MODELS OF BILATERAL TRADE IN NETWORKS

MIHAI MANEA

Department of Economics, MIT, manea@mit.edu

Abstract. This chapter reviews the research on bilateral trade in markets with a network structure. The survey focuses on the following questions: how do local prices depend on global network architecture? When does the law of one price hold? How do competitive forces in different segments of the market influence the bargaining power of each position in the network? Is trade efficient? How does the network evolve as traders reach agreements and exit the market? If forming links is costly for the traders, what networks will emerge? How does the underlying market mechanism affect payoffs and allocations?

1. Introduction

The study of prices and allocations in markets constitutes a central part of economics. The classical paradigm of general equilibrium presumes large economies in which no individual trader has market power. Under this approach, goods are homogeneous and infinitely divisible, and the law of one price holds—that is, identical goods have the same price. Moreover, traders take prices as given. The market clearing conditions implicitly assume that all buyers can trade freely with all sellers without delays or transaction costs. The actual trading procedure and the price formation mechanism are, however, left unmodeled. The theory predicts that efficient allocations emerge in such frictionless competitive economies.

General equilibrium theory, as such, is silent on what the underlying trading process is and how exactly prices emerge. Stahl (1972) and Rubinstein (1982) applied the tools of noncooperative game theory to model explicitly the bargaining process that determines prices. Their pioneering noncooperative bargaining models consider bilateral monopoly situations in which both parties have market power. Building on the noncooperative bargaining approach, Rubinstein and Wolinsky (1985, 1990), Gale (1986a,b, 1987), and McLennan and Sonnenschein (1991) later developed the analysis of trade via bilateral exchanges in dynamic markets with a large number of traders. The primary motivation of this early work on bargaining in markets was to provide noncooperative foundations for general equilibrium theory. The

---

Date: March 8, 2016.

The Oxford Handbook of the Economics of Networks (Eds: Y. Bramoulle, A. Galeotti, and B. Rogers), Oxford University Press. I thank Ben Golub, Dilip Abreu, and the editors for comments as well as Aubrey Clark and Will Rafey for editing and proofreading. Section 5 benefited from discussions with Matt Elliott.

research identifies conditions on the composition of the economy and the matching process that ensure convergence to a competitive equilibrium as bargaining frictions vanish.

However, the assumptions that deliver noncooperative foundations for general equilibrium theory do not reflect the realities of many economic activities. Price-taking behavior is unrealistic in markets that involve only a small number of traders. Social and business relationships, geography, and technological compatibility determine which pairs of traders can engage in exchange. Some trades may entail prohibitive transaction costs. Certain goods and services are tailored for specific segments of the market. Trade is dynamic and market conditions change over time. These departures from the foundations of general equilibrium theory give rise to complex market interactions. Local demand and supply influence the market power of every trader, yet trading activity in remote areas of the economy may have significant spillover effects. Hence, the bargaining power of every trader depends on both the local topology of trading opportunities and the global market architecture. The varying nature of competition in different parts of the economy leads to systematic deviations from the law of one price. Furthermore, decentralized bargaining may generate inefficient outcomes because local incentives for trade are not necessarily aligned with global efficiency.

The asymmetries among both traders and goods described above naturally call for a network formulation. A link between a pair of traders indicates that they can trade (or form a partnership) with each other. Many questions emerge: how do local prices depend on the global network structure? When does the law of one price hold? How do competitive forces in different segments of the market determine the bargaining power of every position in the network? Are trading outcomes efficient? How does the network evolve over time as traders reach agreements and exit the market? What networks are expected to form if traders have to invest in their links? How does the underlying market mechanism affect payoffs and allocations? These questions constitute the focus of an active area of research in network economics. In this chapter we survey some important contributions to the literature, focusing mainly on noncooperative models of bilateral bargaining in networks.

Section 2 defines some key network concepts and reviews two classic results from graph theory, Hall's Marriage Theorem and the Gallai-Edmonds Decomposition Theorem, which are needed for the analysis. In Section 3 we discuss bargaining models in which every pair of traders that reaches an agreement is removed from the network without replacement. Thus the pool of bargaining partners for each trader active in the market shrinks over time. In this setting, traders need to anticipate how the network of trading opportunities will evolve and how their future bargaining position will improve or deteriorate when agreements are forged in different parts of the network. The initial work of Corominas-Bosch (2004) and Polanski (2007) in this area assumes that linked traders are matched in pairs according to a centralized mechanism that maximizes the total surplus that can be achieved from exchange.

As a consequence, these models predict that all agreements occur simultaneously, at the beginning of the game, and that equilibrium outcomes are efficient. Both Corominas-Bosch and Polanski find a close relationship between equilibrium payoffs and the Gallai-Edmonds decomposition. In contrast to the research of Corominas-Bosch and Polanski, Abreu and Manea (2012a,b) study markets with decentralized matching in which a single pair of linked traders bargains at a time. Decentralized bargaining leads to richer market dynamics: not all matches result in agreement, a pair of linked players may refuse to trade at some stage yet agree to trade at a later one, and multiple equilibria may coexist. Moreover, decentralization creates incentives for inefficient trade when we restrict attention to Markov perfect equilibria. A construction of a subgame perfect equilibrium developed by Abreu and Manea (2012b) demonstrates that efficiency may nonetheless be attained in every network using a complex design of punishments and rewards. We compare the balance of bargaining power across the models covered in Section 3 and reflect on sources of divergence.

Section 4 presents the stationary bargaining model of Manea (2011a). In that model, players who reach agreements exit the market and new players fill the positions left vacant, so that the network of trading opportunities does not change over time. A partition of the network into oligopolies emerges endogenously in equilibrium. The oligopolies with the lowest seller-to-buyer ratio have the greatest market power and drive market outcomes. The tractability of the model permits a detailed analysis of both price dispersion in the network and strategic network formation. We conclude the section with Nguyen's (2014) alternative characterization of equilibrium payoffs based on convex programing.

Section 5 reviews the celebrated assignment game of Shapley and Shubik (1971) along with its network applications. After summarizing the key results of Shapley and Shubik, we describe two price formation mechanisms—one proposed by Crawford and Knoer (1981), the other by Demange, Gale, and Sotomayor (1986)—both of which converge to the buyer-optimal core allocation in the assignment game. We then turn to a seminal contribution to the literature on buyer-seller networks due to Kranton and Minehart (2001). In that setting, buyers have private information about their valuations and sellers run simultaneous price-ascending auctions for their linked buyers à la Demange, Gale, and Sotomayor. For every network, the buyer-optimal core outcome of the associated assignment game emerges in equilibrium. Furthermore, a network that maximizes total expected welfare forms if buyers make linking investments before their values are realized. We briefly discuss work by Elliott (2015) showing that the latter conclusion relies critically on the assumption that one side of the market incurs all linking costs. Lastly, we analyze the noncooperative model of Elliott and Nava (2014), which introduces bargaining frictions in the assignment game. In contrast to the other work explored in Section 5, Elliott and Nava find that trade is typically inefficient. Section 6 provides concluding remarks and suggests directions for future research.

2. Framework

A finite set N of *players* interacts in a market. In examples, we assume that $N = \{1, 2, \ldots, n\}$ for some positive integer n. Each player can forge at most one agreement with some other player. Agreements may entail forming a partnership or trading a single unit of an indivisible good. The set of feasible bilateral agreements is described by a network. Formally, a *network* formed by the set of players (or *nodes*) N is a collection of *links* $G \subseteq \{(i,j) \mid i \neq j \in N\}$ with $(i,j) \in G$ if and only if $(j,i) \in G$. To avoid double counting in summations and set cardinalities, we identify the pairs (i,j) and (j,i) and use the shorthand ij for the corresponding undirected link between i and j. When $ij \in G$, we say that i is *linked* to (or is a *neighbor* of) j in G. A node $i \in N$ is *isolated* in G if it is not linked to any node in G. For any network G and subset of nodes $M \subseteq N$, let $\mathcal{L}^G(M)$ denote the set of all neighbors in G of nodes from M, that is, $\mathcal{L}^G(M) = \{i \mid \exists j \in M, ij \in G\}$. A set $M \subseteq N$ is called *G-independent* if no pair of nodes in M is linked in G, i.e., $\mathcal{L}^G(M) \cap M = \emptyset$.

