Markov-Perfect Network Formation

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

Robin S. Lee† and Kyna Fong‡

April 19, 2011

Abstract

We develop a dynamic model of network formation with transfers, and detail the specification and computation of a Markov-Perfect equilibrium. Applications include the study of bilateral oligopoly and buyer-seller networks. The framework takes as primitives each agent's static profits as a function of the current network state, and computes both the equilibrium recurrent class of networks and transfers between agents. In an example drawn from healthcare, we model the negotiations between insurance firms and hospitals. Simulations reveal the efficient network may not always be realized in equilibrium, and that static Nash equilibrium conditions leveraged in previous applied work may not hold in dynamic environments leading to different predicted determinants of negotiated payments. We also demonstrate how agents' "bargaining power" and associated transfers between agents can be partially identified from observing equilibrium network structures. Finally, we use the framework to explore the impact of hospital mergers on negotiated payments, insurer premiums, and consumer welfare controlling for subsequent changes in the network and contracting partners.

1 Introduction

Network formation, or the creation of relationships, trading partnerships, or *links*, is a general phenomenon underlying many areas of economic exchange and interaction. Capturing the notion that the value of a relationship between two parties can be influenced by relationships involving other parties is crucial in understanding common economic settings such as buyer-seller networks and bilateral oligopoly. Justifiably, given their importance, networks have attracted a great deal of interest and research; however, there have been few models that explicitly yield sharp predictions about equilibrium network structures and the rents that accrue to agents when there are

*We thank John Asker, Adam Brandenburger, Allan Collard-Wexler, Ignacio Esponda, Rachel Kranton, Ariel Pakes, Andy Skryzpacz, and seminar participants at Duke, NYU Stern, UCLA, U. Toronto, U. Warwick, the UBC Summer IO Conference, and the Utah Winter Business Economics Conference for helpful comments. All errors are our own.

†Stern School of Business, New York University. Contact: rslee@stern.nyu.edu.

‡Department of Economics, Stanford University. Contact: kyna.fong@stanford.edu.

externalities and transfers. The complexity of network formation, especially when relationships are constantly renegotiated in a dynamic setting, have so far limited previous approaches.

This paper offers a new dynamic model of network formation with endogenous transfers. In light of the complexity of the problem, the primary emphasis lies on tractability and computability. Our goal is to provide an applied framework for the analysis and computation of both *equilibrium surplus division* and *network composition* when agents engage in repeated interaction. Our primary application is bilateral contracting between firms in vertical markets (e.g., negotiations between: manufacturers and retailers; health insurers and medical providers; content providers and distributors; or software developers and hardware providers), but the framework can be extended to other settings (e.g., horizontal mergers or alliances). By providing predictions as to who contracts with whom, the model is also useful for counterfactual analysis and can be used to predict how trading or contracting relationships between firms might change as a result of a policy intervention, merger, or other industry shock. Furthermore, we demonstrate how dynamics can help (partially) identify both agents' "bargaining power" and contracted transfers – objects which are often unobserved in applied work – by using observed equilibrium network structures.

We consider an infinite horizon, discrete time network formation game between N agents, and do not impose restrictions on the set of feasible networks that may arise. In each period, there is a current *network structure* or *state* reflecting the current set of relationships between agents; additionally, at most one agent is randomly selected to be a "proposer," who may offer to form new links or maintain existing links with other agents. Associated with these offers may be lump-sum transfers, and the proposer may also be required to pay a cost to break existing links. Through this process, the network evolves over time.

Within each period, there is a separate subgame where agents' receive period payoffs as a function of the current network structure. Importantly, the subgame and realized payoffs are assumed to be primitives of the analysis; since these payoffs are a function of the entire network, the formulation admits any general form of externalities across agents. Within the subgame, all agents that are linked engage in bilateral Nash bargaining to determine *period contracts*.[1] Our model thus admits two sources of what may be interpreted as asymmetric "bargaining power." First, the probability with which an agent is selected to be a proposer determines how often he may propose new links or open up negotiation with other agents; a higher probability of being selected gives the agent the ability to move towards more desirable or profitable network structures. Second, each agent has an explicit Nash bargaining parameter which is used when negotiating period contracts, which may reflect her individual "negotiating ability" in capturing a larger share of potential gains from any bilateral trade.

As a solution concept, we restrict attention to Markov-Perfect equilibria (MPE) (Maskin and Tirole, 1988) in which agents' strategies are only a function of the current network state. This allows us to capture the importance of dynamics with forward-looking agents and endogenous outside options while maintaining tractability. One conceptual issue often discussed in the network forma-

[1]E.g, these may represent wholesale prices, capitation rates, carriage fees, royalty payments, etc.

tion and bilateral contracting literatures is the question of which equilibrium concept or refinement is most appropriate, particularly governing beliefs of opponents' actions in out-of-equilibrium play. Commonly used approaches include relying upon static Nash conditions, "passive beliefs," and/or pairwise stability, all of which are inherently static concepts that do not allow for further deviations by others (c.f. McAfee and Schwartz (1994); Jackson and Wolinsky (1996)).[2] Such assumptions are not innocuous, as they determine outside options and disagreement points of bargaining, which in turn affect surplus division and equilibrium play. One of the key features of our model is agents' disagreement points from any bilateral bargain are endogenously determined, and equal the continuation values of moving to a network state without that link. We thus ensure that our model is internally consistent with respect to the complete infinite horizon game, and capture the idea that agents repeatedly renegotiate (albeit perhaps with some delay) while anticipating future changes to the network.

We provide assumptions which guarantee the existence of an MPE, and provide a means of computation. Determining the equilibrium network distribution – i.e., who contracts with whom over time – and the associated division of surplus given static period profits is the first step towards predicting counterfactual network structures in dynamic applied settings.[3] Additionally, as static Nash conditions may fail to hold once dynamics are incorporated, relying on them for inference may result in inaccurate predictions regarding transfers and other underlying fundamentals of interest. Thus, we believe there may be significant value that can be derived from an applicable dynamic framework.

One key observation we make is that the dynamic nature of our model introduces an additional source of identification that is not present in many static settings: bargaining power directly affects agents' outside options and value functions, impacting both which networks are sustainable in equilibrium and the steady state distribution over those networks. Thus, we can estimate values of bargaining power and contracted transfers that are consistent with observed network structures given agents' static profits and discount factors. This is in contrast to previous empirical work in bilateral contracting which has required assuming some notion of bargaining power (e.g., one side makes take-it-or-leave-it offers) or requiring knowledge of contracted transfers in order to estimate parameters of interest, as static Nash conditions utilized therein do not allow these factors to affect outside options and disagreement points (e.g., Ho (2009), Crawford and Yurukoglu (2010)).[4]

We explore the implications of our model for applied analysis through a simulated example. We study a stylized bilateral contracting game between hospitals and health maintenance organizations (HMOs); this environment and the determinants of negotiated transfers between medical providers

[2]There also exist extensive form games where all agents may re-contract immediately without delay following any breakdown in negotiations (e.g., Stole and Zweibel (1996)).

[3]Lee (2010) computes an equilibrium of a dynamic network formation game without transfers in which links can be made unilaterally but not dropped; Crawford and Yurukoglu (2010) computes the equilibrium reallocation of a contracting game in a static environment.

[4]This is not strictly a function of dynamics, per se; utilizing 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.

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). In our example, period payoffs are generated from both an underlying demand system for insurers and providers as well as a pricing game among firms, and we compute MPE networks and transfers for a large number of simulated markets.

Results indicate that 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. We provide Monte-Carlo evidence illustrating that agents' bargaining powers can be partially identified off of observed equilibrium network structures for even a small numbers of markets. Furthermore, 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 yield different predicted equilibrium transfers than a static model. Finally, we use our model to study the impact of hypothetical hospital mergers on negotiated transfers and profits, crucially controlling for post-merger changes in network structure and focusing solely on induced changes in bargaining power. 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. 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 a higher cost, less efficient hospitals.

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/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 literature the estimation of dynamic games – particularly Hotz and Miller (1993), Hotz, Miller, Sanders, and Smith (1994) and Aguirregabiria and Mira (2007) – for the computation and estimation of our model.

