin{cases} g & \text{if } \exists \mathbf{t}_g \text{ s.t. } \; \forall\, ij \in g,\; \Delta S_{ij}(g;\eta,\mathbf{t}_g) \geq 0, \\ \Gamma(g';\eta,V) & \text{otherwise, where } g' = g \setminus \{ij \in g : \exists \mathbf{t}_{-ij;g} \text{ s.t.} \\ & \qquad \forall t_{ij;g},\; \Delta S_{ij}(g;\eta,\{t_{ij;g},\mathbf{t}_{-ij;g}\}) < 0\} \end{cases} \tag{9}$$

Note that $\Gamma(g;\cdot)$ may also be the empty network.

3.4 Markov Perfect Network Formation Game

We can parametrize our model by the tuple $(N,\mathbf{G},\pi,\mathbf{b},\beta,\mathbf{f},\mathbf{c})$, where $\pi = \{\pi_i\}$, $\mathbf{b} = \{b_{ij}\}$, $\beta = \{\beta_i\}$, $\mathbf{f} = \{f^{\eta},\mathbf{f}^{\epsilon} \equiv \{f^{\epsilon}_i\}\}$, and $\mathbf{c} = \{c_i(\cdot)\}$.

[16] This is also why agents would anticipate moving to $g' \equiv \Gamma(g-ij;\cdot)$ as opposed to simply $g-ij$ under disagreement if contracts could be renegotiated within a period.

3.5 Equilibrium

A (pure-strategy) Markov-Perfect equilibrium (MPE) of this game is a set of strategies σ^* such that for any proposer i, network g, and payoff shocks ϵ_i:

$$\sigma_i^*(g, \epsilon_i) = \arg\max_{a_i \in A_i} \left[\epsilon_{a_i,i} + v_i^{\sigma^*}(a_i, g)\right]. \tag{10}$$

In addition, given V^{σ^*} for any η, period contracts $\mathbf{t}_g^{\sigma^*}$ satisfy (8) for all stable g and $\Gamma(g;\cdot)$ satisfies (9) for all $g \in \mathbf{G}$.

Existence Following Milgrom and Weber (1985) and Aguirregabiria and Mira (2007), we find that an MPE of this game can be re-expressed in probability space. Let σ^* be an MPE, and P^{σ^*} be the associated conditional choice probabilities. Note that P^{σ^*} is a fixed point of the the following *best response probability function* function $\Lambda(P) = \{\Lambda_i(a_i|g; P_{-i})\}$, where

$$\Lambda_i(a_i|g; P_{-i}) = \int I\left\{a_i = \arg\max_{a \in A_i}\left(\epsilon_{a,i} + v_i^P(a, g)\right)\right\} f_i^{\epsilon}(\epsilon_i) d\epsilon_i \tag{11}$$

where v^P are choice specific value functions derived from v^σ in (6) defined explicitly in terms of conditional choice probabilities P.

To prove there exists at least one fixed point of Λ and hence at least one MPE, it is sufficient to prove that Λ is continuous in the compact space P; the existence of at least one fixed point will then follow from Brouwer's theorem. From (6), we see that the continuity of v_i^P is guaranteed if the expression $E_\eta\left[\pi_i(g, \eta, \mathbf{t}_g^P) + \beta V_i^P(g) : g = \Gamma(g'; \eta, \boldsymbol{V}^P)\right]$ is continuous in P for all g, g', where V_i^P is the value function given by (7) expressed in conditional choice probabilities. This in turn follows given assumptions on the per-period payoff functions, the continuity of the density function f^η, and the following additional assumption:

Assumption 3.1. For all g and continuation values $\{V_i^P(\cdot)\}$, there exists $\bar{\eta}$ in the support of f^η s.t. $\forall \eta > \bar{\eta}$, (i) there is a unique set of per period contracts $\mathbf{t}_g^P(\eta)$ that solves (8); (ii) each contract in $\mathbf{t}_g^P(\eta)$ is continuous in P and η; and (iii) $\pi_i(g, \mathbf{t}_g)$ is continuous in $\mathbf{t}_g$.

A.3.1 rules out the possibility that a small change in conditional choice probabilities P_i for some agent results in a discontinuous change in $E_\eta\left[\pi_i(g, \eta, \mathbf{t}_g^P) + \beta V_i^P(g) : g = \Gamma(g'; \eta, V^\sigma)\right]$ due to either the set of negotiated transfers changing discontinuously, the network g becoming unstable for any realization of η, or $\pi_i(\cdot)$ changing discontinuously. A sufficient condition that guarantees there is always a realization of η which ensures g is stable is that f^η has full support. Finally, since f_i^ϵ is assumed to be continuous, Λ is continuous in P, and the existence claim follows.

A.3.1 is stronger than necessary to guarantee existence; additionally, whether or not it holds depends on the specification of π, and will be application dependent. In our applied example discussed in the next section, we were always able to compute and find an MPE of the game even with $\eta = 0$.

3.6 Discussion

Multiplicity and Pairwise Stability Within a period, there is the potential for multiplicity in both the network formation stage and the bargaining stage. We do not address the latter since it will depend on the period profit functions π; we are inherently relying on an assumption such as 3.1 or an ability to consistently choose a unique set of transfers if there are many that satisfy (8).

At the network formation stage, one might imagine that our reliance on the Myerson (1991) network formation game—through its requirement that two agents simultaneously announce a link in order for it to form—would admit a similar kind of multiplicity issue as discussed therein: e.g., the empty network could always be a Nash equilibrium since if everyone announced the empty network, no agent could profitably deviate. However, our use of period shocks ϵ_i partially addresses this issue since it rules out the possibility that *pairwise profitable* links (given the actions of others) would not be negotiated due to miscoordination. To see why this is the case, consider a set of strategies in which g is the realized negotiation and bargained network, but $g \cup \{ij\}$ would have been preferred by both i and j.[17] If link ij is not negotiated since agent j never proposes it, then (net the payoff shock) i should be indifferent between proposing ij and not; thus, there will be strictly non-zero probability that i proposes link j to negotiate. If this is the case, then j should place non-zero probability on negotiating with i, and the original set of strategies could not be an equilibrium.[18] Nonetheless, insofar that there are multiple pairwise stable networks which can be proposed (given profits and continuation values), there may still an issue of which negotiation network(s) will be reached; whether this is the case will depend on the underlying primitives of the model.

Finally, we believe that this issue of multiplicity is mitigated since we also enforce optimality at all states, even those that might not be contained in the recurrent class (see Fershtman and Pakes (2012) for a weaker equilibrium condition in dynamic games); however, we acknowledge that there still may exist multiple MPE, each with potentially different recurrent classes of networks, and/or different equilibrium strategies.

Relationship to Static Bargaining Models

The model as specified only introduces dynamics through frictions in forming and breaking links. If $c_i(g'|g) = 0 \; \forall \; g', g$ (or if $\beta = 0$), then the continuation value functions do not enter the determination of period-contracts (given by (8)) in any meaningful way. Consequently, if agents propose a given network g which is stable (i.e., there exists $\mathbf{t}_g$ such that a solution to (8) exists for each link $ij \in g$), the contracts negotiated will coincide with those from a static Nash-in-Nash bargains solution, similar to Horn and Wolinsky (1988) if contracts are not allowed to be renegotiated upon

[17] I.e., period profits and continuation values in $g \cup \{ij\}$ are higher than under g given negotiated transfers for both i and j.

[18] This reasoning is similar to Jackson and Watts (2002) who study a game in which links are randomly chosen and can be formed if both agents prefer to do so, and broken if at least one agent chooses to do so. With some small probability, however, the opposite occurs. They show that if a pairwise stable network exists, any *stochastically stable* network—i.e., a network with steady-state probability bounded away from 0 as the probability of error goes to 0—must be pairwise stable. Our use of errors ϵ and η perform similar functions in avoiding non-pairwise stable states.

disagreement, or Stole and Zweibel (1996) if they are allowed to be immediately renegotiated (but not re-formed) in a given period. Even without dynamics, however, this model provides guidance as to which network(s) are reached in equilibrium as agents are allowed to propose which links can form in each period.[19]

3.7 Example Revisited

At this point, we return to the example discussed in Section 2 and depicted in Figure 1. Recall under both static bargaining procedures we discussed, $t_{1,1}(g_1) = t_{1,2}(g_2) = 6$ under equal bargaining power. When contracts among other agents are binding under disagreement, $t_{1,1}(g_3) = t_{1,2}(g_3) = 3$. As noted previously, if either $\beta = 0$ or $c_i(g'|g) = 0 \; \forall \; g', g$, the predicted transfers in our dynamic model would coincide with these as well. Under these assumptions, our game would essentially become a repeated version of the static bargaining game, sharing these equilibrium divisions of surplus. Equilibrium strategies for U would be to mix evenly between proposing a link to either D_1 or D_2 (but never both); and D_1 and D_2 would always propose linking with U. Thus, the model yields the prediction that as long as the variance of the ϵ shocks are small, only networks g_1 and g_2 should occur (with equal probability); as the variance increases, the probability of g_3 being reached increases as well.

However, once $\beta > 0$ and there are costs to changing the network structure, results are different. First, let us consider the case where $\beta = .9$, agents each bear a period cost of 1 to form a link that was not present in the previous period, and ϵ is distributed extreme value type I with variance $\pi^2/8$. In equilibrium, $t_{1,1}(g_1) = t_{1,2}(g_2) \approx 7.6$, and $t_{1,1}(g_3) = t_{1,2}(g_3) \approx 4.4$. Note that not only are the transfers significantly higher than the static Nash results, but also those that would be received under the Stole and Zweibel (1996) assumptions as well. To understand why U can obtain higher transfers even with equal bargaining power, consider the bargain in g_1. Since there are costs of forming a new link, if U and D_1 come to an agreement the next period network will likely be g_1 as well. However, upon disagreement, U will just as likely choose to link with D_2 in the next period as it would relink with D_1, which would result in D_1 being potentially without a trading partner for several periods. This provides U with additional leverage and improves its outside option.