In the models of Sections 3 and 4, the existence of a link ij in G indicates that players i and j can generate a *unit* of surplus by reaching an agreement with each other. The assumption of symmetric link values allows us to focus exclusively on how the structure of the network G affects the bargaining power of every node in G. For some of the models we discuss, players are partitioned into *buyers* and *sellers*, with the corresponding sets denoted by B and S ($B \cup S = N$ and $B \cap S = \emptyset$). A network G is *bipartite* with the partition (B, S) if every link in G connects a buyer to a seller or, using the notation above, $\mathcal{L}^G(B) \subseteq S$ and $\mathcal{L}^G(S) \subseteq B$. The underlying assumption is that buyers have unit demand and sellers have unit supply.

A network G' is a *subnetwork* of G if it consists of a subset of the links in G, i.e., $G' \subseteq G$. The subnetwork G' of G *covers* a set of nodes $M \subseteq N$ if every node in M has at least one link in G'. The *subnetwork* of G *induced* by a subset of nodes M is the network $\{ij \in G \mid i, j \in M\}$. Players i and j are (path) *connected* in a network G if there exists a sequence $i_0 = i, i_1, \ldots, i_{\bar{k}} = j$ such that $i_k i_{k+1} \in G$ for $k = \overline{0, \bar{k} - 1}$. Connectedness is an equivalence relation; its equivalence classes are called *connected components* of G. Thus all players in a given connected component are connected by a path in the network, and there are no links between distinct components.

A *matching* in a network G is a subnetwork H of G in which every node has at most one link (i.e., $H \subseteq G$ and there do not exist three distinct nodes i, j, k in G such that both links ij and ik belong to H). If $ij \in H$ then we say that i and j are *matched* (with each other) under H. A matching of G is *perfect* if it covers all the nodes in G. Matchings of G that contain the largest number of links are called *maximum matchings*. The cardinality of the maximum matchings of G defines the *maximum total surplus* in G. For example, in

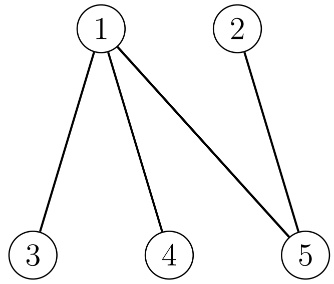

FIGURE 1. A network example

the 5-player network from Figure 1, there are two maximum matchings, $\{(1,3),(2,5)\}$ and $\{(1,4),(2,5)\}$, each generating a maximum total surplus of 2.

A simple condition due to Hall (1935) characterizes the bipartite networks which admit perfect matchings. More generally, Hall's theorem provides a necessary and sufficient condition for such networks to contain matchings that cover all the buyers (or sellers). The condition requires that every group of buyers be collectively linked to a set of sellers that has at least the same cardinality.

Theorem 1 (Hall's Marriage Theorem [26]). *Suppose that G is a bipartite network with the partition (B, S). Then there exists a matching of G that covers B if and only if*

$$|\mathcal{L}^G(M)| \geq |M|, \forall M \subseteq B. \tag{1}$$

The network G has a perfect matching if and only if it has an equal number of buyers and sellers ($|B| = |S|$) and satisfies condition (1).

Maximum matchings determine the total surplus that can be generated when bilateral trades are organized efficiently. The following characterization of maximum matchings, independently discovered by Gallai (1964) and Edmonds (1965), proves useful in the welfare analysis of the models we consider. The result relies on the following partition of the set of nodes N in a network G. The set of *under-demanded* nodes in G, denoted by U, consists of all nodes i with the property that there exists at least one maximum matching of G in which i is not matched. The set of *over-demanded* nodes in G, denoted O, consists of nodes that do not belong to U and are linked to at least one node in U. The set of *perfectly matched* nodes in G is formed by the remaining nodes, $P = N \setminus (O \cup U)$. For an illustration, in the network from Figure 1, nodes 3 and 4 are under-demanded, 1 is over-demanded, and 2 and 5 are perfectly matched.

Theorem 2 (Gallai-Edmonds Decomposition [26]). *Fix a network G with the sets of perfectly matched, over-demanded, and under-demanded nodes denoted by $P, O,$ and U, respectively.*

(1) The following statements hold for every maximum matching H of G:
- *every node in O is matched under H to a node in U;*

- *every element of P is matched under H to another element of P;*
- *every connected component of the subnetwork induced by U in G has an odd number of nodes, and in each such component all nodes except at most one are covered by H.*

(2) The network obtained by removing any single node from every connected component of the subnetwork induced by U in G admits a perfect matching.

(3) If G is bipartite, then U is G-independent.[1]

A few special networks come up in our analysis and examples. In a *line* network, the set of nodes is ordered linearly and every node is linked only to its immediate predecessor and successor in the order (if any); we refer to the first and last nodes in the order as the *endpoints* of the line. A *cycle* is a network obtained by adding a link between the endpoints of a line network. A *star* is a network in which one node, called the *center*, is linked to all the other nodes, the *spokes*, and there are no links between spokes.

3. Non-Stationary Bargaining Models

In this section we discuss several noncooperative bargaining models in which pairs of players linked in a network G forge agreements and exit the market. Every model assumes that players bargain with their neighbors in G at discrete dates $t = 0, 1, \ldots$ over an infinite time horizon and have a common discount factor $\delta \in (0, 1)$. Pairs of players that reach agreements trade and leave the network. The remaining players continue to bargain in the resulting subnetwork. Unless otherwise stated, at every stage, players have perfect information of all past actions (including moves by nature).

Let v_i^δ denote the (expected equilibrium) *payoff* of player $i \in N$ for a given discount factor δ. We say that a payoff profile $(v_i^\delta)_{i\in N}$ (and the equilibrium behavior that generates it) is *efficient* for the underlying network G if $\sum_{i\in N} v_i^\delta$ is equal to the maximum total surplus in G. Efficiency requires that all agreements take place without any delay. In settings where matching is decentralized and agreements cannot occur simultaneously, the following welfare criterion is more suitable. The family of payoffs $((v_i^\delta)_{i\in N})_{\delta\in(0,1)}$ (and the corresponding equilibrium outcomes) is *asymptotically efficient* if $\sum_{i\in N} v_i^\delta$ converges to the maximum total surplus in G as $\delta \to 1$.

3.1. **Bargaining with Public Offers.** Corominas-Bosch (2004) studies a game on a bipartite network in which buyers and sellers alternate in making public offers. In periods

[1]Part (3) of the result is not regularly stated with the decomposition theorem but follows immediately from part (2). For a proof, assume without loss of generality that one connected component C of the subnetwork induced by U in G includes more buyers than sellers. If C consists of more than one buyer, then its connectedness implies that it includes at least one seller i. Then removing seller i from C leads to a set of nodes with a different number of buyers and sellers, which clearly cannot have a perfect matching. This contradicts part (2).

$t = 0, 2, \ldots$ each (remaining) seller simultaneously posts a price in the interval $[0, 1]$ at which he is willing to trade the good, and then every buyer observes the posted prices and simultaneously announces a price level he is willing to accept (without naming a specific trading partner). For every $p \in [0, 1]$, consider the subnetwork induced by the buyers and sellers who expressed willingness to trade at (the exact) price p. A maximum matching H of this subnetwork is selected according to a deterministic procedure, and the buyer-seller pairs matched under H trade at price p and exit the market. Players who have not completed a transaction remain in the game for period $t + 1$. In periods $t = 1, 3, \ldots$ roles are reversed, with buyers posting prices and sellers responding. Players have a common discount factor $\delta \in (0, 1)$. The main finding of Corominas-Bosch (2004) is that this game has an efficient subgame perfect equilibrium for every discount factor, which yields payoffs closely tied to the Gallai-Edmonds decomposition.

Theorem 3 (Corominas-Bosch, 2004)**.** *In any bipartite network, there exists an efficient subgame perfect equilibrium for every discount factor δ in which over- and under-demanded players receive payoffs of 1 and 0, respectively, while perfectly matched sellers and buyers obtain payoffs of $1/(\delta + 1)$ and $\delta/(\delta + 1)$, respectively.*

To prove this result, Corominas-Bosch constructs an equilibrium in which all agreements take place in the first period and the trades carried out form a maximum matching. In the first period of the constructed equilibrium, over- and under-demanded sellers post prices of 1 and 0, respectively. Perfectly matched sellers ask for a price of $1/(\delta + 1)$, which is derived from Rubinstein's (1982) two-player bargaining game with alternating offers. Given these equilibrium offers, buyers who are perfectly matched, over-demanded, and under-demanded accept prices of $1/(\delta + 1), 0$, and 1, respectively. The specification of strategies off the equilibrium path is more elaborate.