Our paper contributes to the general literatures on network formation and contracting with externalities. Several papers focus on static environments and study the existence of stable and efficient networks (Jackson and Wolinsky, 1996; Kranton and Minehart, 2001; Prat and Rustichini, 2003; Bloch and Jackson, 2007), vertical contracting when there is when 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), and surplus division (de Fontenay and Gans, 2007). Dynamic papers include Manea (2009) and Abreu and Manea (2010a,b) which analyze surplus division over unit surplus in a repeated game on networks (but restricts how the network can change), Gomes (2005) which studies dynamic coalition formation games, and Dutta, Ghosal, and Ray (2005) which provides a model of dynamic network formation without transfers, and provides conditions under which the efficient network will be realized. Our

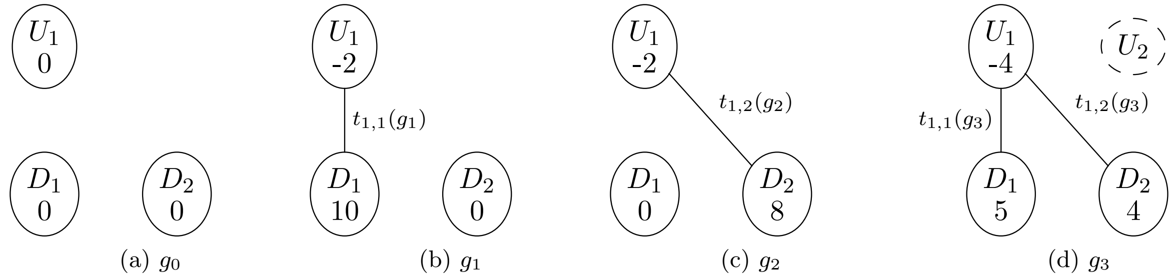

Figure 1: Potential Networks g_0, g_1, g_2, g_3 between firms U_1, D_1, D_2. Period payoffs contained within circles; $t_{ij}(g_k)$ represents payment between U_i and D_j under network g_k.

use of bilateral Nash bargaining with endogenous disagreement points and random proposers also relates to stochastic bargaining models (Binmore (1987); Merlo and Wilson (1995); c.f. Muthoo (1999)).

Finally, our paper is in the spirit of 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.

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 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; 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 $\pi_{U_1}(\mathbf{t}(g_3)) = -4 + t_{1,1}(g_3)+t_{1,2}(g_3)$, $\pi_{D_1}(\mathbf{t}(g_3)) = 5-t_{1,1}(g_3)$ 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 (1986). 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; furthermore, 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 bilateral Nash bargaining in each period: e.g., in network g_3, the negotiated transfer $t_{1,2}(g_3)$ will maximize the Nash product of U_1 and D_2's gains from trade *given* U_1 *and* D_1*'s trade* $t_{1,1}(g_3)$:

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

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

Regardless of the particular bargaining protocol used, there is still a need to specify each agent's outside option ($\widetilde{\pi_k}$). Some static models have assumed that all other agreements are binding given disagreement and do not change (e.g., Cremer and Riordan (1987); Horn and Wolinsky (1988)). 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: $\widetilde{\pi_{U_1}}(D_2; g_3) = -2 + t_{1,1}(g_3)$. Note that what U_1 receives upon disagreement with D_2 from D_1 is what would have been negotiated under g_3. This is similar to Nash equilibrium reasoning in that agents do not account for future adjustments by other agents. 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 $t_{1,1}(g_3) = 3.5$ and $t_{1,2}(g_3) = 3$.

Another approach, as utilized in Stole and Zweibel (1996), relaxes the assumption that transfers do not change upon disagreement, and allows all other agents to immediately renegotiate while preventing any disagreeing pair from recontracting again. 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 $t_{1,2}(g_2) = 5$ 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 negotiating $t_{1,1}(g_3)$.

Our paper builds on this latter approach with two important extensions: (i) agents can form new links in the future (including with agents with whom they previously disagreed), and (ii) adjustments to the network may not be immediate and there may be some delay, including after disagreement. We feel that these are particularly important features of any model which attempts to capture network formation and bargaining in real world, repeated settings. We thus specify dynamic link formation game whereby each period only one (randomly selected) agent can propose new links; 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:

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

where each agent k's outside option has been replaced by $\beta_k V_k(g_0)$, which represents the expected (discounted) *continuation value* to agent k of moving to network g_0 (which occurs under disagreement), 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 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 3.2, we return to this example and illustrate how results are substantively different once we account for dynamics.

3 Model

We study an infinite horizon, discrete time, dynamic network formation game 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*. Let $\mathbf{G} \subset \{0,1\}^{n\times(n-1)}$ represent the set of all "feasible" networks: it may include all possible networks, or only 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; we denote this by $ij \in g$. Let $N_i(g)$ denote the set of agents i is connected to in network structure 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.[5] Every agent i discounts future payoffs at constant rate $\beta_i \in (0, 1)$.

Associated with any given network structure g are *per-period payoffs* $\pi(g, \eta, \mathbf{t}_g) \equiv \{\pi_i(g, \eta, \mathbf{t}_g)\}_{i\in N}$, where η represents payoff shocks and $\mathbf{t}_g \equiv \{t_{ij;g}\}_{ij\in g}$ represents negotiated payments between all linked agents in each period referred to as *per-period contracts*.[6] We denote the space of feasible per-period contracts as τ, which is determined by the particular application being modeled; e.g., $\tau \equiv \mathcal{R}$ if contracts can only be simple linear fees, or $\tau \equiv \mathcal{R}^n$ for n-part tariffs.

We will assume per-period payoffs are primitives of the analysis 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. In addition, we assume that the payoffs $\pi_i(g, \cdot)$ are continuous in per-period contracts $\mathbf{t}_g$ for any given network g. [7]

Timing. Assume at the beginning of period k the current network structure is g^k. We first provide a brief overview of the timing for each period before detailing specifics:

1. Agents engage in bilateral Nash bargaining to determine period contracts and period payoffs:

   (a) First, each pair $ij \in g^k$ observes a publicly observable pair-specific shock $\eta_{ij} \in \eta$, where η is drawn from continuous density f^η and η_{ij} represents the additional period profit i and j realize if they reach an agreement and maintain their link in the current period. We assume η_{ij} enters profits additively.

   (b) Any pair $ij \in g^k$ 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 $\tilde{g}$ then

[5] 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.

[6] E.g., in a buyer-seller setting, period payoffs may arise from a price competition game among downstream buyers for consumers given wholesale prices $\mathbf{t}_g$. Note these payoffs also contain any per-period costs of maintaining links to potential trading partners.

[7] In applied work, these payoff functions $\pi(\cdot)$ (even for states never observed) can be recovered from separate analysis (c.f. Ackerberg, Benkard, Berry, and Pakes (2007)); e.g., in a manufacturer-retailer setting, they can be obtained from a structural model of consumer demand for products and retailer pricing.

has links that are unstable, those dissolve as well. This repeats until a stable network $g \subseteq \tilde{g} \subseteq g^k$ (i.e., every pair $ij \in g$ has gains from trade) is reached.

   (c) Per-period contracts $\mathbf{t}_g$ are determined via bilateral Nash bargaining, and each agent i obtains per-period profits $\pi_i(g, \eta, \mathbf{t}_g)$.

2. Subsequent to the distribution of per-period payoffs, at most one agent is randomly chosen to be a "proposer" according to the state-dependent probability distribution $\lambda : \mathbf{G} \to \Delta^n$. Denote this current proposer by i, where $i \in N \cup \{0\}$ and $i = 0$ means that no proposer was chosen.

   (a) If $i \neq 0$ is chosen to be a proposer, i can "propose" to move from the current network structure g to a new network $g' \in \chi_i(g)$, where $\chi_i(g)$ denotes i's set of *reachable networks* from network g. We assume $\chi_i(g) \equiv \{g' : g'_{-i} = g_{-i}\}$ – i.e., $\chi_i(g)$ contains only networks that differ from g by which agents i is linked to.[8]

   (b) When an agent is selected to be a proposer, he observes a vector of random proposer-payoff shocks $\epsilon_i \equiv \{\epsilon_{i,1}, \ldots, \epsilon_{i,|\mathbf{G}|}\}$ drawn independently from a continuous density $f_i^{\epsilon}(\epsilon_i)$ with finite first moments. Agent i receives additional payoff $\epsilon_{i,g'}$ if he proposes network g' and it is accepted.[9]

   (c) In proposing network g', i offers or demands *lump-sum transfers* $T^{\sigma}(g'|g) \equiv \{T^{\sigma}_{i,j}(g'|g)\}$ to or from each agent $j \in N_i(g')$ if j accepts link ij, where $T^{\sigma}_{i,j}(\cdot)$ will depend on agent strategies σ; both lump-sum transfers and strategies will be defined later. As penalties (potentially zero) for breaking links, agent i also may be required to pay lump-sum transfers $T^{\sigma}_{i,k}(g'|g)$ to all agents k such that $k \in N_i(g)$ and $k \notin N_i(g')$.

3. All agents $j \in N_i(g')$ simultaneously choose whether or not to accept lump-sum transfers; if j accepts, link ij is formed (or maintained). This results in the creation of a new network g'', where $g''_{-i} = g'_{-i}$, and $ij \in g''$ if and only if j has accepted i's offer. Transfers $\{T^{\sigma}_{i,j}(g'|g)\}$ are exchanged between i and all $j \in N_i(g'')$, and i realizes additional profits $\epsilon_{i,g'}$. The starting network structure for the next period $k+1$ will be $g^{k+1} = g''$. Note that this means a proposer can unilaterally cut links (subject to a potential penalty), but requires agreement from the other party to form a new link.