The steady-state distribution across states $[g_0, g_1, g_2, g_3] \approx [.00, .43, .43, .14]$, and the induced equilibrium transition probabilities give an 80% chance of staying in either g_1 or g_2 conditional on having been there in the previous period. Note that both D_1 and D_2 are receiving negative period profits under g_3 but g_3 is still stable; this is because both downstream firms realize there is sufficient chance that the network will transition from g_3 to either g_1 or g_2 (which would be extremely profitable for D_1 or D_2), thereby rationalizing the short-term losses incurred.

For $\beta = .9$, as the variance of ϵ goes to 0 and the cost of forming a new link becomes arbitrarily small (but non-zero), the equilibrium network will tend to stay in either g_1 or g_2 in perpetuity. However, U will be able to negotiate a higher transfer in either case, such that $t_{1,1}(g_1)$ and $t_{1,2}(g_2)$

[19] The model also can easily be extended to allow for additional state variables (e.g., investment, capacity) which would introduce additional sources of dynamics.

approaches $(24 - 14\beta)/(4 - 3\beta) = 114/13 \approx 8.8$.[20]

Bertrand Convergence Using the same logic, it can be shown that as $\beta \to 1$, the negotiated transfers approach full surplus extraction by U of 10.

Even for fixed β, as the number of potential competitors grows, the negotiated transfer also increases. Consider the same example with one upstream firm, but now allow for multiple downstream firms $D_1, \ldots, D_n$. Let payoffs remain the same: if U is linked exclusively with any D_i, D_i receives 10, U receives -2; if U links with $k \geq 2$ agents, payoffs are $-2k$ for U and $8/k$ for all linked D_i. All unlinked agents receive 0. Using the same logic, it can be shown *for any set of Nash bargaining parameters* and an arbitrarily small but non-zero cost of forming a new link, as $\sigma_\epsilon \to 0$, $n \to \infty$, U will be able to extract $(12-2\beta)/(2-\beta)$ from any D_i it exclusively contracts with in our model and will only contract with one agent. The reason it cannot extract full surplus for $\beta < 1$ is that D_i can still destroy one period's worth of stage profits upon disagreement, and thus retains some leverage.

On the other hand, for equal Nash bargaining parameters, any static model will still yield the result that U can at most extract 6 from an exclusive relationship no matter how many potential alternative contracting partners it has.

3.8 Extensions to the Model

3.8.1 Endogenous Mergers

The model can be extended to allow for the possibility that certain agents can "merge" with other agents, where merging implies that the two agents are permanently linked (ignoring the possibility for dissolving a merger), and that they jointly propose new links and bargain over all future contracts. This adaptation contributes to the literature on dynamic endogenous mergers (Gowrisankaran (1999), Gowrisankaran and Holmes (2004)) within a framework of endogenous link formation and contracting.

To simplify exposition, we focus on settings in which the set of agents who can merge is exogenously given, and that if these agents link, they can only merge. For example, in the case of buyer-seller networks, although buyers can link and engage in trade with sellers, buyers are only allowed to merge with other buyers (and sellers with other sellers). Exogenously we assume that for any pair ij that can only merge, one agent (say i) will always be deemed the "acquirer" and the other agent the "target." If ij merge, j no longer retains any strategic actions: all future profits

[20]To see this, note equilibrium strategies will be that D_i will always propose a link under $\{g_0, g_i, g_3\}$ and will propose a link with equal probability otherwise; U will always propose to stay in g_1 or g_2 if that was the previous network, and propose g_1, g_2 with equal probability otherwise. In this case, note negotiated transfers must solve:

$$t_{1,1}(g_1) = \arg\max_t [(10 - t + \beta C_{D_1}) - (0 + .5\beta(C_D + c_D))]^{b_D} [(-2 + t + \beta C_U) - (0 + \beta(C_U + c_U))]^{b_U}$$

where $C_{D_1} = (10 - t)/(1 - \beta)$ and $C_U = (-2 + t)/(1 - \beta)$ reflect the continuation values for D_1 remaining in g_1 and U remaining in either g_1 or g_2 in perpetuity, and c_i is the cost of forming a new link for agent i. Note that under disagreement, under equilibrium strategies U re-contracts with either D_1 or D_2 with equal probability. Letting $c_i \to 0$ delivers the result.

that would otherwise accrue to j are captured by i, and any links that are formed with i are also formed with j: i.e., for any network g where $ij \in g$, $N_i(g) = N_j(g)$. Since the merger payment is not assumed to be positive or negative, which firm is the acquirer and which is the target is only for expositional purposes.

We maintain the timing and structure of the main model, but when two agents i and j that can merge propose links with one another, they do not negotiated over per-period contracts during the bargaining stage, but rather over a lump-sum acquisition payment $T_{ij;g}$, which satisfies the following (given negotiated transfers $\mathbf{t}^{\sigma}_{-ij;g}$ for all other agent pairs):

$$T_{ij;g} \in \arg\max_{\tilde{T}} \left[\left[\pi_i(g, \eta, \mathbf{t}^{\sigma}_{-ij;g}\}) - \tilde{T} + V_i^{\sigma}(g)\right] - \left[\pi_i(g', \eta, \mathbf{t}^{\sigma}_{-ij;g}) + V_i^{\sigma}(g')\right] \right]^{b_{ij}} \\ \times \left[[\tilde{T}] - \left[\pi_j(g', \eta, \mathbf{t}^{\sigma}_{-ij;g}) + V_j^{\sigma}(g')\right] \right]^{b_{ji}} \tag{12}$$

where again, $g' = g - ij$, and profits and continuation values for i in all networks where i and j are merged (which includes network g) is the sum of profits for the merged entity (both i and j). Note that if there exists no $T_{ij;g}$ such that both i and j have gains from trade (i.e., both terms in the expression are positive), then the merger is not feasible, and does not form.

We note that more complicated models—e.g., in which two firms can either be merged or contractually linked (where the distinction is whether or not they have to repeatedly contract, and whether or not they internalize joint profits in their decisions), or if mergers can be dissolved—will likely require expanding the state space. The advantage of this parsimonious specification in which two firms can only be either contractually linked or merged is that the state space still remains only the space of all potential feasible networks $\mathbf{G}$.

3.8.2 Exclusive Contracts

In certain applications, an agent i may wish to propose an *exclusive contract* to agent j in which i will only link with j if j does not link with any one else. We can extend our model to incorporate this possibility by expanding each agent's choice set to include the possibility of proposing exclusive links: e.g., each agent i can announce a set of non exclusive links a_i and exclusive links a_i^e to form. In each period, given the set of announcements $\mathbf{a} \equiv \{a_i, a_i^e\}$, the negotiation network $\tilde{g}(\mathbf{a})$ is formed where $ij \in \tilde{g}(\mathbf{a})$ if $ij \in a_i$ and $ij \in a_j$, or $ij \in a_i^e$ and $ij = \{a_j\}$. That is, if i proposes link ij exclusively, link ij is open for negotiation if only if j announces only ij. The model proceeds as before.

Notice that our model does not support imposing specific penalties to agents for breaking an exclusive contract (either by no longer announcing it, or by contracting with other agents in future periods). The reason is that in any given state g, the model cannot distinguish between an agent who has voluntarily chosen to contract with only one other agent, and an agent who has accepted an exclusive contract in a previous period.

4 Computation & Estimation

4.1 Computation of Equilibrium

Although finite per-period payoff shocks η are used to guarantee existence, we detail computation of the equilibrium assuming the support of f^η is arbitrarily small so that η essentially does not affect computation. This section follows closely the discussion in Aguirregabiria and Mira (2007).

Let σ^* be an equilibrium, P^* be the associated conditional choice probabilities, and $\{V_i^{P^*}\}$ the equilibrium value functions for all agents. Note that in equilibrium, we can rewrite any agent i's value function as:

$$V_i^{P^*}(g) = \sum_{a_i \in A_i} P_i^*(a_i|g) \left[\tilde{\pi}_i^{P^*}(a_i|g) + e_i^{P^*}(a_i, g)\right] + \beta \sum_{g'} V_i^{P^*}(g') Q^{P^*}(g'|g), \tag{13}$$

where $\tilde{\pi}_i^{P^*}(a_i|g)$ is i's expected period profits (including costs of negotiating new and existing links) from action a_i, $e_i^{P^*}(a_i|g)$ is the expected choice-specific payoff shock to agent i choosing a_i, and $Q^{P^*}(g'|g)$ are the induced transition probabilities between states g and g'. Formally:

$$\tilde{\pi}_i^{P^*}(a_i|g) = \sum_{g'} q^{P^*}(g'|a_i, g) \left[c_i(g'|g) + \pi_i(\Gamma(g'; V^{P^*}), \mathbf{t}_g^{P^*})\right] , \tag{14}$$

$$e_i^{P^*}(a_i, g) = E\left[\epsilon_{a_i,i}|\sigma_i^*(g, \epsilon_i) = a_i\right] = \int \epsilon_i I\left\{\sigma_i^*(g, \epsilon_i) = a_i\right\} f_i(\epsilon_i) d\epsilon_i , \tag{15}$$

$$Q^{P^*}(g'|g) = \sum_{\mathbf{a} \in \Pi_k A_k} \Pi_{j=1}^N P_j^*(\mathbf{a}[j]|g) I\left\{\Gamma(\tilde{g}(\mathbf{a}); V^{P^*}) = g'\right\} . \tag{16}$$

Even though σ^* is referenced in (15), $e_i^{P^*}(a_i, g)$ is a function only of P^* and choice specific value functions $v_i^{P^*}(a_i, g)$ (c.f. Hotz and Miller (1993)).[21] Consequently for any equilibrium, the computation of value functions in (13) can be obtained via matrix algebra and rewritten in matrix notation as:

$$\mathbf{V}_i^{P^*} = \left(\mathbf{I} - \beta_i \mathbf{Q}^{P^*}\right)^{-1} \left(\sum_{a_i \in A_i} \mathbf{P}_i^*(a_i) * \left[\tilde{\boldsymbol{\pi}}_i^{P^*}(a_i) + \mathbf{e}_i^{P^*}(a_i)\right]\right) \tag{17}$$

where $\mathbf{I}$ is the identity matrix, $\mathbf{Q}^{P^*}$ and $\mathbf{P}_j^*$ are matrices of transition probabilities $Q^{P^*}(g'|g)$ and $P_i^*(a_i|g)$, and $\mathbf{V}_i^{P^*}$, $\pi_i^{P^*}$, and $\mathbf{e}_i^{P^*}(a_i)$ are vectors across all network states g; and $*$ denotes the Hadamard, or element-by-element product. Importantly, values at any state g which are unstable are replaced with those at state $\Gamma^{P^*}(g; V^{P^*})$.