We now provide some intuition for the constructed equilibrium. Consider a bipartite network G with buyer set B and seller set S. Let P, O, and U denote the sets of perfectly matched, over-demanded, and under-demanded nodes in G, respectively. Under the strategies (partially) described above, every under-demanded player receives a zero payoff (under-demanded players who do not reach a first-period agreement become isolated). To see why buyer $i \in U$ has no profitable deviations in the first period, suppose that i does not accept a price of 1. First, note that by part (3) of Theorem 2, the under-demanded buyer i is linked only to over-demanded sellers, so i does not have any neighbor who is willing to accept a price different from 1 in the first period under the prescribed strategies. Hence buyer i does not trade in the first period if he deviates from his prescribed strategy. We next argue that i becomes isolated following the deviation. Given the prescribed strategies and i's deviation, the set of sellers who post a price of 1 and buyers who accept price 1 in the first period is $(O \cap S) \cup (U \cap B) \setminus \{i\}$. Let G' denote the subnetwork of G induced by this set of nodes. Since $i \in U$, there exists a maximum matching of G that does not cover $\{i\}$. By the Gallai-Edmonds decomposition theorem, such a matching links every node in $O \cap S$ to one in $(U \cap B) \setminus \{i\}$. It follows that maximum matchings of G' consist of at least $|O \cap S|$ links. Since every link in the bipartite network G includes one seller, no matching of G' may contain more than $|O \cap S|$ links, and every matching with $|O \cap S|$ links must cover $O \cap S$. It follows that every maximum matching of G' consists of exactly $|O \cap S|$ links and covers $O \cap S$. Hence, under the assumed market clearing mechanism, the entire set of over-demanded sellers $O \cap S$ trades and exits the market following i's deviation. Since i can be linked in G only to over-demanded sellers, the deviation leaves i isolated, which makes i indifferent between following the prescribed strategy and deviating from it.

To understand the relationship between the constructed equilibrium and competitive outcomes in a corresponding economy where buyers and sellers can trade freely, fix a bipartite network G in which there are more buyers than sellers. Then demand exceeds supply and general equilibrium theory predicts that all trades take place at a price of 1. The subgame perfect equilibrium identified by Corominas-Bosch converges to a competitive outcome as $\delta \to 1$ if and only if all sellers are over-demanded and all buyers are under-demanded. Combining Theorems 1 and 2, we can prove that this condition is equivalent to $|\mathcal{L}^G(M)| > |M|$ for every subset of sellers M. Therefore, Corominas-Bosch's equilibrium implements a competitive outcome if and only if every group of traders on the short side of the overall economy has collective market power in the sense that they have access to a group of trading partners of larger size.

Charness, Corominas-Bosch, and Frechette (2007) test the predictions of a simplified version of Corominas-Bosch's model in an experimental setting. Although the payoff asymmetries observed in the laboratory are consistent with the direction of relative bargaining power suggested by the theory, the division of surplus in the experiment is not as extreme. We next discuss bargaining models that predict a more moderate balance of bargaining power.

3.2. **Bargaining with Efficient Centralized Matching.** Polanski (2007) considers a model of bargaining in a network with an efficient centralized matching procedure. In every period a maximum matching of the remaining network is drawn with equal probability and for each matched pair either party is selected as the proposer with probability $1/2$. In each match, the proposer makes an offer to his partner specifying a split of the unit surplus generated by their link. If the partner accepts the offer, then the two players receive the agreed shares and are removed from the network. The game proceeds to the next round in the subnetwork induced by the players who have not previously reached an agreement. Players have a common discount factor $\delta \in (0, 1)$. Polanski assumes that at the bargaining stage each player observes only his own partner (if matched) and the offer made by his partner (when the partner acts as the proposer). Players learn all past actions at the beginning of every new period.

Polanski states that this bargaining game has a unique subgame perfect equilibrium, but the proof he provides is incomplete. However, this does not substantially change the message of his paper. For the sake of accuracy, we prove in the Appendix that the game admits a unique Markov perfect equilibrium (MPE).[2] [3] Polanski shows that in bipartite networks the unique MPE payoffs are related to the Gallai-Edmonds decomposition in a similar fashion to the equilibrium identified by Corominas-Bosch (2004). However, the payoffs of over- and under-demanded players are usually less extreme.

Theorem 4 (Polanski, 2007)**.** *For every network and discount factor, the bargaining game has a unique MPE, in which all matches result in agreement at every stage. In bipartite networks, the MPE payoffs belong to the interval* $(0, 1/2)$ *for under-demanded players and to* $(1/2, 1)$ *for over-demanded players, while those of perfectly matched players are equal to* $1/2$*.*

The first part of the result is proven in the Appendix. The proof relies critically on the assumption that a maximum matching is selected for bargaining in the subnetwork remaining at every stage. We will see that the main conclusions of the argument—the uniqueness and the efficiency of the MPE, along with the fact that every match results in an agreement in the MPE—fail to extend to the setting with decentralized matching analyzed by Abreu and Manea (2012a,b), where it is assumed that a single match forms at a time.

For a sketch of the proof for the second part of Theorem 4, fix a bipartite network G and let P, O, and U denote the sets of perfectly matched, over-demanded, and under-demanded nodes in G, respectively. Consider first a player $i \in P$. By Theorem 2, every maximum matching links player i to another player $j \in P$. Suppose that i and j do not reach an agreement when matched to bargain with each other in the first period. All players in $P \setminus \{i, j\} \cup O$ are covered by every maximum matching and, by the first part of the result, must trade in the first period of the MPE. By the construction of the Gallai-Edmonds decomposition, the perfectly matched players i and j do not have any links to under-demanded players. Thus, i and j are left in a subnetwork that contains the single link ij following their disagreement. Then the continuation payoffs of both i and j conditional on failing to reach an agreement with each other in the first period are equal to $\delta/2$. In the unique MPE, the proposer in

[2]In this setting, the natural notion of a Markov state for a player i active in the market at some stage is given by the set of players who have not reached an agreement by that stage, along with the identity of i's match and the offer received by i in the current period. An *MPE* is a sequential equilibrium in which each player i's behavior after any history depends only on i's state induced by that history. This nonstandard adaptation of the definition of MPEs based on private states reflects the nature of imperfect information in Polanski's game.

[3]The uniqueness of subgame perfect and sequential equilibria is an open problem.

the match (i, j) offers a payoff of $\delta/2$ to his partner and the partner accepts the offer with probability 1. Hence player i's expected payoff conditional on being matched to bargain with j in the first period equals $1/2(1 - \delta/2 + \delta/2) = 1/2$. Since i is always matched to a player in P in the first period and expects a payoff of $1/2$ conditional on being matched with any such player, player i's MPE payoff must be $1/2$.

Consider next a player $i \in O$. By Theorem 2, every maximum matching pairs player i with some $j \in U$. Assume as before that i and j do not reach an agreement when matched with each other in the first period. Theorem 2 implies that player j's connected component C in the subnetwork induced by U in G consists of an odd number of vertices and that every maximum matching that contains the link ij covers all the nodes in C. Hence all players different from j in C—including all of j's neighbors within U—trade in the first period of the MPE. As argued above, all players in $O \setminus \{i\}$ also reach agreements in the first period of the MPE. Therefore, ij constitutes j's only link left in the aftermath of his disagreement with i. Since any maximum matching of the ensuing subnetwork has a single link, this subnetwork consists of the link ij and possibly other links ik for $k \in N$. Then the second-period expected MPE payoffs conditional on i being matched with j and not trading in the first period are at least $1/2$ for i and at most $1/2$ for j. This immediately implies that player i receives an expected payoff of at least $1/2$ by forging an agreement with j in the first period of the MPE. The bound carries over to i's MPE payoffs in the original network because i is always matched to an under-demanded player in the first period.[4]

Finally, consider a player $i \in U$. Since G is bipartite, part (3) of Theorem 2 implies that i is not linked to any player in U. By Theorem 2, in every maximum matching, player i is either matched with a player $j \in O$ or not matched with any player. In the former event, an argument that mirrors the steps establishing the payoff bounds for over-demanded players proves that i's conditional expected payoff cannot exceed $1/2$. In the latter event, all first-period matches result in agreement, leading to a subnetwork in which i is isolated and receives a zero payoff.