Strategies and Value Functions We first focus on the strategies faced by agents when chosen as a proposer or when proposed to, which we call being a "proposee."[10] We restrict attention to Markov strategies denoted by $\sigma = \{\sigma_i(g, \epsilon_i), \tilde{\sigma}_i(g'|g)\}$, where $\sigma_i : \mathbf{G} \times \mathbb{R}^{|\mathbf{G}|} \to \mathbf{G}$ are proposer strategies, $\tilde{\sigma}_i : \mathbf{G} \times \mathbf{G} \times N \to \{0, 1\}$ are proposee strategies, and the space of all strategies σ is given by $\mathbf{\Sigma}$. In other words, proposers only condition on the current network state and their draw

[8] The definition of $\chi_i(g)$ can vary depending on the application; we provide other examples of $\chi_i(g)$ later.
[9] $|\mathbf{G}|$ represents the dimensionality of $\mathbf{G}$.
[10] I.e., if i proposes g', all $j \in N_i(g')$ are proposees.

of payoff shocks when determining their actions, whereas proposees condition on the current and proposed network state as well as who is proposing.

For exposition, we will focus on cases where $\tilde{\sigma}_j(g'|g;i) = 1 \;\forall\; i, j, g, g'$, and will in the next subsection provide conditions on the underlying lump-sum transfers such that it will always be an equilibrium in the subgame for all proposees to always accept any proposed network.

Following related literature on dynamic discrete games (c.f., Hotz and Miller (1993), Hotz, Miller, Sanders, and Smith (1994)), given a vector of strategies σ, we define $P_i^\sigma(g'|g)$ to be the *conditional choice probabilities* of state g' from state g given that agent $i \in N$ is chosen as a proposer:

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

Given the current network structure is g and i is the current proposer who utilizes strategy σ_i, $P_i^\sigma(g'|g)$ is the probability an opponent that does not observe ϵ_i assigns to the event that the next period network will be g'. For convenience, if no one is selected as a proposer such that $i = 0$, we define $P_0^\sigma(g'|g) = 1$ if $g' = g$ and $P_0^\sigma(g'|g) = 0$ otherwise.

Let $V_i^\sigma(g)$ represent agent i's expected value function at the beginning of a period before any agent is selected as a proposer and before any payoff shocks ϵ and η are realized:

$$V_i^\sigma(g) = E_\eta\big[\pi_i(g', \eta, \mathbf{t}_{g'}^\sigma) + \hat{V}_i^\sigma(g') : g' = \Gamma_g(\eta, \hat{V}^\sigma, \mathbf{t})\big], \tag{5}$$

The value function takes expectations over the per-period contracting shocks η, and $g' = \Gamma_g(\cdot)$ 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. $\hat{V}^\sigma \equiv \{\hat{V}_i^\sigma(g)\}$ represents the set of value functions after unstable links have dissolved and per-period payoffs and contracts have been realized, and for each agent i is given by:

$$\begin{aligned}\hat{V}_i^\sigma(g) = &\sum_{j \in N\cup\{0\}, j\neq i} \lambda_j(g) \sum_{g' \in \chi_j(g)} P_j^\sigma(g'|g)\big(T_{j,i}^\sigma(g'|g) + \beta_i V_i^\sigma(g')\big) \\ &+ \lambda_i(g) \int \Big[\max_{g' \in \chi_i(g)} \Big(\epsilon_{i,g'} - \big(\sum_{j\neq i} T_{i,j}^\sigma(g'|g)\big) + \beta_i V_i^\sigma(g')\Big)\Big] f_i(\epsilon_i) d\epsilon_i.\end{aligned} \tag{6}$$

The first line in (6) represents i's expected future profits if agent $j \neq i$ is selected to be a proposer: j chooses to propose network g' with the transition probabilities given by $P_j^\sigma(g'|\cdot)$, which provides i with potential lump sum transfers and the future value function from transitioning to state g'. The second line, by Bellman's principle of optimality, is agent i's portion of the value function when i is selected as proposer. Again given our restriction on proposee strategies, we are implicitly assuming that an agent i can induce any network $g' \in \chi_i(g)$ through the appropriate transfers. Although all value functions in this section are functions of period contracts $\mathbf{t}$, in the next subsection we define $\mathbf{t}$ endogenously as a function of strategies and value functions.

Note that for a given vector of π_i, the right hand side of (5) is a contraction mapping (c.f. Aguirregabiria and Mira (2002)), and thus there is a unique V_i^σ which solves (5) for any given σ.

Lump-sum Transfers. As mentioned, we assume that lump-sum transfers are pre-defined for our game, and $T^{\cdot}_{i,j}(\cdot) : \mathbf{G} \times \mathbf{G} \times \Sigma \to \mathbb{R}$. We will assume that $\forall i \neq j$ and $g, g' \in \mathbf{G} \times \mathbf{G}$:

$$T^{\sigma}_{i,j}(g'|g) = \begin{cases} \beta_j \big(V^{\sigma}_j(g' - ij) - V^{\sigma}_j(g') \big) & \text{if } j \in N_i(g'), \\ c_{i,j}(g'|g) \geq 0 & \text{if } j \in N_i(g) \text{ and } j \notin N_i(g'), \\ 0 & \text{otherwise,} \end{cases} \tag{7}$$

where $T^{\sigma}_{i,i}(g'|g) = -\sum_{j \neq i} T^{\sigma}_{i,j}(g'|g)$, and $g' - ij$ denotes the network realized by subtracting link ij from g'.[11]

What this means is that i either pays j or demands payment from j depending on whether or not j is worse or better off accepting an offer from i to form a link *given j believes all other agents $j' \in N_i(g')$ will accept*. The amount offered or demanded makes j exactly indifferent between accepting and rejecting.[12] Consequently, it will be a subgame Nash equilibrium for all proposees to accept: i.e., $\tilde{\sigma}_j(g'|g; i) = 1$; in the case that there are multiple equilibria in the subgame, we assume that this is the particular one that is selected. Thus, given our definition of transfers, i essentially *chooses* a new network whenever he proposes g'. Depending on the application, however, one can restrict i from offering negative transfers (i.e., demanding a lump-sum payment in order to accept) so that proposers can only make payments to proposees.

The second line in the definition of $T^{\sigma}_{i,j}$ allows for the possibility that i may need to pay any j whom he previously was linked to in g but will no longer be linked to in g' a *termination fee* $c_{i,j}(g'|g) \geq 0$ to break that link; i.e., although links can be broken unilaterally, it might only be possible to do so with some penalty. Depending on the application, this fee can be interpreted as a breach of contract penalty (e.g. liquidated damages, legal costs, damages), the cost of redirecting downstream consumers, or of dissolving a merger. Note that this termination fee can depend on the specific pair i, j as well as the specific networks g and g'. It can be exogenously given, or depending on the application, endogenous as well: e.g., one specification that captures expected damages would be $c_{i,j}(g'|g) = \beta_j \big(V^{\sigma}_j(g' \cup ij) - V^{\sigma}_j(g') \big)$.

Although we define the transfer functions to be exogenous for expositional convenience, we note that the offered transfers for forming a link can arise endogenously for an appropriately defined subgame.[13]

[11] Formally, transfers will only depend on strategy profiles σ in prescribing play for all agents *in the future*; i.e., what agent i does at time t does not enter into the determination of transfers at time t. This is important, since even though σ_i may place 0 probability on g' given the network is g, it will be the case that $T^{\sigma}_{i,j}(g'|g)$ is well defined.

[12] Note any proposee $j \in N_i(g')$ will always accept i's proposal as long as $T^{\sigma}_{i,j}(g'|g) + \beta V^{\sigma}_j(g') \geq \beta V^{\sigma}_j(g'')$, where g'' is what j believes next-period's network would be if j rejected.

[13] E.g., consider the following timing: a proposer publicly announces both the proposed network g' and the set of lump-sum transfers to each proposee, and then all proposees simultaneously decide to accept or reject the offered link. Lump-sum transfers are only paid out if the realized network is g'. In any equilibrium, it can be shown that i will not be able to reach g' from g any more cheaply than by proposing transfers in (7) (assuming that there exists a means of selecting among multiple potential proposee responses for any feasible menu of transfers).

Per-period Contracts 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 which arises from g as *unstable* links are broken.