Define $\boldsymbol{\Upsilon}_i(P) \equiv \{\Upsilon_i(g, P) : g \in \mathbf{G}\}$ to be the solution to (17) for an arbitrary set of probabilities P; i.e., $\boldsymbol{\Upsilon}_i(P)$ is the expected current and future profits of i given all firms (including i) behave according to conditional choice probabilities P. We can follow the same procedure in Aguirregabiria

[21] Certain distributions of ϵ yield closed form expressions for $e_{i,i}^{P^*}(\cdot)$; e.g., if distributed type I extreme value with variance $(\sigma\pi)^2/6$, $e_i^P(a_i, g) = \gamma - \sigma \ln(P_i(a_i|g))$ where γ is Euler's constant.

and Mira (2007) to show that any fixed point of the mapping $\Psi(P) \equiv \{\Psi_i(g'|g;P)\}$, where:

$$\Psi_i(a_i|g;P) = \int I\left\{a_i = \arg\max_{a\in A_i}\left(\epsilon_{a,i} + \tilde{\pi}_i^P(a,g) + \beta_i \sum_{g'} \Upsilon_i(\Gamma(g';\mathbf{\Upsilon}(P)),P)q_i^P(g'|a,g)\right)\right\} f_i(\epsilon_i)d\epsilon_i$$

will also be a fixed point of Λ, and hence will be an MPE.

This suggests a natural computational algorithm to compute an equilibrium: start with initial values for strategies σ^0, transition probabilities P^0, per-period contracts $\mathbf{t}^{P^0}$, value functions $\mathbf{V}^{P^0}$, and at each iteration τ:

1. Obtain updated implied transition probabilities $P^\tau(g')$ from strategies $\sigma^{\tau-1}$, given by (4);
2. For each agent i, update $V_i^{P^\tau}$ either by value function iteration, or setting $V_i^{P^\tau} = \Upsilon_i(P^\tau)$ using (17), where modified policy iteration can be utilized to approximate the matrix inversion (c.f. Judd (1998));
3. Update per-period contracts $\mathbf{t}^{P^\tau}$ and $\Gamma(g;V^{P^\tau})$ using (8);
4. Update agents' optimal strategies given $V_i^{P^\tau}$ to obtain σ^τ.

The algorithm is repeated until convergence. In practice, we stop when $|\mathbf{V}^{P^{\tau+1}} - \mathbf{V}^{P^\tau}| < \rho$, where ρ is some prespecified tolerance, and $|\cdot|$ denotes the sup-norm.[22]

On the "curse of dimensionality" One issue with using the entire network structure as the state space is that the size of the space grows exponentially: the dimensionality is $2^{n\times(n-1)}$ in general network formation games; in bi-partite network formation games, the dimensionality is $2^{B\times S}$ where one side has B agents and the other S. For small n, computation is not problematic, as the entire state space can be traversed rapidly. For larger games, reinforcement learning algorithms (c.f. Pakes and McGuire (2001)) and other approximation techniques may be applicable.

4.2 Estimation of Nash Bargaining Parameters

In this section, we describe how unobserved parameters for agents—in particular, Nash Bargaining parameters—can be estimated as a function of observed actions. In turn, this implies that unobserved transfers $\mathbf{t}$ can also be recovered. Intuitively, if there are gains from trade between two agents who form a link (given the actions of others), a static model would predict that the link should form regardless of which agent obtains a larger share. However, in a dynamic model, different values of Nash bargaining parameters will change each agent's respective outside options through their continuation values, and hence only certain parameter values will be consistent with a link forming in equilibrium.

[22] Note steps 2 and 3 could be iterated to find the true value of $\Upsilon(P^n)$ for any given P^n; in practice, we found this unnecessary.

Let the data consist of $m = 1 \dots M$ markets, each with primitives $(N^m, \mathbf{G}^m, \pi^m, \beta^m, \mathbf{f}^m, \mathbf{c}^m)$ which are either observed, assumed, or can be separately estimated. We assume Nash bargaining parameters $\mathbf{b} \equiv \{b_{ij}\}$ can be parameterized as a function of observable market characteristics $z^{m,t}$ and parameters to be estimated θ.

In this section, we detail two approaches which vary in data requirements and computational complexity. The first requires recomputing the MPE for each evaluation of θ, and will rely on either uniqueness of equilibrium or an ability to compute all potential equilibria; however, it only requires observing a sequence of realized network structures over time for each market. The second approach does not require computing equilibria nor is constrained by potential equilibrium multiplicity; however, it requires observing actions (i.e., links proposed) by each agent.

In section 5.3, we show how these estimators can be used for inference in simulated markets.

4.2.1 Estimation with Full Equilibrium Computation

We assume the econometrician observes a sequence network structures $\{g^{m,t}\}$ for markets $m = 1 \dots M$ and periods $t = 1 \dots T$. There are two potential cases: (i) the econometrician observes the complete sequence of networks in a market without gaps; or (ii) the econometrician sees a potentially incomplete sequence of networks.

For case (i), we can define a *pseudo-likelihood* based on transition probabilities between networks $g^{m,t}$ and $g^{m,t-1}$:

$$R_M^1(\theta, P) = \sum_{m=1}^{M} \sum_{t=2}^{T} \ln Q^P \left(g^{m,t} | g^{m,t-1}; \theta, \mathbf{b}(\theta, z^{m,t}) \right), \tag{18}$$

On the other hand, for case (ii) where there may be gaps between $g^{m,t}$ and $g^{m,t-1}$, we can instead rely on the following *pseudo-likelihood*:

$$R_M^2(\theta, P) = \sum_{m=1}^{M} \sum_{t=2}^{T} \ln \tilde{Q}^P_{g^{m,1}} \left(g^{m,t}; \theta, \mathbf{b}(\theta, z^{m,t}) \right), \tag{19}$$

where $\tilde{Q}(\cdot)$ is the steady-state equilibrium ergodic distribution over *all* networks given P and initial network $g^{m,1}$.

Since there is the possibility of multiple equilibria for a given θ, we define the MLE as:

$$\hat{\theta}_{MLE} = \arg\max_{\theta} \left[\sup_{P \in (0,1)^{N \times |\mathbf{G}|}} R_M^i(\theta, P) \qquad \text{subject to} \qquad P = \Psi(\theta, P) \right], \tag{20}$$

where $R_M^i(\theta, P)$ is either given by (18) or (19) depending on the nature of the data. If all equilibria P can be computed for every θ and compared, then the estimator will be consistent, asymptotically normal, and efficient (Aguirregabiria and Mira (2007)). For small n, this computation may be plausible; however, the ability to compute all equilibria (or reliance on uniqueness) is a strong assumption and will depend on the application.

4.2.2 Two-Step Estimation

The second approach we propose does not require recomputing the equilibrium, nor does it require an assumption on the uniqueness of equilibria. Rather, it assumes that data is generated by a single Markov equilibrium. This method follows closely the two-step approach for estimating dynamic games (c.f. Hotz and Miller (1993), Benkard, Bajari and Levin (2007)): first, policy functions are estimated off of observed data and used to estimate both equilibrium value functions as well as value functions from unilateral deviations by a single agent; second, estimates of structural parameters are obtained by assuming observed data is generated from equilibrium play, and minimizing violations of equilibrium restrictions. There are two substantive differences in applying the approach to our setting. First, we nest the estimation of value functions within a fixed-point routine in order to determine the set of transfers which are consistent with these value functions (as transfers are assumed not to be observed in the data, yet are anticipated by agents for any course of play). Second, as opposed to forming moments based on differences between value functions from observed policies and alternative policies as in Benkard, Bajari and Levin (2007), we instead compute the optimal policy for each agent for any set of parameters (which is feasible given the nature of our problem), and find the set of parameters which minimize deviations between the computed optimal policies and those observed in the data.

For the expositional purposes, we will describe the estimator for a single market. In this market, we assume we observe a sequence of actions $\{\mathbf{a}_t\}$ for every agent (i.e., whom each agent negotiates with each period), and the sequence of network states $\{g_t\}$. We describe the estimation procedure for a single market, noting that moments can be pooled across markets for inference.

First-stage Policy Estimates and Simulated Value Functions In the first-stage, we obtain estimates of each agents' policy functions $\sigma_i : \mathbf{G} \times \mathbb{R}^{|\mathbf{G}_i|} \to A_i$. These will allow us to estimate value functions for each agent in any given state via forward simulation.