The proof that perfectly matched players in G receive payoffs of exactly $1/2$ and over-demanded players in G obtain payoffs greater than $1/2$ in the MPE does not invoke the hypothesis that G is bipartite, so these conclusions hold for general networks. However, Polanski provides an example of a non-bipartite network in which an under-demanded player receives a payoff above $1/2$. In networks that are not bipartite, maximum matchings may contain links between pairs of under-demanded players. Since under-demanded players can have asymmetric payoffs, the balance of bargaining power in matches between such players relative to the even split in which the involved parties receive payoffs of $1/2$ is ambiguous.

[4]In certain events that occur with positive probability, we obtain a strict lower bound of $1/2$ for player i.

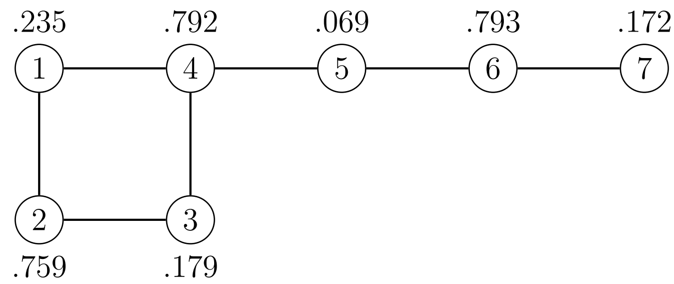

FIGURE 2. One of three profiles of MPE limit payoffs

3.3. **Bargaining with Decentralized Matching.** In the models of both Corominas-Bosch and Polanski, all agreements take place in the first period and equilibrium outcomes are efficient. The efficiency result is driven by the assumption of maximum matchings embedded in the two models. By contrast, Abreu and Manea (2012a) study a network setting with decentralized matching and sequential bargaining. In every period a single pair of linked players is matched according to some probability distribution (which depends on the set of remaining players), and each of the two parties is chosen with probability 1/2 to make an offer to his partner specifying a division of the unit surplus. If the partner accepts the offer, then the two parties exit the game with the proposed shares. If the offer is rejected, then the match dissolves and the two players remain in the game for the next period. As in the earlier models, players have a common discount factor $\delta \in (0, 1)$.

Abreu and Manea (2012a) prove that the bargaining game admits an MPE for every discount factor (the payoff-relevant variables are the set of remaining players, along with nature's selection of a link and a proposer, and the proposer's offer). However, in contrast to the centralized bargaining model of Polanski, multiple MPEs, which yield distinct payoffs, may exist. Moreover, not all matches lead to agreements. It is also possible that a matched pair of players declines to trade at some stage but forges an agreement thereafter (following a sequence of equilibrium trades). The following example demonstrates these possibilities.

Consider the bargaining game on the network from Figure 2, in which each link is equally likely to be activated for bargaining in every subnetwork. Abreu and Manea (2012a) show that for high δ this game has an MPE in which players 1 and 4 do not trade when matched with each other in the original network but an agreement emerges when the first link different from $(1, 4)$ is activated for bargaining. The limit MPE payoffs as $\delta \to 1$ are represented next to corresponding nodes in the figure.

In the constructed equilibrium, the first agreement may induce the following subgames. If players 1 and 2 (2 and 3) reach the initial agreement, then players $3, 4, 5, 6, 7$ $(1, 4, 5, 6, 7)$ are left in a bargaining game on the 5-player line network. If instead players 3 and 4 (4 and 5) reach the first agreement, the remaining subnetwork has two connected components, $\{1, 2\}$ and $\{5, 6, 7\}$ $(\{1, 2, 3\}$ and $\{6, 7\})$. Players 1 and 2 are then involved in a two-player game

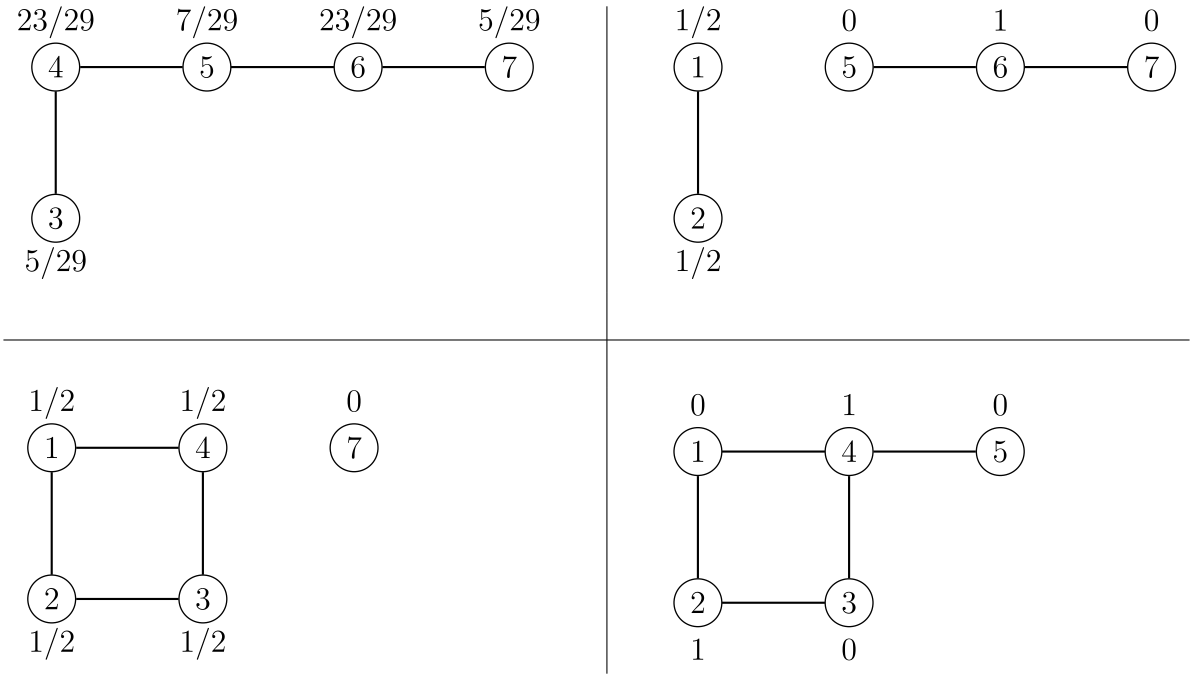

FIGURE 3. MPE payoffs in four subgames

in which they are matched to bargain with probability 1/3 in any period before another agreement is reached (in the complementary event the link $(5,6)$ or $(6,7)$ is selected for bargaining) and with probability 1 following an agreement across the link $(5,6)$ or $(6,7)$. Similarly, players $5, 6$, and 7 bargain in a version of the game on a 3-player line network (or, alternatively, a star with player 6 at the center) in which matching frequencies at every stage depend on whether players 1 and 2 traded up to that stage. The subnetworks ensuing after an agreement takes place across the links $(5,6)$ and $(6,7)$ in the original network are depicted at the bottom half of Figure 3. The limit MPE payoffs as $\delta \to 1$ in these subgames as well as other subnetworks that may arise following an agreement are represented next to each node. We compute the limit MPE payoffs in the 5-player line network in Section 3.4. The reader can consult Abreu and Manea (2012a) for more details on MPE outcomes in the other subgames. One robust property emerging from the inspection of these payoffs is that even indexed players are substantially stronger than odd indexed ones.

The intuition for this equilibrium specification is that even though players 1 and 3 occupy symmetric positions in the original network, their asymmetric behavior in matches with player 4 enhances the bargaining power of player 1. Indeed, when 3 and 4 forge the first agreement, players 1 and 2 are left in a bilateral bargaining game (see the top-right panel in Figure 3), in which player 1 obtains a limit continuation payoff of 1/2, the highest payoff

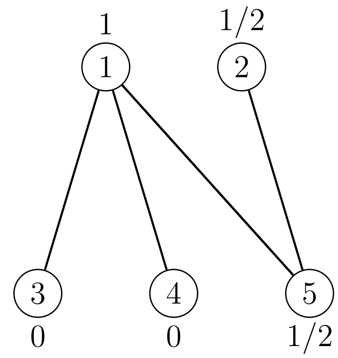

FIGURE 4. Common limit equilibrium payoffs for the three models

available to odd indexed players in any subgame. This boosts the payoff of player 1 relative to other odd indexed players. In particular, player 3 does not benefit from a similarly favorable scenario because 1 does not trade with 4 in the original network. The fact that players 3 and 4 trade when matched in the original network further lowers 3's bargaining power since player 4 has a high payoff relative to player 2. The substantial difference between the payoff of player 2 and those of 4 and 6 can be attributed to the initial agreement between players 3 and 4, which undermines player 2's strong position by leaving him stranded in the bilateral monopoly with player 1.