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

$$t_{ij;g}(\eta) \in \arg\max_{\tilde{t}} \Bigg[\underbrace{\big[\pi_i(g, \eta, \{\tilde{t}, \mathbf{t}^\sigma_{-ij;g}\}) + \hat{V}^\sigma_i(g)\big]}_{S^\sigma_{i,j}(g;\eta,\mathbf{t}^\sigma_g)} - \underbrace{\big[\pi_i(g', \eta, \mathbf{t}^\sigma_{g'}) + \hat{V}^\sigma_i(g')\big]}_{S^\sigma_{i,j}(g-ij;\eta,\mathbf{t}^\sigma_{g'})} \Bigg]^{b_{ij}}$$
$$\times \Bigg[\underbrace{\big[\pi_j(g, \eta, \{\tilde{t}, \mathbf{t}^\sigma_{-ij;g}\}) + \hat{V}^\sigma_j(g)\big]}_{S^\sigma_{j,i}(g;\eta,\mathbf{t}^\sigma_g)} - \underbrace{\big[\pi_j(g', \eta, \mathbf{t}^\sigma_{g'}) + \hat{V}^\sigma_j(g')\big]}_{S^\sigma_{j,i}(g-ij;\eta,\mathbf{t}^\sigma_{g'})} \Bigg]^{b_{ji}} \tag{8}$$

where $g' = \Gamma_{g-ij}(\cdot)$, $\mathbf{t}^\sigma_{-ij;g} \equiv \{\mathbf{t}^\sigma_g \backslash t^\sigma_{ij;g}\}$, 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 $\Delta S^\sigma_{i,j}(g;\eta,\mathbf{t}^\sigma) = S^\sigma_{i,j}(g;\eta,\mathbf{t}^\sigma_g) - S^\sigma_{i,j}(g - ij;\eta,\mathbf{t}^\sigma_{g'})$ and $\Delta S^\sigma_{j,i}(g;\eta,\mathbf{t}^\sigma) = S^\sigma_{j,i}(g;\eta,\mathbf{t}^\sigma_g) - S^\sigma_{j,i}(g-ij;\eta,\mathbf{t}^\sigma_{g'})$), given the contracts of all other linked pairs of agents and the strategies of agents employed in the larger network formation game. This is a variant of the static bilateral bargaining equilibria between an upstream supplier and downstream firms used in Cremer and Riordan (1987) and Horn and Wolinsky (1988), and later adapted by Crawford and Yurukoglu (2010) in applied work to model negotiations between multiple upstream content providers and multiple downstream multichannel video distributors.[14]

Importantly, in a significant departure from these papers and other 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 assume all other pairs immediately renegotiate. Rather, the disagreement point in our model is simply the continuation value from moving to a new network structure in which i and j do not contract. Indeed, our notion of a disagreement point is internally consistent with the larger dynamic network formation game that we propose: when two agents fail to come to an agreement, we assume that they simply move to a new network structure in which they are not linked, but anticipate subsequent changes to the network may occur – furthermore, i and j anticipate that they may even contract again in the future.

There may, however, be instances in which for a certain pairs of linked agents $\{ij \in g\}$, *draws of η, perceived continuation values V and per-period contracts in alternative network structures*, there is no value of $\mathbf{t}_g$ in which both $\Delta S^\sigma_{i,j}(g;\eta,\mathbf{t}^\sigma)$ and $\Delta S^\sigma_{j,i}(g;\eta,\mathbf{t}^\sigma)$ are positive; i.e., there are no

[14] One common motivation for this setup is that each firm sends different agents to bargain simultaneously with each of its linked partners; bargains happen simultaneously, and agents from the same firm do not coordinate with one another. We concede that this particular concept may be unsatisfying as it admits sets of contracts which are not robust to a multi-contract deviation by a single firm; although it is possible to provide more stringent conditions, solving for such contracts is difficult if not infeasible to implement in practice. Nonetheless, we emphasize that (8) still admits contracts which satisfy such stronger conditions.

potential gains from trade. In this case, the Nash bargaining solution is undefined. One reason this may occur in equilibrium is due to the presence of general externalities (i.e., no restrictions on π): 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 previously exhibited gains from trade, to no longer be profitable to maintain.[15]

Consequently, given a set of contracts for other network states $\mathbf{t}^\sigma_{-g} \equiv \{\mathbf{t}^\sigma \setminus \mathbf{t}^\sigma_g\}$, we assume that any network g for which there are no gains from trade between at least one pair of agents $ij \in g$ for *any* set of period-contracts $\mathbf{t}^\sigma_g$ is *unstable*. Similarly, 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 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 $\Gamma_g(\eta, V, \mathbf{t})$ and is defined recursively as:

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

where $\tilde{\mathbf{t}} \equiv \{\tilde{\mathbf{t}}_g, \mathbf{t}_{-g}\}$. Note that $\Gamma_g(\eta, V, \mathbf{t})$ may also be the empty network.

Markov Perfect Network Formation Game We can parametrize our model by the tuple $(N, \mathbf{G}, \chi, \pi, \lambda, \mathbf{b}, \beta, \mathbf{f}, \mathbf{c})$, where $\chi = \{\chi_i\}$, $\pi = \{\pi_i\}$, $\mathbf{b} = \{b_{ij}\}$, $\beta = \{\beta_i\}$, $\mathbf{f} = \{f^\eta, \{f^\epsilon_i\}\}$, $\mathbf{c} = \{c_{i,j}\}$, and $\mathbf{G}$ and λ are as defined before.

3.1 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:

$$\sigma^*_i(g, \epsilon_i) = \arg \max_{g' \in \chi_i(g)} \left[\epsilon_{i,g'} + T^{\sigma^*}_{i,i}(g'|g) + \beta_i V^{\sigma^*}_i(g') \right].^{16} \tag{10}$$

where lump-sum transfers $T^{\sigma^*}_{i,j}(g'|g)$ are given by (7), and given V^{σ^*} for any η, period contracts $\mathbf{t}^{\sigma^*}_{\mathbf{g}}$ satisfy (8) for all stable g and Γ_g 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

[15]This is why when defining the disagreement point in (8), agents anticipate moving to $g' \equiv \Gamma_{g-ij}(\cdot)$ as opposed to simply $g - ij$.

[16]As discussed before, we choose the equilibrium in which proposees always accept.

best response probability function function $\Lambda(P) = \{\Lambda_i(g'|g; P_{-i})\}$, where

$$\Lambda_i(g'|g; P_{-i}) = \int I(g' = \arg\max_{g'' \in \chi_i(g)} \epsilon_{i,g''} + T^P_{i,i}(g''|g) + \beta_i V^P_i(g''))f^\epsilon_i(\epsilon_i)d\epsilon_i \tag{11}$$

where T^P and V^P are transfers and value functions defined for a given vector 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 well-defined and continuous in the compact space P; the existence of at least one fixed point will then follow from Browser's theorem. We proceed by showing V^P is continuous in P, or equivalently that V^σ is continuous in σ. From equations (5) and (6), we see that the continuity of V^σ is guaranteed if the expression $E_\eta\big[\pi_i(g, \eta, \mathbf{t}^\sigma_g)\big]$ is continuous in σ for a fixed network g. This in turn will follow given the density function f^η is assumed to be continuous and from the following assumption:

Assumption 3.1. Fix η, σ, a stable network g, and a matrix of continuation values $\{\hat{V}^\sigma_i(\cdot)\}$. There exists a unique set of per period contracts $\mathbf{t}^\sigma_g(\eta)$ that solves (8). In addition, each contract in $\mathbf{t}^\sigma_g(\eta)$ is continuous in σ for a fixed η.

Finally, since f^ϵ_i is assumed to be continuous and the continuity of V^P implies continuity of T^P (see (7)), the existence claim follows.

Whether or not the assumption used to guarantee existence 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.

3.2 Example

At this point, we return to the example depicted in Figure 1. If every agent shares a common discount factor $\beta = 0$, equal Nash bargaining parameters, $\lambda_U = .5$, $\lambda_{D_1} = \lambda_{D_2} = .25$, and shocks ϵ and η are sufficiently small, the only MPNE involves g_1 and g_3 being visited in equilibrium with equal probability. Equilibrium strategies for U would be to mix evenly between proposing g_1 and g_3; D_1 would always propose g_1 or g_3 (depending on which is reachable from the current network state); and D_2 would propose g_2 under networks g_0 and g_2, and otherwise mix evenly between g_1 and g_3. Transfers would be the same as in the static setting under the Stole and Zweibel (1996) assumption that upon disagreement, all other contracts are immediately renegotiated: $t_{1,1}(g_1) = 6$, $t_{1,2}(g_2) = 5$, and $t_{1,1}(g_3) = t_{1,2}(g_3) = 4$.

If $\beta = .9$, however, g_1 would be the only network visited in equilibrium. Importantly, g_3 is not a *stable network*, and hence is never reached in equilibrium. U would always propose g_1; D_1 would always propose g_1 if reachable, and g_2 otherwise; and D_2 would propose g_2 if reachable, and g_1 otherwise. Equilibrium transfers would be $t_{1,1}(g_1) \approx 6.2$ and $t_{1,2}(g_2) \approx 5.15$. These are strictly higher than if $\beta = 0$: since U_1 is able to recontract with either D_2 or D_1 upon reaching g_0, it has a higher outside option when bargaining with either D_1 under g_1 or D_2 under g_2, and hence can extract higher rents. Thus, a dynamic model ($\beta > 0$) once outside options are endogenized leads to substantively different predictions regarding equilibrium network and surplus division.