We assume for every agent i, the probability i chooses a given action a_i in each state g, $P_i(a_i|g)$, can either directly be observed or estimated from the data, and the distribution of i's payoff shocks f_i^ϵ is known. In this case, as shown in Hotz and Miller (1993), differences in choice-specific value functions $v_i(a_i, g)$ (given by (6)) can be recovered. For example, if payoff shocks are distributed type I extreme value, then

$$v_i(a, g) - v_i(a', g) = \ln(P_i(a|g)) - \ln(P_i(a'|g))$$

for any two actions $a, a' \in A_i$ and the estimated policy function is given by:

$$\hat{\sigma}_i(g, \epsilon_i) \equiv \arg\max_{a \in A_i} \{v_i(a, g) + \epsilon_{a,i}\} = \arg\max_{a \in A_i} \{\ln(P_i(a|g)) + \epsilon_{a,i}\}$$

Next, given any set of policy functions $\sigma \equiv \{\sigma_i, \sigma_{-i}\}$ (including those estimated from the data), consistent estimates of value functions from agent i playing σ_i and all other agents playing σ_{-i} can be obtained by approximating:

$$\hat{V}_i(g;\sigma;\theta) = \mathbb{E}\left[\sum_{t=0}^{\infty}\beta_i^t\left(\pi_i(g^t,\mathbf{t}) - c(g^t|g^{t-1}) + \epsilon^t_{i,\sigma_i(g^{t-1},\epsilon_i^t)}\right) \mid g^0 = g, g^t = \Gamma\left(\tilde{g}(\sigma(g^{t-1},\epsilon^t))\,;\theta\right)\right]$$

where expectations are over current and future values of private shocks ϵ, Γ and $\mathbf{t}$ are consistent with $\{\hat{V}_i\}$ as specified in (8), and $\sigma(g^{t-1},\epsilon^{t-1})$ dictates the profile of actions taken by all agents for a given vector of shocks. To produce the approximation, we nest the forward simulation procedure of Benkard, Bajari and Levin (2007) into fixed point routine in order ensure the consistency of $\mathbf{t}$ and Γ with respect to the estimated set of value functions. The procedure is as follows:

1. For each state g, fix $\mathbf{t}^0$ and Γ^0 at initial values.
2. For each iteration τ, update $\{\hat{V}_i(g;\sigma;\mathbf{t}^\tau,\Gamma^\tau,\theta)\}_i$ as follows:
   (a) Start at state $g^0 = g$. Draw error shocks ϵ_i^0 for each agent i.
   (b) For each agent i, calculate actions $a_i^0 = \hat{\sigma}_i(g^0,\epsilon_i^0)$. Obtain the predicted negotiation network $\tilde{g}(\mathbf{a}^0)$ and new network $g^1 = \Gamma^\tau(\tilde{g}(\mathbf{a}^0))$. For each i, compute stage profits $\pi_i(g^1,\mathbf{t}^\tau) - c_i(\tilde{g}(\mathbf{a}^0)|g^1) + \epsilon_{a_i^0,i}$.
   (c) Repeat (a)-(b) starting at the new state for an additional T periods.
   (d) Average each agent i's discounted stream of payoffs for multiple simulated paths of play to obtain an estimate of $\hat{V}_i(g;\sigma;\mathbf{t}^\tau,\Gamma^\tau,\theta)$ for each i.
3. Update $\mathbf{t}^{\tau+1}$ and $\Gamma^{\tau+1}$ according to (8), given $\mathbf{b}(\theta)$ and values of $\{\hat{V}_i(g;\sigma;\mathbf{t}^\tau,\Gamma^\tau,\theta)\}_{i,g}$.
4. Repeat steps 2-3 until $|\hat{V}_i(g;\sigma;\mathbf{t}^\tau,\Gamma^\tau,\theta) - \hat{V}_i(g;\sigma;\mathbf{t}^{\tau-1},\Gamma^{\tau-1},\theta)| < \rho$ for all agents and states, where ρ is again some pre-specified tolerance. Denote this value function $\hat{V}_i(g;\sigma;\theta)$.

There are a few crucial points to mention. First, note that for a given Γ and set of actions $\mathbf{a}$, the state transition from g is deterministic and is given by $\Gamma(\tilde{g}(\mathbf{a}))$. If, however, we allow for pair-specific payoff shocks η, then there would be a distribution over potential states that could be reached for a given set of actions, and the state transition probabilities would also have to be estimated. Second, estimation of $\hat{V}_i(g;\sigma;\theta)$ relies not only on Assumption 3.1 in that there exists a unique $\mathbf{t}$ for any set of continuation values $V_i(\cdot)$, but also that for any set of policies σ, there is at most one set of continuation values and transfers that are consistent with one another for a given θ.

Second-stage Parameter Estimation In the second stage, we describe the procedure used to estimate θ. For each candidate θ:

1. For each agent i, compute optimal policy $\tilde{\sigma}_i(\cdot;\theta)$ given all other agents are employing equilibrium strategies $\hat{\sigma}_{-i}$:
   (a) Start with candidate policy $\tilde{\sigma}_i^0 = \hat{\sigma}_i$.

(b) For iteration τ, let $\bar{\sigma}^\tau \equiv (\tilde{\sigma}_i^\tau, \hat{\sigma}_{-i})$. For every state g, obtain estimated value functions $\hat{V}_i(g; \bar{\sigma}^\tau; \theta)$ as described in the first-stage.

(c) Update conditional choice value functions $v_i^{\bar{\sigma}}(\cdot)$ for all actions and states given $\{\hat{V}_i(g; \bar{\sigma}^\tau; \theta)\}$ and updated transfers, which in turn generate updated optimal policies $\tilde{\sigma}_i^{\tau+1}(\cdot)$.

(d) Repeat (b)-(c) until the conditional choice probabilities implied by the optimal policies converge up to a pre-specified tolerance.[23] Let the optimal policy for each agent i (given all other agents play $\hat{\sigma}_{-i}$) be denoted $\tilde{\sigma}_i(\cdot; \theta)$.

2. Obtain an estimate of θ by minimizing the sum of squared differences between each agent's conditional choice probabilities implied by the optimal policy $\tilde{\sigma}_i(\cdot; \theta)$ and the policy $\hat{\sigma}_i$ observed in the data:

$$\hat{\theta} = \arg\min_{\theta} \left\{ \sum_g \sum_i \sum_{a \in A_i} \left(P_i^{\{\tilde{\sigma}_i, \hat{\sigma}_{-i}\}}(a|g) - P_i^{\hat{\sigma}}(a|g) \right)^2 \right\}$$

5 Application: Insurer-Provider Negotiations

To apply our framework, we analyze a stylized network formation game between health insurers (e.g., HMOs) and medical providers (e.g., hospitals), where insurers negotiate with providers over reimbursement rates for serving patients. Health insurers typically offer potential customers access to a "network" of providers: if an insurer and a provider are able to agree on a payment scheme under which the provider is willing to treat the insurer's members, then the provider becomes part of the insurer's "network." In general, little is known about how payments between insurers and providers are negotiated in the private insurance sector, yet those determine over 40% of healthcare spending (approximately $1 trillion). We analyze simulated markets to understand the networks and negotiated payments that arise in a dynamic equilibrium; importantly, we allow for forward looking agents and allow the link structure between insurers and hospitals to change over time.

We first detail the stylized stage game which is used to provide the underlying period-profit functions π accruing to each agent for any given network structure. We then describe summary statistics of simulated equilibria across several markets as bargaining power and the number of agents change, study the relationship between observable market characteristics and negotiated per-patient transfers, detail how the variation in equilibrium network structures across markets can be used to identify bargaining power (and hence transfers) if they are unobserved, and finally, in the next section, examine the impact of hospital mergers on industry profits, negotiated transfers, premiums, and consumer welfare.

5.1 Stage Game Timing

For a given network structure g, the basic timing every period is as follows:

[23] E.g., if ϵ_i is distributed type I extreme value, $P_i^\sigma(a_i|g) = \exp(v_i^\sigma(a_i|g))/(\sum_{a \in A_i} \exp(v_i^\sigma(a|g)))$.

1. HMOs set premiums to consumers based on a constant markup over the expected average cost of an enrollee;[24]

2. Each consumer in a market chooses to join at most one HMO and pays the premium for that insurer;

3. A certain proportion of consumers get sick, and choose to attend their most preferred hospital in their HMO network.

Furthermore, we assume that hospitals must serve any patient who visits, and incurs a constant marginal cost in doing so; any HMO without a hospital on its network does not enroll any patients; and there is an outside option to an HMO that consumers may choose.

In our numerical analysis, we use a discrete choice model of consumer demand for insurance plans and hospitals that follows the structure of Capps, Dranove and Satterthwaite (2003) and Ho (2006); as in Pakes (2010), we draw market characteristics from distributions that mimic those in Ho (2006). Further details, including the exact distributional assumptions for firm and consumer characteristics and specification of profit functions, are provided in the appendix.

5.2 Equilibria: Network Structure and Transfers

We first simulate multiple markets of HMOs and hospitals with random draws from firm and consumer characteristics. We vary the Nash bargaining parameters (referred to here as "bargaining power") so that the division of surplus between agents may differ. We consider 3 different scenarios: equal bargaining power, in which $b_{ij} = .5 \; \forall \; ij$; hospitals having greater bargaining power, given by setting $b_{ij} = .8$ when i is a hospital and .2 otherwise; and HMOs having greater bargaining power, given by $b_{ij} = .8$ when i is an HMO, and .2 otherwise.

Network Distributions Table 1 reports summary statistics of equilibria under different specifications. The first column lists the average number of networks which occur more than 10% of the time in the equilibrium network distribution. Although the number of players and therefore the total number of possible networks increases, the average number of networks remains small. With 3 hospitals and 2 HMO's, there are $2^6 = 64$ potential networks, yet the average number of networks that occur frequently in equilibrium is less than 2.

The second and third columns indicate the frequencies with which the full network and the efficient network occur more than 10% of the time in the equilibrium network distribution, where efficient refers to the network which maximizes *industry* profits (i.e., combined HMO and hospital profits). The probability of a full network being reached never occurs for just one hospital, which is partially due to the fact that industry profits and premiums are higher due to greater downstream insurer differentiation when the hospital is exclusive to one HMO. Indeed, the full network occurs rarely even with multiple hospitals: one contributing factor is there is an incentive to not include

[24]Previous empirical work (e.g., Einav, Finkelstein and Cullen (2010); Handel (2013)) have used similar pricing assumptions.