The difference in the bargaining power of players 1 and 3 induced by their asymmetric treatment of player 4 turns out to be sufficiently large to make the conjectured agreements and disagreements incentive compatible. Players 1 and 4 do not have incentives to trade with each other in the original network because either of them can benefit from waiting to be matched with a weaker neighbor. All other initial matches result in agreements since no link different from $(1, 4)$ connects a pair of similarly strong players (in the constructed MPE) with odd and even indices. An interesting feature of this equilibrium is that players 1 and 4 do not trade when matched in the original network, but reach an agreement with each other in subgames that arise with positive probability (for instance, following an agreement between either 2 and 3 or 5 and 6).

We can obtain another MPE for high δ by simply interchanging the roles of players 1 and 3 in the equilibrium described above. For sufficiently high δ, Abreu and Manea argue that the game has a third MPE, in which the matches $(1, 4)$ and $(3, 4)$ lead to trade with a common probability close to 0.529, while all other matches result in agreement with probability 1, in the original network. Therefore, for high δ the bargaining game has at least three MPEs with distinct payoffs.

3.4. **Model Comparisons.** It is instructive to compare the equilibrium outcomes in the three models discussed in this section thus far. For the network depicted in Figure 4, all three models predict that players 1 and 5 do not trade with each other.[5] The intuition is that player 1 has monopoly power over players 3 and 4, while player 5 acts a monopolist for player 2. Then players 1 and 5 are better off trading with their weaker neighbors than forging an agreement with each other. The three models yield identical limit payoffs as players become patient. The common limit payoffs are given by 1 for player 1, 0 for players 3 and 4, and 1/2 for players 2 and 5.

However, each model generates substantially different payoffs in the 5-player line network shown in Figure 5. This can be viewed as a bipartite network if we label players $\{1, 3, 5\}$ as buyers and $\{2, 4\}$ as sellers. In Corominas-Bosch's equilibrium, the under-demanded players $1, 3$, and 5 receive payoffs of 0, while the over-demanded players 2 and 4 obtain a payoff of 1. These extreme payoffs are explained by the strong competitive forces induced by public offers.

To compute the limit MPE payoffs in Polanski's model, note that there are three maximum matchings in this network: $\{(1, 2), (3, 4)\}, \{(1, 2), (4, 5)\}, \{(2, 3), (4, 5)\}$. Recall that all matches lead to trade in the MPE. When players 1 and 2 are matched with each other, they believe that the other match is equally likely to be $(3, 4)$ or $(4, 5)$ and anticipate that either of these matches results in an agreement. Thus if players 1 and 2 fail to reach an agreement, then they are left in a bilateral bargaining game or in a game on the line network formed, in order, by players $1, 2$, and 3. Either subgame arises with probability 1/2 conditional on the match $(1, 2)$ being formed. The limit continuation payoffs of player 1 are 1/2 in the former subgame and 0 in the latter. Then player 2 offers 1 an amount equal to 1's continuation payoff conditional on the match $(1, 2)$ ending in disagreement, which converges to $1/2 \times 1/2 = 1/4$ as players become patient; player 1 accepts the offer for every discount factor in the MPE. Similarly, player 1 offers 2 an amount that converges to 3/4, which 2 accepts. Hence the limit expected payoff of player 1 conditional on being matched (with his sole neighbor, player 2) in the first round is 1/4. Since player 1 becomes isolated when the matching $\{(2, 3), (4, 5)\}$ is drawn in the first period, an event which occurs with probability 1/3, his limit equilibrium payoff is $2/3 \times 1/4 = 1/6$. By symmetry, player 5 obtains the same payoff. Similar calculations show that players 2 and 4 receive limit MPE payoffs of 5/6, while player 3 gets 0 in the limit.

Abreu and Manea (2012a) provide a computational method for identifying MPEs for high discount factors in their model. This method proves that if each remaining link is activated for bargaining with equal probability at every stage, then the decentralized bargaining game on the 5-player line network admits an MPE in which all matches result in agreement for

[5]Note that in Polanski's model, players 1 and 5 never get the opportunity to bargain with each other.

Corominas-Bosch	0	1	0	1	0
Polanski	1/6	5/6	0	5/6	1/6
Abreu and Manea	5/29	23/29	2/29	23/29	5/29

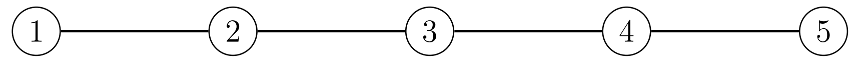

FIGURE 5. Distinct limit equilibrium payoffs in the three models

high δ.[6] The limit MPE payoffs (v_i) as $\delta \to 1$ solve the following system of equations:

$$\begin{aligned}
v_1 &= \frac{1}{4}\frac{1}{2}(1 - v_2 + v_1) + \frac{1}{4}0 + \frac{1}{4}\frac{1}{2} + \frac{1}{4}0 \\
v_2 &= \frac{1}{4}\frac{1}{2}(1 - v_1 + v_2) + \frac{1}{4}\frac{1}{2}(1 - v_3 + v_2) + \frac{1}{4}\frac{1}{2} + \frac{1}{4}1 \\
v_3 &= \frac{1}{4}0 + \frac{1}{4}\frac{1}{2}(1 - v_2 + v_3) + \frac{1}{4}\frac{1}{2}(1 - v_4 + v_3) + \frac{1}{4}0 \\
v_4 &= v_2 \\
v_5 &= v_1.
\end{aligned}$$

For example, the first term in the payoff equation for player 1 captures the fact that the link $(1, 2)$ is activated for bargaining in the first period with probability $1/4$, in which case player 1 obtains a limit payoff of $1 - v_2$ or v_1, depending on nature's selection of a proposer; the second term corresponds to the probability $1/4$ event that players 2 and 3 are matched to bargain with each other in the original network and forge an agreement that isolates player 1. The last two equations reflect network (and equilibrium) symmetries. The unique solution to the system of linear equations is given by $v_1 = v_5 = 5/29, v_2 = v_4 = 23/29, v_3 = 2/29$.

The discrepancy between the payoffs in Polanski's model and Abreu and Manea's in the example above can be assigned to the following strategic differences. In the model of Polanski, two pairs of players are matched simultaneously and an agreement in one of the matches influences the outside options of the players involved in the other. By contrast, in the model of Abreu and Manea, only one match forms at a time and the matched players continue to bargain in the same network if they fail to reach an agreement. In other words, the assumption of centralized efficient matching speeds up the process of isolating under-demanded players and reduces their bargaining power.

3.5. **Efficiency with Decentralized Bargaining.** In both the equilibrium analyzed by Corominas-Bosch and the unique MPE of Polanski, the centralized organization of matches for trade ensures efficient outcomes. Under the decentralized bargaining protocol of Abreu

[6]In this example, one can show that for any discount factor there exists a unique MPE and that every match results in agreement in the MPE.