3.3 Computation

Although finite per-period payoff shocks η are necessary for 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 transition 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:

$$V_i^{P^*}(g) = \pi_i(\tilde{g}, \mathbf{t}^{P^*}) + \sum_{j \in N \cup \{0\}} \lambda_j(\tilde{g}) \sum_{g' \in \chi_j(\tilde{g})} \Big[P_j^*(g'|\tilde{g}) \big(T_{j,i}^{P^*}(g'|\tilde{g}) + \beta_i V_i^{P^*}(g') + e_{j,i}^{P^*}(g', \tilde{g}) \big) \Big], \quad (12)$$

where $\tilde{g} = \Gamma_g^{P^*}$, and $e_{j,i}^{P^*}(g', \tilde{g}) = 0$ if $j \neq i$, and

$$e_{i,i}^{P^*}(g', \tilde{g}) = E[\epsilon_i | \sigma_i^*(\tilde{g}, \epsilon_i) = g'] = \int \epsilon_i I\{\sigma_i^*(\tilde{g}, \epsilon_i) = g'\} f_i(\epsilon_i) d\epsilon_i, \quad (13)$$

which is the expected choice-specific payoff shock to agent i when he is chosen as a proposer, the network is $\tilde{g}$ and he proposes g'. Even though σ^* is referenced in (13), $e_{i,i}^{P^*}$ is a function only of P^* and *choice specific value functions* $v_i^{P^*}(\cdot|g) \equiv T_{i,i}^{P^*}(\cdot|g) + \beta_i V_i^{P^*}(\cdot)$ (c.f. Hotz and Miller (1993)).[17] Consequently for any equilibrium, the computation of value functions in (12) can be obtained via matrix algebra and rewritten in matrix notation as:

$$\mathbf{V}_i^{P^*} = \left(\mathbf{I} - \beta_i \sum_j [\lambda_j \otimes \mathbf{1}_{|\mathbf{G}|}^T] * \mathbf{P}_j^* \right)^{-1} \left(\pi_i + \sum_j \lambda_j * \sum_{g'} \mathbf{P}_j^*(g') * [\mathbf{T}_{j,i}^{P^*}(g') + \mathbf{e}_{i,j}^{P^*}(g')] \right) \quad (14)$$

where $\mathbf{I}$ is the identity matrix, $\mathbf{1}_{|\mathbf{G}|}$ a vector of ones of length $|\mathbf{G}|$, $\mathbf{P}_j^{\sigma^*}$ is the matrix of j's perceived transition probabilities; $\mathbf{V}_i^{\sigma^*}$, λ_j, $\mathbf{P}_j^{\sigma^*}(g')$, π_i $\mathbf{T}_{j,i}^{\sigma^*}(g')$, and $\mathbf{e}_{i,j}^{\sigma^*}(g')$ are vectors across all network states g (given equilibrium strategies σ^*, where applicable); and $\otimes$ is the Kronecker product and $*$ denotes the Hadamard, or element-by-element product. Importantly, values at any state $g \neq \Gamma_g^{P^*}$ which are unstable are replaced with those at state $\Gamma_g^{P^*}$.

Define $\Upsilon(P) \equiv \{\Upsilon_i(g, P) : g \in \mathbf{G}\}$ to be the solution to (14) for an arbitrary set of probabilities P; i.e., $\Upsilon(P)$ is the vector of value functions for all agents which is consistent with the transition probabilities given by P: $\Upsilon_i(g, P) = V_i^{P^*}(g)$. Another interpretation is $\Upsilon_i(g, P)$ represents firm i's expected value function of being in network state g given i perceives all firms (including i itself) to behave according to conditional choice probabilities P. We can follow the same procedure in Aguirregabiria and Mira (2007) to show that any fixed point of the mapping $\Psi(P) \equiv \{\Psi_i(g'|g; P)\}$,

[17] Certain distributions of ϵ, including type I extreme value and normal, yield closed form expressions for $e_{i,i}^{P^*}(\cdot)$.

where:

$$\Psi_i(g'|g;P) = \int I(g' = \arg\max_{g''\in\chi_i(g)} \epsilon_{i,g''} + T^P_{i,i}(g''|g) + \beta_i \Upsilon_i(g'';P)) f_i(\epsilon_i) d\epsilon_i$$

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, transfers $\mathbf{T}^{P^0}$, per-period transfers $\mathbf{t}^{P^0}$, value functions $\mathbf{V}^{P^0}$, and at each iteration n:

1. Obtain updated implied transition probabilities $\mathbf{P}^n(g')$ from strategies σ^{n-1}, given by (4);
2. For each agent i, update $V_i^{P^n} = \Upsilon_i(P^n)$ using (14), where modified policy iteration can be utilized to approximate the matrix inversion (c.f. Judd (1998));
3. For each pair of agents i, j across all pairs of states, update lump-sum transfers $\mathbf{T}^{P^n}_{i,j}$ using $V_i^{P^n}$ and (7);
4. Update Γ^n_g for each network state g based on values of $\mathbf{t}^{n-1}$ and $V_i^{P^n}$ and (9);[18]
5. Update per-period contracts $\mathbf{t}^{P^n}$ using (8);
6. Update agents' optimal strategies given $V_i^{P^n}$ and $\mathbf{T}^{P^n}_{i,j}$ to obtain σ^{n+1}.

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

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. Further aiding computation is the fact that the dimensionality of $\chi_i(\cdot)$ grows much slower than $|\mathbf{G}|$. For larger games, we are exploring the use of other approaches including reinforcement learning algorithms (c.f. Pakes and McGuire (2001)).

3.4 Extensions to the Model

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

[18] If a network g is determined to be unstable, we reset $\mathbf{t}_g = 0$.

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: the probability that j is ever selected to be a proposer becomes 0, and all future profits that would otherwise accrue to j are captured by i. Additionally, i and j will have the same set of links upon merging: i.e., for any network g where $ij \in g$, $N_i(g) = N_j(g)$.

We maintain the timing and structure of the main model, but assume when an agent is selected as a proposer, he may propose at most *one* merger in addition to offering new links or dropping existing links as before. E.g., if i is the proposer, who for the sake of discussion is a buyer, his set of reachable networks $\chi_i(g)$ includes the possibility of linking with other sellers as well as linking with at most one additional buyer. Instead of offering a lump-sum or per-period contracts, a proposer who offers a merger negotiates the purchase price T_{ij}^M which satisfies the following Nash bargaining problem:

$$T_{ij}^M(g) \in \arg\max_{\tilde{T}} \left[\hat{V}_i^\sigma(g) - \tilde{T} - \hat{V}_i^\sigma(g-ij)\right]^{b_{ij}} \times \left[\hat{V}_j^\sigma(g) + \tilde{T} - \hat{V}_j^\sigma(g-ij)\right]^{b_{ji}} \tag{15}$$

Consequently, a merger is only feasible if for the network g that is proposed, $\hat{V}_i^\sigma(g) + \hat{V}_j^\sigma(g) \geq \hat{V}_i^\sigma(g-ij) + \hat{V}_i^\sigma(g-ij)$.

Assumptions on how λ, bargaining parameters, or profit functions change upon merging will depend on the application.

Exclusive Contracts In certain applications, an agent may be able to propose an *exclusive contract* in which he demands another agent contract only with him. We can extend our model to incorporate this possibility by expanding an agent's set of reachable networks when he is chosen as a proposer. Let $\chi_i^e(g)$ denote agent i's set of reachable networks from g when exclusive contracts are permitted, where:

$$\chi_i^e(g) \equiv \chi_i(g) \cup \left\{g' : \forall j \in N_i(g'),\ (g'_{-i} \cap g'_{-j} = g_{-i} \cap g_{-j})\ \wedge\ N_j(g') = \{i\}\right\}.$$

In words, agent i can propose to form a link with any set of agents, and in addition, he has the option of proposing to any of those agents that she drop all her other links, i.e. form an exclusive contract. For any proposed network $g' \in \chi_i^e(g)$, let $k_i(g')$ denote the set of agents to whom i offers an exclusive contract.

To accommodate exclusivity, we also modify the transfers that accompany proposals. Specifically, we let $T_{i,j}^{e,\sigma}(g'|g)$ denote the new transfers and define it as follows:

$$T_{i,j}^{e,\sigma}(g'|g) = \begin{cases} \max\{\beta_j\big(V_j^\sigma(g') - V_j^\sigma(g'')\big), 0\} & \text{if } j \in N_i(g'), \\ c_{i,j}(g,g') & \text{if } j \in N_i(g) \text{ and } j \notin N_i(g'), \\ 0 & \text{otherwise,} \end{cases}$$

where $g'' = (g' \cup g_j) - ij$. As before, $T_{i,i}^{e,\sigma}(g'|g) = -\sum_{j\neq i} T_{i,j}^{e,\sigma}(g'|g)$.

If an agent j is offered an exclusive contract, she has two options: she can either *accept*, thereby forming link ij, receiving transfer $T_{i,j}^{e,\sigma}$ and dropping all of her other links, or she can *reject* so that ij is not formed, no transfer is made, and she keeps her existing links.

Notice that our model does not support imposing specific penalties to agents for breaking an exclusive contract. The reason is that in any given state, 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.