Table 1: Simulated Equilibrium Network Distributions

	"B-Pow"	# Eq Net	Full Net	Eff. Net	Single (90%)	Single (50%)	Single & Full	Single & Eff	Active Hosp	Exp. Links
1 Hosp	Equal	1.03	0.01	0.88	0.97	1.00	0.01	0.88	1.00	1.00
2 HMOs	Hospitals	1.01	0.00	0.91	0.99	1.00	0.00	0.91	1.00	0.99
	HMOs	1.02	0.00	0.80	0.98	1.00	0.00	0.80	1.00	0.99
2 Hosp	Equal	3.36	0.39	0.90	0.01	0.17	0.04	0.14	2.00	2.65
2 HMOs	Hospitals	3.57	0.22	0.83	0.00	0.23	0.00	0.23	2.00	2.49
	HMOs	2.67	0.01	0.92	0.01	0.73	0.01	0.67	1.99	2.30
3 Hosp	Equal	1.92	0.00	0.72	0.01	0.05	0.00	0.01	2.99	2.88
2 HMOs	Hospitals	1.89	0.00	0.54	0.01	0.15	0.00	0.10	2.94	2.55
	HMOs	1.53	0.00	0.63	0.00	0.45	0.00	0.36	2.91	2.42

Summary statistics from 100 market draws for each specification. *"B-Pow"*: Equal - $b_{ij} = .5 \; \forall \; ij$; Hospitals - $b_i j = .8$ when i is a hospital, .2 otherwise; HMOs - $b_i j = .8$ when i is an HMO, .2 otherwise. *# Eq Net*: Average number of networks that occur more than 10% in the equilibrium network distribution (E.N.D.). *Full Net / Eff Net* : % of runs in which full / efficient network occurs more than 10% in E.N.D. *Single (x%)*: % of runs in which a single network occurs more than x% in E.N.D. *Single & Full / Eff*: % of runs in which a single network occurs more than 90% in E.N.D., and that network is full / efficient. *Active Hosp*: average number of hospitals that have contracts with at least one HMO more than 10% of the time in E.N.D. *Expected Links*: expected number of bilateral links in E.N.D.

high cost hospitals (particularly insofar they lead to higher negotiated prices) since HMOs cannot influence which hospital its own patients visit. In addition, the "efficient" network which maximizes industry profits is not always reached, which should not be surprising given the limited contracting space and presence of contracting externalities.

The third and fourth columns indicate the percentage of markets in which there is a single network structure which occurs more than 90% or 50% of the time in the equilibrium network distribution. Across specifications, this probability falls are more agents are present in the market. The next two columns indicate the percentage of markets in which either the full network or the efficient network occurs at least 50% of the time.

Active Hosp indicates the average number of hospitals that have a contract with at least one HMO 10% of the time in the equilibrium network distribution. Very few markets have one hospital excluded from contracting. Finally, the last column indicates the expected number of links that are sustained in equilibrium.

Clearly, these statistics are dependent on the underlying primitives (e.g., the variance of the idiosyncratic shock will change the number of equilibrium networks and probability of a single network occurring); furthermore, as the number of firms increases, the number of total potential networks and states does as well. Nonetheless, we stress that even with the same underlying primitives, adjusting the Nash bargaining parameters changes equilibrium outcomes in substantive ways, even as the number of agents and network states are held fixed. We continue to explore this point in the following exercises.

Table 2: Regression of Hospital Margins on Observables / Characteristics

Timing:	Dynamic						Static					
	Equal		Hospital		HMO		Equal		Hospital		HMO	
	Coeff	s.e.	Coeff	s.e.	Coeff	s.e.	Coeff	s.e.	Coeff	s.e.	Coeff	s.e.
Const.	-2.40	1.33	0.72	1.43	1.96	1.48	21.77	0.73	23.94	0.63	18.31	0.69
Avg. Cost	-0.94	0.05	-0.96	0.05	-0.77	0.07	-0.65	0.06	-0.56	0.05	-0.70	0.05
Cost-AC	-0.23	0.07	-0.20	0.07	0.10	0.10	-0.23	0.08	-0.36	0.07	-0.16	0.07
# Patient	-0.01	0.08	0.05	0.06	0.18	0.10	0.41	0.05	0.38	0.05	0.31	0.06
Total # Patients	-0.04	0.04	-0.11	0.03	-0.12	0.05	-0.30	0.03	-0.27	0.02	-0.31	0.02
HMO Marg	12.03	0.52	11.58	0.49	8.67	0.68	2.04	0.33	1.66	0.27	3.86	0.37
R^2	0.77		0.79		0.50		0.57		0.62		0.65	

Projection of simulated equilibrium expected per-patient margins between hospital i and HMO j onto equilibrium market observables as bargaining power varies (Equal - $b_{ij} = .5 \ \forall \ ij$; Hospitals - $b_ij = .8$ when i is a hospital, .2 otherwise; HMOs - $b_ij = .8$ when i is an HMO, .2 otherwise). Results pool across 2x2 and 3x2 settings. Av. Cost: average hospital marginal cost in the market; Cost-AC: difference between hospital's marginal cost and average cost in the market; # Patient (Total # Patients): expected number of patients of HMO j (from all HMOs) served by hospital i; HMO Marg: expected HMO margins (premiums minus marginal cost). Extra Hospital: indicator for whether there are 3 hospitals (instead of 2) in the market.

Predicted Transfers We next examine equilibrium hospital per-patient margins computed across different specifications, and project them on market characteristics, where margins are defined as per-patient transfers from HMOs to hospitals minus hospital costs for serving a patient. This exercise is in the spirit Ho (2009) and Pakes (2010) to examine the role of various factors in determining negotiated per-patient transfers between hospitals and insurers.

In our dynamic model, we focus on the expected per-patient margins received by each hospital from each HMO markets with 2 HMOs and either 2 hospitals (where expectations are taken over the equilibrium network distribution). We also examine a static specification ($\beta = 0$) where agents do not anticipate future changes to the network, and disagreement points are given by static profits in the the stable network (i.e., a network in which there exist gains to trade between all contracting agents) which arises if two agents fail to contract. A purely-static model, importantly, cannot determine which network arises in equilibrium, and only can determine which networks are stable. We nonetheless use our dynamic network formation model to determine the equilibrium network distribution given $\beta = 0$ to obtain expected values for margins and other market observables.

Table 2 reports results. We first focus on the results from the dynamic specification. As expected, hospitals receive higher per-patient margins (on average) when they have higher bargaining power, and lower when HMOs have higher bargaining power. Both being in a market with higher average costs and having higher costs than one's competitor negatively impacts predicted hospital margins (though the latter effect is not significant when HMOs have greater bargaining power). Furthermore, consistent with previous findings and anecdotal evidence on quantity discounts, the number of total patients served by a hospital reduces margins. Generally, higher HMO margins are also associated with higher negotiated hospital margins.

The static specification generally shares similar signs as the dynamic model, with significant differences in magnitudes. In particular, it underestimates the sensitivity of negotiated transfers on average market costs. A static model also predicts a greater effect of having a greater number of patients served from a given HMO (holding fixed total patients served); this measure proxies for the HMO having a worse outside option upon losing that hospital, which is mitigated in a dynamic model as an HMO can either recontract with that hospital in the future, or with others. This is consistent with the idea that by anticipating future renegotiation and continuation values, hospitals' outside options are slightly weaker and HMO outside options are slightly stronger in a dynamic setting; indeed, once HMOs have greater bargaining power, a dynamic model predicts much lower hospital margins.[25] Indeed, we find, a static model predicts 5% - 19% higher transfers on average than a dynamic model depending on whether there is equal or HMOs have greater bargaining power.

5.3 Estimation of Nash Bargaining Parameters

One important takeaway from the simulated equilibrium network distributions is that differences in the allocation of bargaining power has a significant impact on equilibrium networks. In this section, we describe how observed network variation allows for the estimation of unobserved Nash bargaining parameters and associated transfers.

We focus on 2x2 markets and assume Nash bargaining parameters are parameterized by $\theta = b_H \in [0,1]$, where $b_{ij} = b_H$ if i is a hospital and $b_{ij} = (1 - b_H)$ otherwise. For all other parameters used to generate the data, we assume them to be known.

Full Equilibrium Computation We observe 20 network configurations $\{g^{m,1}, \ldots, g^{m,20}\}$ for each market m drawn from the equilibrium steady state distribution (we do not assume that these networks need be sequentially observed). We choose the value of b_H which maximizes the probability of observing the set of networks in the data (where the likelihood is given by (20) for $i = 1$). In the absence of a proof of uniqueness, we are relying on a strong assumption of either the uniqueness of the equilibrium network distribution across all equilibria, or the ability to select the same equilibrium in the presence of multiplicity during the estimation routine. We make the latter assumption given the estimation routine utilizes the same computational algorithm as the data generating process. The degree to which this is problematic will depend on the application.

Table 3 summarizes the estimation results as we vary the size of the sample from 1, 5, and 10 different markets; for computational convenience, we conduct a grid search over $[0,1]$ where we allow b_H to vary by .05. With only a single market per sample, the 95% confidence interval is not particularly informative; this partly occurs because there are some markets in which a single network exists, which is consistent with a wide range of values for b_H. However, once the sample

[25] This comes from the demand specification, in which consumers perceive HMOs are more or less bundles of hospitals. E.g., consider the full network between 2 HMOs and 2 hospitals; if an HMO and hospital fail to come to an agreement, the HMO may lose many of its patients to the other HMO while the hospital may not lose many patients at all. Thus, accounting for dynamics—and the potential that the HMO can recontract again with the hospital—strengthens its outside option.