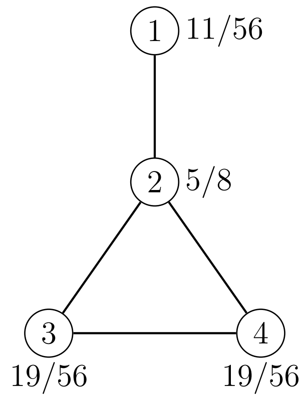

FIGURE 6. Asymptotically inefficient MPE

and Manea, links that are not part of any maximum matching may be activated for bargaining and result in inefficient agreements. Indeed, incentives for bilateral agreements are not always aligned with global welfare maximization in a setting with bargaining frictions created by decentralized random matching. The following example from Abreu and Manea (2012a) illustrates this problem. Consider the game induced by the 4-player network from Figure 6 in which each link is equally likely to be selected for bargaining as long as no agreement has been reached. In this network there exists a single maximum matching, formed by the links $(1, 2)$ and $(3, 4)$, which generates the maximum total surplus of 2. Trade between player 2 and either 3 or 4 leaves the other two players isolated, so agreements in the matches $(2, 3)$ and $(2, 4)$ are inefficient. However, Abreu and Manea prove that there exists a unique MPE for every discount factor, in which every match materializes in an agreement.[7] The unique MPE payoffs are shown to converge as $\delta \to 1$ to 11/56 for player 1, 5/8 for player 2, and 19/56 for players 3 and 4. The limit payoffs sum to 3/2, which reflects the fact that with probability 1/2 one of the inefficient matches $(2, 3)$ and $(2, 4)$ forms and in this event only one unit of surplus is created along the equilibrium path. Hence the MPE is asymptotically inefficient, generating a total welfare of 3/2, which is smaller than the available total surplus of 2.[8]

This example raises the following question: if we allow for non-Markovian behavior, is it possible to structure incentives in order to construct an asymptotically efficient subgame perfect equilibrium? The tension between the global derivation of maximum matchings and

[7] We can immediately rule out any MPE in which only efficient agreements take place. In such an equilibrium, all players would expect payoffs below 1/2 (the MPE could be decomposed into two separate bilateral bargaining games in which each of the pairs $(1, 2)$ and $(3, 4)$ is matched with a probability that evolves stochastically), creating strict incentives for the pairs $(2, 3)$ and $(2, 4)$ to trade when matched. Ruling out MPEs in mixed strategies, which could generate asymptotically efficient outcomes if the probability of inefficient agreements vanishes as $\delta \to 1$, requires more meticulous arguments.

[8] In this example, Polanski's model assumes that only the maximum matching can form and predicts an efficient equilibrium outcome in which every player expects a payoff of 1/2.

the local nature of bilateral interactions makes this a challenging question. Pairs of players who reach inefficient agreements are removed from the network permanently, so cooperative behavior that generates the maximum total surplus cannot be enforced via standard repeated game threats. An additional complication is that, when multiple maximum matchings exist, links that are part of a maximum matching in the network prevailing at some stage may cease to have this property after a series of efficient agreements takes place. Thus the notion of efficient agreements is history dependent. Despite these obstacles, Abreu and Manea (2012b) establish that an asymptotically efficient equilibrium always exists.

Theorem 5 (Abreu and Manea, 2012b)**.** *The bargaining game admits a family of asymptotically efficient subgame perfect equilibria.*

The proof constructs asymptotically efficient equilibria in which players who resist the temptation to reach inefficient agreements are rewarded by certain neighbors and players who do not conform to the rewarding procedure are punished via the threat of a sequence of trades that isolates them from the network. The evolving nature of the network structure as prescribed agreements take place complicates the design of punishment and reward schemes. Unlike in Polanski's model, it is a priori unclear which matches should lead to an agreement (the example of Section 3.3 speaks to this issue) and how the resulting stochastic evolution of the network affects the outside options in every match. Since it is difficult to specify any subgame perfect equilibrium, the construction of asymptotically efficient equilibria starts from an implicitly defined Markov strategy profile. The idea is to modify the payoffs in the original game by imposing prohibitive fines for players who reach inefficient agreements. Fines are also used to incentivize players to forge efficient agreements in certain matches. Such "forced" agreements pave the way for the isolation of deviators. An MPE of the modified game is then employed as the reference point for punishments and rewards in the original game. For matches unaffected by the payoff modifications, the strategies prescribed by this MPE are sequentially rational. The incentives to deviate resulting from the modifications of the original game need to be adjusted via explicit constructions of punishments and rewards.

While the precise calibration of reward and punishment paths for any given network structure is intricate, it is worth pointing out a key step in the proof that relates to the other models discussed in this section. Abreu and Manea prove that the limit MPE payoffs of perfectly matched and over-demanded players in the modified bargaining game for every network are greater than or equal to $1/2$.[9] To understand the role that this finding plays in the equilibrium construction, consider a link ij that is not part of a maximum matching in the subnetwork remaining at some stage, so an agreement between i and j would be

[9]This observation is reminiscent of Polanski's result but is technically more involved because, as already argued, under the decentralized matching process it is difficult to predict which matches lead to an agreement and how the network evolves.

inefficient. Then neither i nor j is under-demanded in the Gallai-Edmonds decomposition for the remaining subnetwork. By the finding mentioned above, the limit MPE payoffs of both i and j at the stage under consideration in the modified game are at least 1/2. It follows that the gains that i and j can obtain by reaching an inefficient agreement with each other, relative to the reference payoffs derived from the modified game, converge to 0 as $\delta \to 1$. Thus a small reward bounded away from 0 is sufficient to deter i from accepting a tempting offer from j (and vice versa). The proof shows how such a reward can be delivered by one of i's neighbors, a suitably chosen player k, following a series of agreements that occur with probability bounded away from 0 if i turns down a tempting offer (relative to the reference payoffs) from j. Player k is incentivized to offer the reward to i via the threat of a sequence of agreements that leaves k vulnerable to isolation. Players i and j deliver the ultimate punishment by forging an agreement (that is incentive compatible with respect to the reference payoffs) following such a history, which isolates k in the remaining network.

We lastly record that Abreu and Manea (2012b) test the robustness of their conclusion with respect to the bargaining protocol. They extend the construction of asymptotically efficient equilibria to an alternative model, which assumes that in every period a player, rather than a link, is selected stochastically and that the selected player can activate any of his links for bargaining. Once a link is activated, either party is chosen with equal probability to propose a division of the surplus as in the benchmark model.

4. Bargaining in Stationary Networks

In each of the models discussed in the previous section, there is a finite number of players and a finite set of feasible trades. Players who reach agreements are removed from the network and the pool of potential trading partners for each remaining player shrinks over time. These assumptions are realistic in small specialized markets, in which the entry of new traders is impossible (for the relevant time horizon). In this section, we consider the distinct setting of an economy that is continually replenished with new traders, so that the market composition is constant over time. Rubinstein and Wolinsky (1985) studied the first such bargaining model in the case of a stationary market with symmetric buyers and sellers. Gale (1987) expanded the analysis to markets with heterogeneous buyers and sellers. The focus of this research was to provide noncooperative foundations for general equilibrium theory. In particular, Gale found that the law of one price holds in the sense that all prices at which trade takes place in equilibrium converge to the same level as bargaining frictions disappear.

As discussed in the introduction, we should not expect the law of one price to apply to economies with network asymmetries (even as players become patient). It is nevertheless interesting to understand how the network architecture shapes local prices and bargaining power in a stationary market. Manea (2011a) pursues this matter in the context of the following bargaining model. Consider a network G connecting the finite set of nodes N, and let $(p_{ij} > 0)_{ij \in G}$ be a probability distribution over the links in G. In every period $t = 0, 1, \ldots$, a link ij in G is selected with probability p_{ij}, and each of the players i and j is chosen with equal probability to make an offer to the other player specifying a division of the unit surplus available to the pair. If the other player accepts the offer, then the two players exit the game with their agreed shares. The market is maintained in a steady state as follows: if i and j reach an agreement, then in period $t + 1$ two new players assume the same positions in the network as i and j. If the offer is rejected, then the match is dissolved and the two players remain in the market for the next matching and bargaining round. Players share a discount factor $\delta \in (0, 1)$.

The stationarity of the economy simplifies the equilibrium analysis considerably as we do not need to keep track of payoffs in subnetworks that arise endogenously following agreements in the original network. A result of Manea (2014a) implies that every subgame perfect equilibrium of the stationary bargaining game generates the same expected payoffs, which we denote by v_i, for each player active at position i in the network at the beginning of any period.[10] The equilibrium payoffs constitute the unique solution to the following system of equations:

$$v_i = \sum_{\{j|ij\in G\}} \frac{p_{ij}}{2} \max(1 - \delta v_j, \delta v_i) + \left(1 - \sum_{\{j|ij\in G\}} \frac{p_{ij}}{2}\right) \delta v_i, \forall i \in N.$$

The payoff equations capture the following equilibrium conditions. In any subgame perfect equilibrium, when i is selected to propose to j, if $1 - \delta v_j > \delta v_i$ then i offers δv_j to j, an offer which j accepts with probability 1. If $1 - \delta v_j < \delta v_i$ then i makes an offer that j rejects with probability 1.[11] Given the stationarity of the environment and the fact that all subgame perfect equilibria are payoff equivalent, player i expects a continuation payoff of δv_i if he does not forge an agreement in the current period. Therefore, in any subgame following the selection of a link and a proposer, player i obtains a payoff different from δv_i only in the event that he is selected to make an offer to a player j for which $1 - \delta v_j > \delta v_i$. This event occurs with probability $p_{ij}/2$ and yields a payoff of $1 - \delta v_j$ for i.