Strategic Lump-sum Transfers In the appendix, we present a variant of our model in which lump-sum transfer offers are within each agents' strategy space, but links that can be negotiated arise randomly. We show that the equilibrium lump-sum transfers proposed will be identical to those prescribed before.

3.5 Estimation of Nash Bargaining Parameters

Assume there are $m = 1 \dots M$ markets, each with primitives $(N^m, \mathbf{G}^m, \chi^m, \pi^m, \beta^m, \mathbf{f}^m, \mathbf{c}^m, \lambda^{\mathbf{m}})$ which are either observed, assumed, or can be separately estimated. 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.[19]

The econometrician observes a sequence network structures $\{g^{m,t}\}$ for markets $m = 1 \dots M$ and $t = 1 \dots T$. There are two potential cases: (i) the econometrician observes all changes to the network structure and, although the econometrician does not observe who a "proposer" is, she knows (or can make an assumption about) when a network does not change as a result of an agent actively choosing to keep the same network (i.e., whenever $g^{m,t} = g^{m,t-1}$); or (ii) the econometrician only knows that the sequence of networks are separated by at least one possible network change, but the sequence can contain gaps.

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

$$Q_M^1(\theta, P) = \sum_{m=1}^{M} \sum_{t=2}^{T} \ln \Phi(g^{m,t}|g^{m,t-1}; \theta, P, \mathbf{b}(\theta, z^{m,t})), \tag{16}$$

where $\Phi(g^{m,t}|g^{m,t-1}; \theta, P) = \sum_{j\in N^m} \lambda_j(g^{m,t-1})\Psi_j(g'|g; P, \mathbf{b}(\theta, z^{m,t}))$. On the other hand, for case

[19]Ongoing research involves determining whether both λ and $\mathbf{b}$ can both be separately identified.

(ii) where we are not certain the sequence of observed network structures does not contain gaps, we can instead rely on the following *pseudo-likelihood*:

$$Q_M^2(\theta, P) = \sum_{m=1}^{M} \sum_{t=2}^{T} \ln \tilde{\Phi}(g^{m,t}; \theta, P, \mathbf{b}(\theta, z^{m,t})), \tag{17}$$

where $\tilde{\Phi}(\cdot)$ is the steady-state equilibrium ergodic distribution over *all* networks given $\{\Phi(g^{m,t}|g^{m,t-1}; \cdot)\}$.

Since there is the possibility of multiple equilibria, we define the MLE as:

$$\hat{\theta}_{MLE} = \arg\max_{\theta} \big[\sup_{P \in (0,1)^{N \times |\mathbf{G}|}} Q_M^i(\theta, P) \qquad \text{subject to} \qquad P = \Psi(\theta, P) \big], \tag{18}$$

where $Q_M^i(\theta, P)$ is either given by (16) or (17) 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; nonetheless, it is an extremely strong assumption, and it remains to be seen how feasible this estimator is. In section 4.3, we perform a Monte Carlo exercise testing the performance of this estimator on a simulated data set, and rely on the assumption that our estimation routine can select out the same equilibrium as the data generating process.

4 Application: Insurer-Provider Negotiations

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 equlbria; 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 examine the impact of hospital mergers on industry profits, negotiated transfers, premiums, and consumer welfare.

4.1 Stage Game Timing

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

1. HMOs compete Nash-in-prices and choose premiums to charge consumers;
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.

We assume: hospitals must serve any patient which 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 which consumers may choose. Further details, including distributional assumptions for firm and consumer characteristics and specification of profit functions, are provided in the appendix.

4.2 Equilibria: Network Structure and Transfers

We first simulate multiple markets of HMOs and hospitals with random draws from firm and consumer characteristics; we assume that hospitals are the only ones who may be proposers (i.e., $\lambda_i = 0$ if i is not a hospital), motivated by the conversations with negotiators which indicate that hospitals are typically the ones that open negotiations with insurance providers. Nonetheless, although hospitals are the ones assumed to bring up negotiations over new links, we only allow lump-sum transfers to be positive, and vary the Nash bargaining parameters so that the division of surplus between agents may differ. Thus, in this section, we use "bargaining power" to refer to an agent's Nash bargaining parameter. 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 4 hospitals and 2 HMO's, there are $2^8 = 256$ potential networks, yet the average number of networks that occur frequently in equilibrium is approximately 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 is relatively low for most of the 1x2 markets as there is an incentive for a hospital to be exclusive to one HMO: reducing downstream HMO

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	2.35	0.00	0.02	0.09	0.64	0.00	0.02	1.00	0.86
2 HMOs	Hospitals	2.18	0.00	0.01	0.15	0.76	0.00	0.01	0.99	0.89
	HMOs	2.83	0.01	0.02	0.01	0.28	0.01	0.01	1.00	0.80
2 Hosp	Equal	2.38	0.77	0.44	0.14	0.60	0.44	0.23	1.96	3.23
2 HMOs	Hospitals	2.03	0.82	0.48	0.37	0.72	0.61	0.31	2.00	3.47
	HMOs	2.94	0.69	0.56	0.04	0.15	0.04	0.11	1.96	2.68
3 Hosp	Equal	2.10	0.45	0.14	0.03	0.27	0.10	0.03	2.92	4.55
2 HMOs	Hospitals	1.95	0.69	0.20	0.12	0.58	0.46	0.10	3.00	5.16
	HMOs	1.92	0.05	0.25	0.00	0.02	0.00	0.01	2.88	3.38
4 Hosp	Equal	-	-	-	-	-	-	-	-	-
2 HMOs	Hospitals	-	-	-	-	-	-	-	-	-
	HMOs	-	-	-	-	-	-	-	-	-

Summary statistics from 100 market draws for each specification (250 draws for 2x2). *"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.

price competition allows for greater industry rents to be extracted, and hence split among industry players. However, once an additional hospital is added to the market, this probability rises dramatically. Nonetheless, a full network is still not always reached once there are multiple hospitals: one contributing factor is there is an incentive to not include high cost hospitals (particularly insofar they lead to higher negotiated prices) since HMOs cannot influence which hospital its own patients visit. The efficient network 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, more than half the markets have a network that is visited more than 50% of the time when hospitals have greater bargaining power. 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. Except for cases where hospitals have higher bargaining power, markets visit networks that exclude at least one hospital from contracting; on average, less than one hospital is expected to be excluded. Finally, the last column indicates the expected number of links that are sustained in equilibrium.

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.	7.32	1.57	6.03	3.79	4.73	0.39	4.76	1.47	-2.59	3.91	4.37	0.32
Av. Cost	-0.41	0.05	-1.04	0.08	-0.03	0.02	-0.50	0.05	-1.02	0.07	-0.05	0.01
Cost-AC	-0.27	0.06	-0.06	0.10	-0.14	0.02	-0.18	0.05	-0.11	0.09	-0.13	0.02
# Patient	0.48	0.09	0.32	0.21	0.15	0.03	0.08	0.10	-0.24	0.23	0.06	0.02
Total # Patients	-0.41	0.05	-0.71	0.11	-0.08	0.01	-0.21	0.05	-0.42	0.12	-0.03	0.01
HMO Marg	1.67	0.78	7.32	1.47	-1.27	0.23	3.24	0.73	10.35	1.51	-1.01	0.19
Extra Hospital	-1.26	0.14	-1.50	0.28	-0.39	0.05	-1.10	0.13	-1.29	0.25	-0.29	0.04
R^2	0.30		0.31		0.20		0.31		0.34		0.28	

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 - $b_i j = .8$ when i is a hospital, .2 otherwise; HMOs - $b_i j = .8$ 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 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 (2009) 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 or 3 hospitals. 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 a higher per-patient transfer (on average) when they have higher bargaining power, and lower when HMOs have higher bargaining power. Both being in a market with lower average costs and having a lower than average costs positively impacts predicted hospital margins. Furthermore, consistent with previous findings and anecdotal evidence on quantity discounts, the number of total patients served by a hospital reduces margins. However, a dynamic model also predicts across bargaining specifications that the more patients a hospital serves for a particular HMO (holding fixed total patients served), the higher the margins the hospital can demand from that HMO; this measure proxies for the HMO having a worse outside option upon losing that hospital. Generally, higher HMO margins are also associated with higher negotiated hospital margins unless HMOs have stronger bargaining power. Finally, having an additional hospital in the market negatively impacts hospital margins.

The static specification generally shares similar signs as the dynamic model, with some differences in magnitudes. However, across markets, a static model tends to overestimate implied hospital transfers: under equal bargaining power, they are 12% higher; they are 6% or 31% higher on average if hospitals or HMOs had greater bargaining power. 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.[20] This suggests that the impact of certain market characteristics influences bargaining only when agents anticipate future changes to the network structure.

4.3 Estimation, Monte Carlo

One important takeaway from the simulated equilibrium network distributions is that the allocation of bargaining power through the choice of Nash bargaining parameters has a significant impact on which networks are reached in equilibrium, and hence the determinants of equilibrium transfers. We explore how this variation in observed network structure can be used to identify and estimate Nash bargaining parameters used to generate the data, which would then allow us to compute transfers between agents if they were unobserved.