Table 3: Monte Carlo Estimates of b_H

	True b_H	1 Markets / Sample	5 Markets / Sample	10 Markets / Sample
Avg. Estimate:	0.50	0.48	0.47	0.51
95% C.I.:		(0.10,0.90)	(0.20,0.70)	(0.40,0.60)
Avg. Estimate:	0.80	0.60	0.76	0.77
95% C.I.:		(0.10,0.90)	(0.40,0.90)	(0.60,0.80)
Avg. Estimate:	0.20	0.20	0.24	0.23
95% C.I.:		(0.10,0.40)	(0.20,0.50)	(0.20,0.30)

Estimated values of hospital bargaining power b_H for 40 samples of either 1, 5, or 10 markets in 2x2 settings where a sequence of 20 networks were observed. Grid search conducted over b_H in increments of .05.

size increases to include 5 markets, fewer values of b_H are consistent with generating the observed sequence of networks in the data. As a result, the average estimates become extremely close to the true value of b_H; once 10 markets are used, the mean is within .02 and the 95% confidence interval is within .15 of the true value.

Two-step Estimation We examine 40 samples consisting of one market observation, varying hospital bargaining powers as in the last example, and assumed conditional choice probabilities $\{P_i(a|g)\}$ were observed (in addition to market level primitives).[26] In all sample runs, using the algorithm described in the previous section, we exactly recovered b_i (on a .05 grid search) with just one market, even though negotiated transfers—necessary to construct agent value functions—were unobserved. Thus, presuming the optimality of observed conditional choice probabilities within a market provides a great deal of information which can be used to infer both unobserved transfers and relative bargaining power.

6 Application: Hospital Mergers

Consolidation in the US healthcare delivery system has increased dramatically in recent years, and anti-trust regulators have become increasingly concerned with market concentration leading to decreased consumer welfare. However, regulators have historically had difficulty challenging mergers. A key question in such anti-trust analyses is whether the potential benefits of consolidation, such as increased efficiency, reduced excess capacity, lower transactional frictions, and higher risk tolerance outweigh the potential costs of increased market power. The methods introduced in this paper provide a novel way of analyzing these effects by explicitly account for the renegotiation of transfers and contracts following a merger.

In this section, we simulate counterfactual equilibrium networks and negotiated transfers subsequent to hypothetical hospital mergers. In the current analysis, we focus on simulated markets with 2 hospitals and 2 HMOs, and examine what occurs when hospitals merge *exogenously* into one

[26]We do not consider first-stage estimation error (introduced when recovering conditional choice probabilities from the data) in this analysis.

Table 4: Merger Simulations

	"B-Pow"	$+\Delta\pi^H$	$-\Delta\pi^H_{5\%}$	$+\Delta\pi^M$	$-\Delta\pi^M_{5\%}$	$+p^M$	$-p^M_{5\%}$	+Ins	$-\text{Ins}_{5\%}$
(i) Dynamic	Equal	0.72	0.28	0.73	0.25	0.81	0.14	0.19	0.76
	Hospitals	0.59	0.29	0.12	0.29	0.75	0.20	0.25	0.71
	HMOs	0.80	0.17	0.76	0.24	0.85	0.11	0.15	0.77
(ii) Dynamic,	Equal	-	-	0.97	0.01	0.99	0.00	0.01	0.99
$+\Delta\pi^H \geq 0$	Hospitals	-	-	0.15	0.07	1.00	0.00	0.00	0.95
	HMOs	-	-	0.89	0.11	0.99	0.00	0.01	0.90
(iii) Static	Equal	0.12	0.85	0.02	0.91	1.00	0.00	0.00	1.00
	Hospitals	0.04	0.87	0.01	0.98	1.00	0.00	0.00	1.00
	HMOs	0.25	0.71	0.02	0.87	1.00	0.00	0.00	1.00
(iv) Static,	Equal	-	-	0.17	0.25	1.00	0.00	0.00	1.00
$+\Delta\pi^H \geq 0$	Hospitals	-	-	0.25	0.50	1.00	0.00	0.00	1.00
	HMOs	-	-	0.08	0.52	1.00	0.00	0.00	1.00

Summary statistics from merger simulations, where: (i) and (ii) are from a dynamic model ($\beta = .9$), (iii) and (iv) from a static model, and (ii) and (iv) condition also on markets where hospitals find it profitable to merge. *"B-Pow"*: Equal - $b_{ij} = .5 \; \forall \; ij$; Hospitals - $b_i j = .8$ when i is a hospital, .2 otherwise; HMOs - $b_i j = .8$ when i is an HMO, .2 otherwise. $+\Delta\pi^H, -\Delta\pi^H_{5\%}$: percentage of markets in which total hospital profits increases at all or falls by 5%; $+\Delta\pi^M, -\Delta\pi^M_{5\%}$: percentage of markets in which total HMO profits increases at all or falls by 5%; $+p^M, -p^M_{5\%}$: percentage of markets in which both HMO premiums increase or fall by 5%; $+Ins, -Ins_{5\%}$: percentage of markets in which total patients insured increases at all or falls by 5%.

"hospital system" per market. We model mergers in the following a stylized fashion: upon merging, all primitives in the market remain the same (e.g., hospitals do not realize any cost savings), but equilibrium networks, negotiated transfers, and premiums charged by HMOs can change. We assume hospital systems may negotiate different per-patient contracts for each hospital from each HMO, but crucially internalize the joint payoffs across both hospitals; additionally, HMOs cannot opt to only have one hospital on its network and must either have both or none. We thus attempt isolate the impact of mergers on the bargaining, network, and pricing outcomes without taking a stance on potential cost savings, quality improvements, or other efficiencies. This exercise is a first step towards building a more complete framework for merger analysis.

Results Statistics from merger simulations across 100 simulated markets are shown in Table 4. The table provides the percentage of markets in which: total hospital profits (π^H) or total HMO profits (π^M) increase, or fall by 5%; both HMO premiums (p^M) increase, or fall by 5%; and total patients insured in the market increases, or falls by 5%. We again assume $\beta = .9$. Row (i) examines mergers across all markets as bargaining power varies while (ii) conditions only on those markets in which hospital expected profits increase subsequent to a merger (which can be seen as a proxy for "voluntary" hospital mergers); (iii) and (iv) perform the same exercise in a static model ($\beta = 0$). We first focus on the dynamic specification.

Generally, across all markets, approximately half of all hospital mergers tend to increase hospital profits —this is consistent with the idea that mergers strengthen hospitals' outside options with respect to HMOs', and allows them to extract higher transfers. Similar to the example discussed in

Sections 2 and 3.7, the monopolist upstream hospital system can play the two downstream HMOs off one another.

However, what is surprising is that hospitals may be worse off from merging and, if given the choice, may prefer not to merge; this can happen under all three bargaining power specifications, but occurs nearly half the time when hospitals have greater bargaining power. To understand why hospital profits can fall, it is worth noting that mergers which lead to reduced hospital profits occur mainly in markets in which either one or both hospitals were excluded pre-merger from at least one HMO network. This leads to a natural explanation: we find that in markets where one hospital was excluded pre-merger, it is predominately the high cost hospital; post-merger, since an HMO would have to include both hospitals on its network in order to have any patient demand, utilization of the higher cost hospital would rise as it no longer can be excluded. Key to this effect is the inability of HMOs to steer patients towards the lower cost hospital if it had both hospitals on its network.[27] When hospitals already have greater bargaining power, the gains from merging through outside options are smaller and more likely to be offset by inefficient utilization.

Another interesting observation is that when hospitals have greater bargaining power, HMOs are generally opposed to hospital mergers as their profits fall. However, under equal bargaining power or when HMOs have greater bargaining power, HMO profits tend to increase when hospitals merge—as a result, HMOs may actually not oppose hospital mergers in many situations. This may be an artifact of the fixed markup rule: HMO per-patient premiums and hence expected profits are increasing in the transfers paid to hospitals, while transfers paid to hospitals are increasing in hospital bargaining power. Thus, insofar HMO markups were set too low, slightly increasing transfers and premiums can lead to higher HMO profits. But when transfers become too high (as hospital bargaining power increases), HMO profits fall.

We next examine the impact of mergers on other market outcomes, particularly those that influence consumer welfare. First, note hospital mergers may lead to more patients being insured in a market—this primarily occurs as the result of a previously excluded hospital (e.g., a hospital that has only one active contract with an HMO) being included on both HMOs as a consequence of the merger. Secondly, mergers may also occasionally lead to reduced premiums. Although this may arise due to fiercer price competition among HMOs, we have assumed a simple pricing model (fixed markups for premiums) which does not admit that explanation here. Rather, that premiums can fall after hospitals merge—particularly when hospitals have greater bargaining power—can be explained by the fact that hospitals better internalize the impact of increasing their own per-patient rates on total patient flows upon merging. I.e., if a hospital increased its rates for a given insurer, premiums would rise and fewer patients would be insured; pre-merger a hospital only would only care about the reduction in their own patient volume. Consequently, if hospitals possessed greater bargaining power and hence captured a larger share of industry rents, they would have an incentive to lower negotiated per-patient rates in order to increase total patient volume. Nonetheless, for

[27] This has been assumed in the current specification; recent work has studied the ability of insurers to direct patients to certain providers via physician incentives (Ho and Pakes, 2013).

those mergers in which hospitals find it profitable to merge, this effect does not often occur; in general, mergers seem to predominantly lead to higher premiums and fewer patients insured.

Comparison to Static Model We next turn to the simulation results from a static merger model, where both pre and post merger profits are computed assuming $\beta = 0$. Here, negotiated transfers do not internalize what transfers would be in alternative network structures, nor anticipate future changes to the network.

The most striking result, seen in row (iii) from Table 4, is that the vast majority of potential hospital mergers are predicted to lower hospital profits. A static model assumes a merged hospital system receives 0 patients from an HMO upon disagreement; since it does not allow for renegotiation, there is no ability for the hospital to have the HMOs compete against one another when determining rates (as in our stylized example discussed in the previous section). Hence, merging would be predicted to harm hospitals' outside options by reducing their flexibility to selectively contract without providing any significant benefit. This is suggestive that static models understimate the incentives for consolidation by medical providers and understate the degree to which industry rents are captured by the more concentrated part of the market.