Manea (2011a) shows that there exist a discount factor threshold $\underline{\delta} \in (0, 1)$ and a subnetwork G^* of G such that in any equilibrium of the bargaining game for any discount factor $\delta \in (\underline{\delta}, 1)$, trade takes place with probability 1 across all links in G^* and with probability 0 for the remaining links.[12] This finding is used to demonstrate that equilibrium payoffs

[10]More generally, Manea (2014a) establishes the payoff equivalence of equilibria in markets with multiple player types in which the matching frequencies for every pair of types are exogenous and time-dependent.

[11]When $1 - \delta v_j = \delta v_i$, player i is indifferent between offering j the minimum amount δv_j necessary for an agreement and making an unacceptable offer.

[12]Formally, for every $\delta \in (\underline{\delta}, 1)$ and $ij \in G$, equilibrium payoffs satisfy $\delta(v_i + v_j) < 1$ if $ij \in G^*$ and $\delta(v_i + v_j) > 1$ otherwise (it is impossible that $\delta(v_i + v_j) = 1$).

converge to a limit $(v_i^*)_{i\in N}$ as $\delta \to 1$. The following concepts, which rely on the endogenous network of agreements G^*, prove helpful in characterizing the limit equilibrium payoffs. A nonempty set of players is *mutually estranged* if it is G^*-independent. The set of *partners* for a mutually estranged set M is defined as $\mathcal{L}^{G^*}(M)$.

Fix a mutually estranged set M with partner set L. For high δ, the set L spans the relevant bargaining opportunities of the players in M. Then the group M is collectively weak if L has a relatively small cardinality. The relevant measure of M's strength turns out to be the *shortage ratio* of M, defined as the ratio of the number of partners to estranged players, $|L|/|M|$. Manipulating the payoff equations, we find that for every mutually estranged set M with partner set L, the equilibrium payoffs $(v_i)_{i\in N}$ for sufficiently high δ satisfy

$$\sum_{j\in L} v_j \geq \sum_{i\in M} v_i.$$

Loosely speaking, M cannot have more "collective bargaining power" than its partner set L. The inequality above implies that the ratio of the limit equilibrium payoffs of the worst-off player in M and the best-off player in L does not exceed the shortage ratio of M. Based on this observation, Manea (2011a) establishes the following bounds on limit equilibrium payoffs:

$$\begin{aligned} \min_{i\in M} v_i^* &\leq \frac{|L|}{|M|+|L|} \\ \max_{j\in L} v_j^* &\geq \frac{|M|}{|M|+|L|}. \end{aligned}$$

A key step in the analysis shows that the bounds on limit equilibrium payoffs corresponding to a mutually estranged set M and its partner set L need to bind unless the worst-off player in M belongs to a mutually estranged set with a lower shortage ratio. The intuition is that in every connected component of G^* (where not all players have limit payoffs of $1/2$), some mutually estranged players and their partners share the unit surplus according to the corresponding shortage ratio as $\delta \to 1$. Therefore, the bounds above must bind for any mutually estranged set that achieves the lowest shortage ratio.

Define $r_1 = \min_{M \text{ is } G\text{-independent}} |\mathcal{L}^G(M)|/|M|$ and let M_1 be the union of all G-independent sets M that minimize the expression $|\mathcal{L}^G(M)|/|M|$. Set $L_1 = \mathcal{L}^G(M_1)$. Manea (2011a) shows that if $r_1 < 1$, then r_1 constitutes the minimum shortage ratio over all mutually estranged sets, while M_1 represents the largest mutually estranged set minimizing the shortage ratio and L_1 serves as its partner set. This means that when the lowest shortage ratio is smaller than 1, it can be computed by restricting attention to sets that are independent in the original network G instead of the a priori unknown agreement network G^*, with corresponding sets of partners formed by neighbors in G rather than G^*. The analysis reveals that all members of M_1 obtain minimum limit equilibrium payoffs, given by $r_1/(r_1+1)$, and all members of

L_1 achieve the maximum equilibrium limit payoffs of $1/(r_1+1)$. By definition, the players in M_1 are only linked to players in L_1 in the original network G. Moreover, since M_1 consists of all players with minimum limit payoffs and every member of L_1 has at least one link to M_1, players in L_1 do not have incentives to trade with players outside M_1 for high δ. Hence G^* does not contain any links from $L_1 \cup M_1$ to the remaining set of nodes. Then the nodes in $L_1 \cup M_1$ can be removed from G and the network induced by the remaining nodes can be analyzed as a separate submarket.

The arguments above lead to the following network decomposition algorithm. Define the sequence $(r_s, M_s, L_s, N_s, G_s)_s$ recursively as follows. Let $N_1 = N$ and $G_1 = G$. For $s \geq 1$, the algorithm terminates if $N_s = \emptyset$. Otherwise, let

$$r_s = \min_{M \subseteq N_s,\ M \text{ is } G\text{-independent}} \frac{|\mathcal{L}^{G_s}(M)|}{|M|}. \tag{2}$$

If $r_s \geq 1$, then the algorithm stops. Else, let M_s be the union of all minimizers M in (2) and define $L_s = \mathcal{L}^{G_s}(M_s)$. Set $N_{s+1} = N_s \setminus (M_s \cup L_s)$ and let G_{s+1} be the subnetwork of G induced by the nodes in N_{s+1}. Denote by $\overline{s}$ the (finite) step at which the algorithm ends.

A key lemma proves that if $r_s < 1$, then M_s is G-independent and minimizes the expression in (2). Hence, at every step, the algorithm determines the largest mutually estranged set minimizing the shortage ratio in the subnetwork induced by the remaining players and removes the corresponding estranged players and partners. This definition ensures that all players with extremal limit payoffs in the remaining network are identified simultaneously, so that the removed players trade only with one another in equilibrium for high δ. Then we can treat the residual subnetwork as a separate market.

The algorithm ends when all players have been removed or $|\mathcal{L}^{G_s}(M)| \geq |M|$ for every G-independent set $M \subseteq N_s$. The sequence $(r_s)_s$ is strictly increasing and the sets $M_1, L_1, \ldots, M_{\overline{s}-1}, L_{\overline{s}-1}, N_{\overline{s}}$ partition the set of nodes N. The main result of Manea (2011a) establishes that the limit equilibrium payoff of each player is determined by his cell in this partition.

Theorem 6 (Manea, 2011a)**.** *The limit equilibrium payoffs as $\delta \to 1$ in the stationary bargaining game are given by*

$$\begin{aligned} v_i^* &= \frac{r_s}{r_s+1}, \forall i \in M_s, \forall s < \overline{s} \\ v_j^* &= \frac{1}{r_s+1}, \forall j \in L_s, \forall s < \overline{s} \\ v_k^* &= \frac{1}{2}, \forall k \in N_{\overline{s}}. \end{aligned}$$

As an illustration, for the network from Figure 7, the algorithm ends in $\overline{s} = 2$ steps. The relevant outcomes are $r_1 = 1/2, M_1 = \{3, 4\}, L_1 = \{1\}$ and $r_2 = 1, N_2 = \{2, 5\}$. Thus the

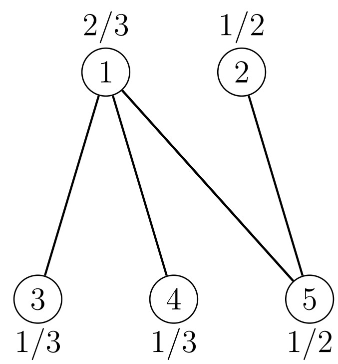

FIGURE 7. Limit equilibrium payoffs in the stationary model

limit equilibrium payoffs are given by 2/3 for player 1, 1/3 for players 3 and 4, and 1/2 for players 2 and 5. Recall that all the models discussed in Section 3 predict that in this network player 1 receives a payoff of 1, while players 3 and 4 obtain a payoff of 0 in the limit $\delta \to 1$. To examine this discrepancy, consider a scenario in which players 1 and 3 are matched and reach an agreement with each other in the first period. Then, in the setting of Section 3, players 1 and 3 are removed from the network without replacement, and player 4 is left isolated, in which case his continuation payoff is 0. By contrast, in the stationary economy, a new trader enters the market at node 1 in the second period, and player 4 may get a chance to bargain with this trader. In equilibrium, player 4 capitalizes on the perpetual rebirth of bargaining partners at node 1 to secure a limit payoff of 1/3. Relatedly, limit payoffs in the stationary case sum to 7/3, which is greater than the maximum total surplus of 2 in this network. The reason is that the entrance of a new player at node 1 following the first-period agreement across the link $(1, 3)$ creates the potential for realizing a unit of surplus in a match with the player present at node 4 from the first period. In general, since newly born players may share surplus with older ones, the total sum of limit payoffs generated by the positions in the network does neither reflect the size of any matching in the static network, nor constitute a suitable welfare measure in the stationary setting.