We initially focus on 2x2 markets where we observe a sequence of 10 networks $\{g^{m,1}, \ldots, g^{m,10}\}$ for each market m, where the initial network drawn from the equilibrium steady state distribution, and subsequent networks are drawn according to the equilibrium transition probabilities. We assume Nash bargaining parameters are constant across markets in each sample, and are parameterized by $\theta = b_H \in [0, 1]$, where we assume $b_{ij} = b_H$ 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. We choose the value of b_H which maximizes the probability of observing the sequence of networks in the data; i.e., the MLE given by (18) 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 are implicitly relying on the latter assumption given the estimation routine utilizes the same computational algorithm as the data generating process. However, 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

[20]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 Market / Sample	5 Markets / Sample	10 Markets / Sample
Avg. Estimate:	0.50	0.50	0.48	0.48
95% C.I.:		[0.15,0.95]	[0.20,0.90]	[0.25,0.60]
Avg. Estimate:	0.80	0.55	0.75	0.79
95% C.I.:		[0.05,0.95]	[0.30,0.95]	[0.60,0.95]
Avg. Estimate:	0.20	0.43	0.28	0.22
95% C.I.:		[0.05,0.95]	[0.10,0.60]	[0.10,0.40]

Estimated values of hospital bargaining power b_H for 40 samples of either 1, 10, or 25 markets in 2x2 settings where a sequence of 10 networks were observed. Grid search conducted over b_H in increments of .05.

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 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 .25 of the true value.

4.4 Hospital Mergers [Preliminary]

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 measuring these potential costs.

In this section, we apply our framework to counterfactual merger analysis by simulating new 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 "hospital system" per market.[21] 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

[21] Voluntary mergers are being examined in ongoing analysis.

Table 4: Merger Simulations

	"B-Pow"	$+\Delta\pi^H$	$-\Delta\pi^H_{5\%}$	$+\Delta\pi^M$	$-\Delta\pi^M_{5\%}$	$+p^M$	$-p^M$	+Patients	$-$Patients$_{5\%}$
(i) All	Equal	0.88	0.08	0.19	0.62	0.79	0.04	0.15	0.77
Markets	Hospitals	0.77	0.12	0.14	0.59	0.79	0.10	0.16	0.44
	HMOs	0.92	0.07	0.44	0.11	0.69	0.06	0.17	0.38
(ii) All,	Equal	-	-	0.11	0.68	0.86	0.02	0.09	0.84
$+\Delta\pi^H \geq 0$	Hospitals	-	-	0.09	0.63	0.81	0.09	0.15	0.47
	HMOs	-	-	0.46	0.08	0.73	0.06	0.14	0.38
(iii) Full	Equal	0.99	0.00	0.01	0.95	0.97	0.00	0.01	0.95
Markets	Hospitals	0.93	0.03	0.02	0.80	0.93	0.02	0.02	0.54
	HMOs	1.00	0.00	0.00	0.29	0.86	0.00	0.00	0.14
(iv) Full,	Equal	-	-	0.00	0.95	0.98	0.00	0.00	0.96
$+\Delta\pi^H \geq 0$	Hospitals	-	-	0.01	0.81	0.94	0.01	0.01	0.54
	HMOs	-	-	0.00	0.29	0.86	0.00	0.00	0.14

Summary statistics from merger simulations, where (i) and (ii) contains all markets in sample, (iii) and (iv) condition on only markets in which the full network originally occurred at least 50% of the time before a merger, and (ii) and (iv) condition also on markets where hospitals find it profitable to merge. *"B-Pow"*: Equal - $b_{ij} = .5 \; \forall \; ij$; Hospitals - $b_ij = .8$ when i is a hospital, .2 otherwise; HMOs - $b_ij = .8$ 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%; $+p^M, -p^M$: percentage of markets in which both HMO premiums increase or fall at all; $+Patients, -Patients(5\%)$: percentage of markets in which total patients insured increases at all or falls by 5%.

the bargaining, network, and pricing outcomes without taking a stance on potential cost savings, quality improvements, or other efficiencies.

Results from merger simulations across 250 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; and total patients insured in the market increases at all or falls by 5%. We again assume $\beta = .9$. Row (i) examines mergers across all markets as bargaining power varies, (ii) conditions only on those markets in which hospital expected profits increase subsequent to a merger, (iii) focuses only on markets in which the full network occurred at least 50% of the time in the equilibrium network distribution pre-merger, and (iv) modifies (iii) by again conditioning only on markets where hospitals find a merger to be profitable.

Generally, across all markets, hospital profits tend to increase when hospitals merge – this is consistent with a merger strengthening hospitals' and/or weakening HMOs' disagreement points. However, what is surprising is that hospitals may occasionally be worse off, and if given the choice, may prefer not to merge. For the most part, however, 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.[22] We can verify this explanation by focusing only on markets where the pre-merger equilibrium network distribution visited the full network in at least 50% of periods (row (iii)), and noting the percentage of mergers which are preferred by hospitals increases dramatically; indeed, if we focus only on markets where the full network occurred 75% or more both pre- and post-merger (not shown), all simulated mergers would lead to higher hospital profits. Hence, in situations where pre-merger all hospitals contracted with all HMOs, under equal bargaining power hospitals unilaterally prefer to merge as the risk of greater post-merger utilization of the higher cost hospital is no longer an issue.

Interestingly, when hospitals have greater bargaining power, it is more likely that hospital profits fall post-merger, even when the market was previously likely to be in the full network state. To help explain this result, it is useful to note that a merger, by bundling hospitals together under one system, reduces the ability of HMOs to differentiate themselves: pre-merger, one HMO might have had two hospitals and the other only one, but post-merger, generally both HMOs each have both hospitals. This leads to fiercer price competition in the premium setting game and reduces total industry profits, leaving less surplus to be split amongst the firms. When hospitals have greater bargaining power, they can extract most of the rents from HMOs and hence generally prefer increasing total industry profits, which occasionally favors the unmerged scenario.

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. Not only can this come about due to fiercer price competition among HMOs, but hospitals, upon merging, also partly internalize the impact of increasing their own per-patient rates on total patient flows. 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, this does not occur often for mergers which occur in markets with already full networks, and in general, mergers seem to predominantly lead to lower HMO profits, higher premiums, and fewer patients insured.

5 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-

[22]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, 2010).

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. Finally, 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.

A Alternative Model

We present an alternative formulation of our model in which lump-sum transfer functions are no longer exogenous but rather part of each agent's strategy space. To avoid the issues raised in section 3, we assume in this setup that links are negotiated one at a time, and the link is chosen randomly.

We modify the timing of the main model such that after a proposer i is chosen at step (2), an agent $j \in N_i(\mathbf{G}) \cup \{0\}$ is chosen randomly by nature according to the distribution $\lambda^{N,i}(\cdot) : \mathbf{G} \to \Delta(N_i(\mathbf{G}) \cup \{0\})$, where $N_i(\mathbf{G})$ denotes the set of all agents j that i could feasibly link with under any network structure; i.e., $N_i(\mathbf{G}) = \{j : \exists g \in \mathbf{G} \text{ s.t. } ij \in g\}$. Note ij can either be an existing link under g, or one that hasn't been formed. After a proposer i and an agent j that i can link with has been selected, i makes a take-it-or-leave it offer to j in the form of a lumpsum transfer. If j accepts, the link is formed (or maintained, if $ij \in g$); otherwise, ij do not link.

Strategies for a proposer will be denoted by $T = \{T_i(g, j)\}$, where $T_i : \mathbf{G} \times N_i(\mathbf{G}) \to \{\emptyset\} \cup \mathbb{R}^+$ specifies the transfer that i offers to j to form link ij, and $\{\emptyset\}$ denotes the null contract. We redefine an propoer i's value function conditional on network state g, potential proposee j, and strategy profile T as:

$$\hat{V}_i^T(g; i, j) = \max_{T_i} \left(\pi_i(g) - \mathbf{1}_{(T_i, g, j)} T_i + \beta_i V_i^T(g' + \mathbf{1}_{(T_i, g, j)} \cdot ij) \right)$$

where $g' = g$ if $ij \notin g$, and $g' = g - ij$ otherwise; $\mathbf{1}_{(T_i,g,j)}$ is the indicator for whether j accepts offer T_i to form a link; and

$$V_i^T(g) = \sum_{j \in N \cup \{0\}} \lambda_j(g) \sum_{k \in N_j(\mathbf{G})} \lambda_k^{N,j}(g) \hat{V}_i^T(g; j, k).$$

Since $\mathbf{1}(T_i, g, j) = 1$ if and only if $T_i \geq V_j^T(g' + ij) - V_j^T(g')$, it is straightforward to show that one set of equilibrium transfers will be:

$$T_i(g, j) = \begin{cases} -\beta_j \left(V_j^T(g' + ij) - V_j^T(g') \right) & \text{if } V_i^T(g' + ij) + V_j^T(g' + ij) \geq V_i^T(g') + V_j^T(g') \\ \emptyset & \text{otherwise,} \end{cases} \tag{19}$$

The reason that there may be multiple equilibria is that whenever i offers $\emptyset$, he can also offer any transfer that j will refuse. However, whenever ij do link, payments in any equilibrium will be given by the top line of (19).