Summary In this exercise, there are several takeaways. First, the presumption that voluntary hospital mergers tend to reduce consumer welfare absent cost savings and quality improvements seems to have merit; though under certain conditions, consumer welfare can actually increase post-merger (though not when such mergers are voluntary), this occurs infrequently. Second, HMOs and hospitals do not necessarily have opposed interests when evaluating mergers; as shown, with fixed markups, industry profits can rise for both sides of the market while consumer welfare falls. Third, static bargaining models fail to adequately capture hospital incentives for merging, as the vast majority of simulated markets would predict that hospitals do not wish to merge.

7 Concluding Remarks

We have developed a model of dynamic network formation and bargaining in the presence of externalities for use in applied work, and explored its usefulness in a stylized model of health insurance-hospital negotiations. Dynamics are important to consider in such industries where agents interact repeatedly and networks change over time, and incorporating them yields substantively different predictions than static models. We have demonstrated the feasibility of recovering estimates of unobserved bargaining power and transfers using only observed equilibrium networks or actions. The framework is useful for understanding equilibrium surplus division and network formation in bilateral oligopoly, and can help analyze potential policy changes or mergers by predicting future network changes and recontracting decisions among firms. Finally, we have stressed the importance of dynamic considerations in the context of hospital mergers: static approaches drastically understate the incentives for hospitals to consolidate by underestimating the extent to which the concentrated side of the market can capture surplus.

A HMO Hospital Application

A.1 Model Preliminaries

We analyze a market with M HMO plans, H hospitals, and C consumers. Each HMO j and hospital k possesses a vector of characteristics θ_j^M and θ_k^H respectively. Individuals are divided among R different demographic groups, where group r makes up share σ_r^R of the population and values HMO and hospital characteristics according to the coefficients β_r^M and β_r^H respectively. We assume any hospital can contract with any number of different HMOs, and similarly any HMO can contract with any number of hospitals: i.e., $\mathbf{G} \equiv \{0,1\}^{M\times H}$. Recall $N_k(g)$ denotes the set of HMOs of which that hospital k is a member; similarly, $N_j(g)$ represent the set of hospitals that are in HMO j's network of providers.

Individual Choice We assume every individual will be hospitalized with probability γ. If sick, in order to use a particular hospital $k \in H$, an individual needs to have enrolled in an HMO plan j with $j \in N_k(g)$ if the current network structure is g. Each HMO j charges a premium to its members $p_j(g)$. There is also an outside option which provides the individual with necessary health care in the case of illness – the utility of this option is normalized to 0.

Let individual i be part of demographic group r. We define an individual i's utility from using hospital k as

$$u_{i,k}^H = \alpha_r^H \theta_k^H + \omega_{i,k} \tag{21}$$

where ω is distributed iid Type I extreme value. From this formulation, we can define an individual i's utility from enrolling in a given HMO j that has a set of hospitals $N_j(g)$:[28]

$$u_{i,k}^M(g) = \alpha_r^M \theta_k^M - \alpha^P p_k(g) + \gamma \left(ln(\sum_{h\in N_k(g)} exp(\alpha_r^H \theta_h^H)) \right) + \varepsilon_{i,k} \tag{22}$$

where ε is also distributed iid Type I extreme value.

With this linear utility function and distribution on error terms, we can calculate the the (expected) share of the population that chooses HMO j given any particular network structure g as follows:

$$\sigma_j^M(g) = \sum_{r\in R} \sigma_r^R \frac{\exp(\alpha_r^M \theta_j^M - \alpha p_j + \gamma \ln(\sum_{h\in N_j(g)} \exp(\alpha_r^H \theta_h^H)))}{1 + \sum_{m\in M}(\exp(\alpha_r^M \theta_l^M - p_l + \gamma \ln(\sum_{h\in N_m(g)} \exp(\alpha_r^H \theta_h^H))))} = \sum_{r\in R} \sigma_r^R \tilde{\sigma}_{j,r}^M(g)$$

We use $\tilde{\sigma}_{j,r}^M(g)$ to represent the share of demographic group r that chooses HMO plan j.

Define the demographic distribution of individuals within each HMO plan (i.e., the share of people who use HMO plan j who are part of demographic group r) as follows:

$$\tilde{\sigma}_{r,j}^R(g) = \frac{\sigma_r^R \tilde{\sigma}_{j,r}^M(g)}{\sum_{s\in R} \sigma_s^R \tilde{\sigma}_{j,s}^M(g)}$$

Thus, the share of HMO plan j's customers who actually will be sick and need to use hospital $k \in N_j(g)$ can be written as:

$$\sigma_{k,j}^H(g) = \gamma \sum_{r\in R} \tilde{\sigma}_{r,j}^R(g) \frac{\exp(\alpha_r^H \theta_k^H)}{\sum_{h\in N_j(g)} \exp(\alpha_r^H \theta_h^H)}$$

Note that although σ_j^M is a function of the entire network structure g and premiums charged by all HMOs, $\sigma_{k,j}^H$ will just be a function of HMO j's own hospital network $N_j(g)$.

Hospital and HMO Per-Period Profits For expositional convenience, let σ_j^M and σ_k^H will denote values for a given network structure g unless otherwise specified. Let t_{jk} be the negotiated per-patient transfers between hospital k and j in network structure g. For hospital k, profits for a given network structure g are given by the equation

$$\pi_k^H(g) = (t_{jk}(g) - mc_k^H)(\sum_{j\in N_h(g)} N\sigma_j^M \sigma_{k,j}^H)$$

mc_k^H represents the marginal cost of serving each patient at hospital k.

[28]Note $E_\omega(\max_{h\in N_j(g)}(u_{i,h}^H)) = ln(\sum_{k\in N_j(g)} \exp(\alpha_r^H \theta_k^H))$

For any HMO j, its profits are 0 if it has no hospitals (i.e., $N_j(g) = \{\}$); otherwise, profits are:

$$\pi_j^M(g) = N\sigma_j^M(p_j - mc_j^M - \sum_{k \in N_j(g)} \sigma_{k,j}^H t_{j,k}(g)) \qquad \text{if } N_j(g) \neq \{\}$$

where mc_j^M represents the marginal cost of insuring a patient for insurer j.

Timing The timing in each period, given a network structure g, is as follows:

1. Each HMO j chooses a premium $p_j(g) \in \mathbb{R}^+$ that it will charge each consumer that chooses to join its plan. We assume premiums are set to be 15% above the average costs of serving a patient.
2. Each individual i chooses to enroll in an HMO plan, with the utility from choosing HMO j being $u_{i,j}^M$ defined in (22), or chooses to utilize the outside option, thereby deriving a utility of 0.
3. γ proportion of the population becomes sick. Each individual i that is sick chooses the best hospital on the HMO they enrolled in to attend according to their utility given by (21).
4. HMO payoffs ($\pi^M(g)$) and hospital payoffs ($\pi^H(g)$) are realized.

A.2 Parameters

Units, unless otherwise specified, are in thousands.

Demographic Characteristics People are hospitalized with probability $\gamma = .075$. Market size is distributed normally with a mean of 500, standard deviation 300, and a minimum value of 100; thus, if everyone in the mean market subscribes to an HMO plan, 37.5 patients will need to be served. In the current specification, we do not assume there are different demographic groups, and assume $\alpha_r^M = \alpha_r^H = \alpha^P = 1$ for all agents.

HMO & Hospital Characteristics HMO per-patient costs c_j^M are normally distributed with mean .75 and standard deviation .25. HMO quality θ_j^M is distributed normally with mean 0 and standard deviation .25. It is correlated with costs by a value of $\rho^M = .5$. Hospital quality θ_h^H for each hospital are normals with mean $\mu_{\theta^H} = 50$ and standard deviation $\sigma_{\theta^H} = .5$. Hospital constant marginal costs $\bar{c}_j^H$ are normally distributed with mean $\mu_{\bar{c}^H} = 11$ and standard deviation $\mu_{\bar{c}^H} = 3$, with a minimum of 2. Costs and hospital quality index for a particular hospital h are correlated by a value of $\rho^H = .5$. For each pair, these variables are generated first by creating two correlated standard normal random variables, and then appropriately transforming them with the correct mean and standard deviation.

Parameters of the Dynamic Game We assume proposer shocks ϵ_i are distributed iid Type I extreme value with variance $20\pi^2/6$, no profit shocks η (i.e., $\eta_{i,j} = 0$), a common discount factor $\beta = .9$, and costs for forming a new link to be 100, with no costs for breaking a link.

References

Abreu, Dilip, and Mihai Manea. 2013*a*. "Bargaining and Efficiency in Networks." *Journal of Economic Theory*. Forthcoming.

Abreu, Dilip, and Mihai Manea. 2013*b*. "Markov Perfect Equilibria in a Model of Bargaining in Networks." Forthcoming.

Ackerberg, Daniel, C. Lanier Benkard, Steven Berry, and Ariel Pakes. 2007. "Econometric Tools for Analyzing Market Outcomes." In *Handbook of Econometrics*. Vol. 6a, , ed. James .J. Heckman and Edward E. Leamer, Chapter 63. Amsterdam:North-Holland.

Aguirregabiria, Victor, and Pedro Mira. 2002. "Swapping the Nested Fixed Point Algorithm: A Class of Estimators for Discrete Markov Decision Models." *Econometrica*, 70(4): 1519–1543.

Aguirregabiria, Victor, and Pedro Mira. 2007. "Sequential Estimation of Dynamic Discrete Games." *Econometrica*, 75(1): 1–53.

Benkard, C. Lanier, Patrick Bajari, and Jonathan Levin. 2007. "Estimating Dynamic Models of Imperfect Competition." *Econometrica*, 75(5): 1331–1370.

Bernheim, B. Douglas, and Michael D. Whinston. 1998. "Exclusive Dealing." *Journal of Political Economy*, 106(1): 64–103.

Binmore, Ken. 1987. "Perfect Equilibria in Bargaining Models." In *The Economics of Bargaining.* , ed. Ken Binmore and Partha Dasgupta. Basil Blackwell, Oxford.