Note that this game will be equivalent to the setting in which a proposer could choose whether or not to form a link with j, with transfers in the case of forming a link being prescribed in (19). This latter case is almost identical to the game described in our main text, with the restriction on $\chi_i(g)$ being limited to one link deviations, and the caveat that a proposer only makes payments to one agent in each period.

B HMO Hospital Application

B.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., $\mathbf{\Gamma} \equiv \{0, 1\}^{M \times H}$. 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 $j \in N_k(g)$ given the current network structure is g. Each HMO j charges a one-time premium p_j. 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

$$u_{i,k}^H = \alpha_r^H \theta_k^H + \omega_{i,k} \tag{20}$$

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 $N_j(g)$:[23]

$$u^M_{i,k}(g) = \alpha^M_r \theta^M_k - \alpha^P p_k + \gamma \left(ln(\sum_{h \in N_k(g)} exp(\alpha^H_r \theta^H_h)) \right) + \varepsilon_{i,k} \tag{21}$$

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:

$$\sigma^M_j(g) = \sum_{r \in R} \sigma^R_r \frac{\exp(\alpha^M_r \theta^M_j - \alpha p_j + \gamma \ln(\sum_{h \in N_j(g)} \exp(\alpha^H_r \theta^H_h)))}{1 + \sum_{m \in M}(\exp(\alpha^M_r \theta^M_l - p_l + \gamma \ln(\sum_{h \in N_m(g)} \exp(\alpha^H_r \theta^H_h))))} = \sum_{r \in R} \sigma^R_r \tilde{\sigma}^M_{j,r}(g)$$

We use $\tilde{\sigma}^M_{j,r}(g)$ 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:

$$\tilde{\sigma}^R_{r,j}(g) = \frac{\sigma^R_r \tilde{\sigma}^M_{j,r}(g)}{\sum_{s \in R} \sigma^R_s \tilde{\sigma}^M_{j,s}(g)}$$

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:

$$\sigma^H_{k,j}(g) = \gamma \sum_{r \in R} \tilde{\sigma}^R_{r,j}(g) \frac{\exp(\alpha^H_r \theta^H_k)}{\sum_{h \in N_j(g)} \exp(\alpha^H_r \theta^H_h)}$$

Note that although σ^M_j is a function of the entire network structure g and premiums charged by all HMOs, $\sigma^H_{k,j}$ 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^M_j and σ^H_k 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

$$\pi^H_k(g) = (t_{jk} - c^H_k)(\sum_{j \in N_h(g)} N \sigma^M_j \sigma^H_{k,j}) + \sum_{j \in N_k(g)} \eta_{k,j}$$

c^H_k represents the average cost of serving each patient at hospital k.

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

$$\pi^M_j(g) = N \sigma^M_j (p_j - c^M_j - \sum_{k \in N_j(g)} \sigma^H_{k,j} t_{j,k}) + \sum_{k \in N_j(g)} \eta_{k,j} \qquad \text{if } N_j(g) \neq \{\}$$

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

1. Each HMO j chooses a premium $p_j \in \mathbb{R}^+$ that it will charge each consumer that chooses to join its plan.
2. Each individual i chooses to enroll in an HMO plan, with the utility from choosing HMO j being $u^M_{i,j}$ defined in (21), 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 (20).
4. HMO payoffs (π^M) and hospital payoffs (π^H) are realized.

B.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 $\alpha^M_r = \alpha^H_r = \alpha^P = 1$ for all agents.

[23]Note $E_\omega(\max_{h \in N_j(g)}(u^H_{i,h})) = ln(\sum_{k \in N_j(g)} \exp(\alpha^H_r \theta^H_k))$

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 $\mu_{\theta^H} = 25$ 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 $10\pi^2/6$. We assume away profit shocks η; i.e., $\eta_{i,j} = 0$. We set discount factors $\beta_i = .9 \; \forall \; i$, and set termination fees $c_{i,j} = 0$.

References

Abreu, D., and M. Manea (2010a): "Bargaining and Efficiency in Networks," mimeo.

——— (2010b): "Markov Perfect Equilibria in a Model of Bargaining in Networks," mimeo.

Ackerberg, D., C. L. Benkard, S. Berry, and A. Pakes (2007): "Econometric Tools for Analyzing Market Outcomes," in *Handbook of Econometrics*, ed. by J. J. Heckman, and E. E. Leamer, vol. 6a, chap. 63. North-Holland, Amsterdam.

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

——— (2007): "Sequential Estimation of Dynamic Discrete Games," *Econometrica*, 75(1), 1–53.

Bernheim, B. D., and M. D. Whinston (1986): "Menu Auctions, Resource Allocation, and Economic Influence," *Quarterly Journal of Economics*, 51(1), 1–32.

Binmore, K. (1987): "Perfect Equilibria in Bargaining Models," in *The Economics of Bargaining*, ed. by K. Binmore, and P. Dasgupta. Basil Blackwell, Oxford.

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

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

Crawford, G. S., and A. Yurukoglu (2010): "The Welfare Effects of Bundling in Multichannel Television," mimeo.

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

de Fontenay, C. C., and J. S. Gans (2007): "Bilateral Bargaining with Externalities," mimeo.

Doraszelski, U., and A. Pakes (2007): "A Framework for Applied Dynamic Analysis in IO," in *Handbook of Industrial Organization*, ed. by M. Armstrong, and R. Porter, vol. 3, pp. 1887–1966. North-Holland, Amsterdam.

Dranove, D., M. Satterthwaite, and A. Sfekas (2008): "Boundedly Rational Bargaining in Option Demand Markets: An Empirical Application," mimeo.

Dutta, B., S. Ghosal, and D. Ray (2005): "Farsighted Network Formation," *Journal of Economic Theory*, 122, 143–164.

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

Gomes, A. (2005): "Multilateral Contracting with Externalities," *Econometrica*, 73(4), 1329–1350.

Gowrisankaran, G. (1999): "A Dynamic Model of Endogenous Horizontal Mergers," *RAND Journal of Economics*, 30, 56–83.

Gowrisankaran, G., and T. J. Holmes (2004): "Mergers and the Evolution of Industry Concentration: Results From the Dominant Firm Model," *RAND Journal of Economics*, 35, 561–582.

Hart, O., and J. Tirole (1990): "Vertical Integration and Market Foreclosure," *Brookings Papers on Economic Activity, Microeconomics*, 1990, 205–286.

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

Ho, K., and A. Pakes (2010): "Physician Responses to Financial Incentives: Evidence from Hospital Discharge Data," mimeo.

Horn, H., and A. Wolinsky (1988): "Bilateral Monopolies and Incentives for Merger," *RAND Journal of Economics*, 19(3), 408–419.

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

Hotz, V. J., R. A. Miller, S. Sanders, and J. Smith (1994): "A Simulation Estimator for Dynamic Models of Discrete Choice," *Review of Economic Studies*, 61(3), 265–289.

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

Judd, K. L. (1998): *Numerical Methods in Economics*. MIT Press, Cambridge, MA.

Kranton, R. E., and D. F. Minehart (2001): "A Theory of Buyer-Seller Networks," *American Economic Review*, 91(3), 485–508.

Lee, R. S. (2010): "Vertical Integration and Exclusivity in Platform and Two-Sided Markets," mimeo, http://pages.stern.nyu.edu/ rslee/papers/VIExclusivity.pdf.

Manea, M. (2009): "Bargaining in Stationary Networks," mimeo.

Maskin, E., and J. Tirole (1988): "A Theory of Dynamic Oligopoly: I & II," *Econometrica*, 56(3), 549–600.

McAfee, R. P., and M. Schwartz (1994): "Opportunism in Multilateral Vertical Contracting: Nondiscrimination, Exclusivity, and Uniformity," *American Economic Review*, 84(1), 210–230.

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

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

Muthoo, A. (1999): *Bargaining Theory with Applications*. Cambridge.

Pakes, A. (2009): “Alternative Models for Moment Inequalities,” mimeo.

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

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

Prat, A., and A. Rustichini (2003): “Games Played Through Agents,” *Econometrica*, 71(4), 989–1026.

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

Segal, I., and M. D. Whinston (2003): “Robust Predictions for Bilateral Contracting with Externalities,” *Econometrica*, 71(3), 757–791.

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

Stole, L. A., and J. Zweibel (1996): “Intra-Firm Bargaining under Non-Binding Contracts,” *Review of Economic Studies*, 63(3), 375–410.

Town, R. J., and G. Vistnes (2001): “Hospital Competition in HMO Networks,” *Journal of Health Economics*, 20, 733–753.