Binmore, Ken, Ariel Rubinstein, and Asher Wolinsky. 1986. "The Nash Bargaining Solution in Economic Modelling." *RAND Journal of Economics*, 17(2): 176–188.

Bloch, Francis, and Matthew O. Jackson. 2007. "The Formation of Networks with Transfers among Players." *Journal of Economic Theory*, 127(1): 83–110.

Capps, Cory, David Dranove, and Mark Satterthwaite. 2003. "Competiton and Market Power in Option Demand Markets." *RAND Journal of Economics*, 34(4): 737–763.

Collard-Wexler, Allan, Gautam Gowrisankaran, and Robin S. Lee. 2013. "Bargaining in Bilateral Oligopoly: An Alternating Offers Representation of the "Nash-in-Nash" Solution." Unpublished.

Corominas-Bosch, Margarida. 2004. "Bargaining in a network of buyers and sellers." *Journal of Economic Theory*, 115: 35–77.

Crawford, Gregory S., and Ali Yurukoglu. 2012. "The Welfare Effects of Bundling in Multichannel Television." *American Economic Review*, 102(2): 643–685.

Cremer, Jacques, and Michael H. Riordan. 1987. "On Governing Multilateral Transactions with Bilateral Contracts." *RAND Journal of Economics*, 18(3): 436–451.

de Fontenay, Catherine C., and Joshua S. Gans. 2013. "Bilateral Bargaining with Externalities." *Journal of Industrial Economics.* Forthcoming.

Doraszelski, Ulrich, and Ariel Pakes. 2007. "A Framework for Applied Dynamic Analysis in IO." In *Handbook of Industrial Organization.* Vol. 3, , ed. Mark Armstrong and Rob Porter, 1887–1966. Amsterdam:North-Holland.

Dranove, David, Mark Satterthwaite, and Andrew Sfekas. 2008. "Boundedly Rational Bargaining in Option Demand Markets: An Empirical Application." Unpublished.

Dutta, Bhaskar, Sayantan Ghosal, and Debraj Ray. 2005. "Farsighted Network Formation." *Journal of Economic Theory*, 122: 143–164.

Einav, Liran, Amy Finkelstein, and Mark R. Cullen. 2010. "Estimating Welfare in Insurance Markets Using Variation in Prices." *Quarterly Journal of Economics*, 75(3): 877–921.

Elliot, Matthew. 2013. "Inefficiencies in NeNetwork Markets." Unpublished.

Ericson, Richard, and Ariel Pakes. 1995. "Markov-Perfect Industry Dynamics: A Framework for Empirical Work." *Review of Economic Studies*, 62: 53–82.

Fershtman, Chaim, and Ariel Pakes. 2012. "Dynamic Games with Asymmetric Information: A Framework for Empirical Work." *Quarterly Journal of Economics*, 127: 1611–1661.

Galeotti, Andrea, Sanjeev Goyal, and Jurjen Kamphorst. 2006. “Network Formation with Heterogeneous Players.” 54: 353–372.

Gomes, Armando. 2005. “Multilateral Contracting with Externalities.” *Econometrica*, 73(4): 1329–1350.

Gomes, Armando, and Pehilippe Jehiel. 2005. “Dynamic Processes of Scoial and Economic Interactions: On the Persistence of Inefficiencies.” *Journal of Political Economy*, 113(3): 626–667.

Gowrisankaran, Gautam. 1999. “A Dynamic Model of Endogenous Horizontal Mergers.” *RAND Journal of Economics*, 30: 56–83.

Gowrisankaran, Gautam, and Thomas J. Holmes. 2004. “Mergers and the Evolution of Industry Concentration: Results From the Dominant Firm Model.” *RAND Journal of Economics*, 35: 561–582.

Gowrisankaran, Gautam, Aviv Nevo, and Robert Town. 2013. “Mergers When Prices Are Negotiated: Evidence from the Hospital Industry.” Unpublished.

Grennan, Matthew. 2013. “Price Discrimination and Bargaining: Empirical Evidence from Medical Devices.” *American Economic Review*, 103(1): 147–177.

Handel, Benjamin. 2013. “Adverse Selection and Switching Costs in Health Insurance Markets: When Nudging Hurts.” *American Economic Review*. Forthcoming.

Hart, Oliver, and Jean Tirole. 1990. “Vertical Integration and Market Foreclosure.” *Brookings Papers on Economic Activity, Microeconomics*, 1990: 205–286.

Ho, Kate, and Robin S. Lee. 2013. “Insurer Competition and Negotiated Hospital Prices.” Unpublished.

Ho, Katherine. 2006. “The Welfare Effects of Restricted Hospital Choice in the US Medical Care Market.” *Journal of Applied Econometrics*, 21(7): 1039–1079.

Ho, Katherine. 2009. “Insurer-Provider Networks in the Medical Care Market.” *American Economic Review*, 99(1): 393–430.

Ho, Katherine, and Ariel Pakes. 2013. “Physician Responses to Financial Incentives: Evidence from Hospital Discharge Data.” Unpublished.

Horn, Henrick, and Asher Wolinsky. 1988. “Bilateral Monopolies and Incentives for Merger.” *RAND Journal of Economics*, 19(3): 408–419.

Hotz, V. Joseph, and Robert A. Miller. 1993. “Conditional Choice Probabilities and the Estimation of Dynamic Models.” *Review of Economic Studies*, 60(3): 497–529.

Hotz, V. Joseph, Robert A. Miller, Seth Sanders, and Jeffrey Smith. 1994. “A Simulation Estimator for Dynamic Models of Discrete Choice.” *Review of Economic Studies*, 61(3): 265–289.

Inderst, Roman, and Christian Wey. 2003. “Bargaining, Mergers, and Technology Choice in Bilaterally Oligopolistic Industries.” *RAND Journal of Economics*, 34(1): 1–19.

Iozzi, Alberto, and Tommaso Valletti. 2013. “Vertical Bargaining and Countervailing Power.” *American Economic Journal: Microeconomics*. Forthcoming.

Jackson, Matthew O. 2004. “A Survey of Models of Network Formation: Stability and Efficiency.” In *Group Formation in Economics; Networks, Clubs and Coalitions.* , ed. Gabrielle Demange and Myrna Wooders, Chapter 1. Cambridge, U.K.:Cambridge University Press.

Jackson, Matthew O. 2005. “Allocation Rules for Network Games.” *Games and Economic Behavior*, 51(1): 128–154.

Jackson, Matthew O., and Alison Watts. 2002. “The Evolution of Social and Economic Networks.” *Journal of Economic Theory*, 106: 265–295.

Jackson, Matthew O., and Asher Wolinsky. 1996. “A Strategic Model of Social and Economic Networks.” *Journal of Economic Theory*, 71(1): 44–74.

Judd, Kenneth L. 1998. *Numerical Methods in Economics.* Cambridge, MA:MIT Press.

Kranton, Rachel E., and Deborah F. Minehart. 2001. “A Theory of Buyer-Seller Networks.” *American Economic Review*, 91(3): 485–508.

Manea, Mihai. 2011. “Bargaining in Stationary Networks.” *American Economic Review*, 101: 2042–2080.

Maskin, Eric, and Jean Tirole. 1988. “A Theory of Dynamic Oligopoly: I & II.” *Econometrica*, 56(3): 549–600.

McAfee, R. Preston, and Marius Schwartz. 1994. “Opportunism in Multilateral Vertical Contracting: Nondiscrimination, Exclusivity, and Uniformity.” *American Economic Review*, 84(1): 210–230.

Melo, Emerson. 2013. “Bargaining and Centrality in NeNetwork Markets.” Unpublished.

Merlo, Antonio, and Charles Wilson. 1995. “A Stochastic Model of Sequential Bargaining with Complete Information.” *Econometrica*, 63(2): 371–399.

Milgrom, Paul, and Robert Weber. 1985. “Distributional Strategies for Games with Incomplete Information.” *Mathematics of Operations Research*, 10: 619–32.

Muthoo, Abhinay. 1999. *Bargaining Theory with Applications.* Cambridge.

Myerson, Roger B. 1977. “Graphs and Cooperation in Games.” *Mathematics of Operations Research*, 2(3): 225–229.

Myerson, Roger B. 1991. *Game Theory: Analysis of Conflict.* Cambridge, MA:Harvard University Press.

Pakes, Ariel. 2010. “Alternative Models for Moment Inequalities.” *Econometrica.*

Pakes, Ariel, and Paul McGuire. 1994. “Computing Markov Perfect Nash Equilibrum: Numerical Implications of a Dyamic Differentiated Product Model.” *RAND Journal of Economics*, 25(4): 555–589.

Pakes, Ariel, and Paul McGuire. 2001. “Stochastic Algorithms, Symmetric Markov Perfect Equilibrum, and the ‘curse’ of dimensionality.” *Econometrica*, 69(5): 1261–1281.

Polanski, Arnold. 2007. “A Decentralized Model of Information Pricing in Networks.” *Journal of Economic Theory*, 136: 497–512.

Prat, Andrea, and Aldo Rustichini. 2003. "Games Played Through Agents." *Econometrica*, 71(4): 989–1026.

Rubinstein, Ariel. 1982. "Perfect Equilibrium in a Bargaining Model." *Econometrica*, 97–109.

Segal, Ilya. 1999. "Contracting with Externalities." *Quarterly Journal of Economics*, 64(2): 337–388.

Segal, Ilya, and Michael D. Whinston. 2003. "Robust Predictions for Bilateral Contracting with Externalities." *Econometrica*, 71(3): 757–791.

Shapley, Lloyd S. 1953. "A Value for n-person Games." In *Contributions to the Theory of Games.* Vol. 2, , ed. H. W. Kuhn and A. W. Tucker. Princeton University Press.

Sorensen, Alan T. 2003. "Insurer-Hospital Bargaining: Negotiated Discounts in Post-Deregulation Connecticut." *Journal of Industrial Economics*, 51(4): 469–490.

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.

Watts, Alison. 2001. "A Dynamic Model of Network Formation." *Games and Economic Behavior*, 34: 331–341.