An important implication of the characterization of limit equilibrium payoffs is that submarkets emerge endogenously in equilibrium. The group L_s acts as an oligopoly for the players in M_s. The *oligopoly subnetwork* induced by the set of nodes $L_s \cup M_s$ is a union of connected components of the agreement network G^*. For high δ, every equilibrium transaction takes place between members of the same oligopoly subnetwork. The limit prices as $\delta \to 1$ are uniform within every such subnetwork. In equilibrium, each trader self-selects into the most favorable submarket to which he has access.

Note that limit equilibrium payoffs do not depend on the relative probabilities $(p_{ij} > 0)_{ij \in G}$ with which links are activated for bargaining. Furthermore, the conclusions of the model extend to a setting in which multiple disjoint links are activated for bargaining in every period. Relatedly, the model and the results admit an alternative interpretation whereby there is a continuum of players at each node and bargaining proceeds simultaneously for a positive mass of heterogeneous matches. Manea (2014b) provides foundations for the steady state assumption in a setting with a continuum of players where new traders optimally decide whether to enter the market for a (small) cost.

4.1. **Network Formation.** To approach the question of network formation in the stationary bargaining model, it is useful to first define and characterize equitable networks. A network G is called *equitable* if all players receive limit equilibrium payoffs of $1/2$ in G.[13] Theorem 6 implies that a network G is equitable if and only if $|\mathcal{L}^G(M)| \geq |M|$ for every G-independent set M. Manea (2011a) shows that the latter condition is equivalent to G containing a subnetwork that covers all of G's nodes and consists of a disjoint union of a matching and cycles with odd numbers of nodes. Since bipartite networks cannot contain odd cycles, the necessary and sufficient condition for a bipartite network G to be equitable boils down to G admitting a perfect matching (cf. Theorem 1).

Manea (2011b)—the online appendix to Manea (2011a)—shows that if the players active in the first-period market are responsible for forming the network before they engage in the stationary bargaining game, and there are no linking costs, then (1) adding a new link to a network cannot decrease the limit payoffs of either of the players it connects; and (2) a network is pairwise stable[14] if and only if it is equitable. Gauer (2014) studies network formation in the stationary bargaining model under the assumption that every link entails positive costs for both parties involved. Clearly, in this setting equitable networks in which some links are "redundant" for attaining equitability cannot be pairwise stable. In the class of equitable networks, only "skeleton" networks—formed by disjoint unions of a matching and cycles of odd length—can survive pairwise stability. Gauer confirms that such skeleton networks are indeed pairwise stable for small linking costs. However, he finds that some non-equitable networks (in particular, networks formed by odd cycles and a single isolated player) are also pairwise stable for small costs. Gauer then proceeds to partially characterize pairwise stable networks for arbitrary linking costs and identifies a richer collection of stable structures for intermediate costs.

In Manea (2011b), the definition of pairwise stability is adjusted to bipartite networks to account for the fact that only buyer-seller pairs can contemplate forming new links. The modified solution concept is coined *two-sided pairwise stability.* We say that a bipartite network is *non-discriminatory* if all transactions take place at the same limit price (i.e., all buyers obtain identical limit equilibrium payoffs). When linking costs are zero, Manea

[13]Recall that limit equilibrium payoffs depend on the collection of links activated for matching with positive probability but not on the relative activation probabilities.

[14]See Jackson and Wolinsky (1996) for the definition of pairwise stability.

(2011b) proves that a bipartite network is two-sided pairwise stable if and only if it is non-discriminatory. Manea (2011a) develops a simplified version of the network decomposition algorithm for the case of bipartite networks. The payoff characterization derived from that alternative decomposition reveals that a bipartite network is non-discriminatory if and only if for every buyer subset M, the ratio $|\mathcal{L}^G(M)|/|M|$ is greater than or equal to the ratio of seller to buyer nodes in the network. Polanski and Vega-Redondo (2014) extend the result that two-sided pairwise stable networks are non-discriminatory to a setting in which buyers and sellers have heterogeneous valuations. They also generalize the necessary and sufficient condition for non-discriminatory pricing.

4.2. **Coalitional Bargaining.** Nguyen (2014) extends the model of Manea (2011a) to a setting in which coalitions of various sizes can create different amounts of surplus. A coalition along with an ordering of its members is randomly drawn, and then the first player in the order proposes a split of the surplus available to the coalition and the other members decide, in order, whether to accept or reject the offer. An agreement is reached only if the proposal is unanimously accepted. The members of the agreeing coalition exit the market and are replaced by clones. Players have a common discount factor δ.

Nguyen shows that all stationary equilibria are payoff equivalent and characterizes the unique equilibrium payoffs as the solution to a convex optimization problem. For the special case of Manea (2011a), the convex program reduces to

$$\begin{aligned} \min_{(v_i)_{i\in N},(z_{ij})_{ij\in G}} \quad & 2\delta(1-\delta)\sum_{i\in N} v_i^2 + \sum_{ij\in G} p_{ij} z_{ij}^2 \\ \text{s.t.} \quad & \delta(v_i+v_j)+z_{ij} \geq 1, \forall ij \in G. \end{aligned}$$

The main idea for this characterization is that the dual optimization problem boils down to solving the equilibrium payoff equations. Nguyen proves that the limit equilibrium payoffs as $\delta \to 1$ constitute the optimal solution for the following convex program:

$$\begin{aligned} \min_{(v_i)_{i\in N}} \quad & \sum_{i\in N} v_i^2 \\ \text{s.t.} \quad & v_i+v_j \geq 1, \forall ij \in G. \end{aligned}$$

This finding leads to an alternative proof for Theorem 6. Nguyen also applies his result to the study of multilateral bargaining in the context of intermediation in networks and cooperation within overlapping communities.

5. The Assignment Game and Related Noncooperative Models

Shapley and Shubik (1971) provide an elegant and powerful analysis of two-sided markets based on cooperative game theory. In their assignment game, as in the other models discussed in this chapter, buyers have unit demand and sellers have unit supply for a heterogeneous good. The sets of buyers and sellers, which we denote by B and S, respectively, are finite. The good supplied by seller j, which we call *good* j for brevity, is worth nothing to j and $a_{ij} \geq 0$ to buyer i. An *assignment* is a mapping $\mu : B \cup S \to B \cup S$ with $\mu(i) \in S \cup \{i\}$ for $i \in B$ and $\mu(j) \in B \cup \{j\}$ for $j \in S$, which satisfies $\mu(\mu(k)) = k$ for all $k \in B \cup S$. For a buyer-seller pair (i, j), the condition $\mu(i) = j$ is equivalent to $\mu(j) = i$ and indicates that buyer i is *assigned* good j under μ. If $\mu(i) = i$ ($\mu(j) = j$) for buyer i (good j), then i (j) is left *unassigned* by μ. An assignment μ is *efficient* if it maximizes the *total surplus* $\sum_{i \in B, \mu(i) \neq i} a_{i\mu(i)}$. We refer to the achieved maximum as the *maximum total surplus*. Assuming that every coalition can organize trade efficiently, this setting generates a *cooperative game* with transferable utility in which the value of each coalition is determined by the maximum total surplus available in the economy consisting only of its members. Shapley and Shubik establish several facts about the structure of the core of this cooperative game, which is coined the *assignment game*.[15]

Theorem 7 (Shapley and Shubik, 1971)**.**

(1) The core of the assignment game is nonempty.

(2) Any profile of core payoffs corresponds to a competitive equilibrium outcome in an economy where each good j is assigned an individual price p_j.

(3) The set of core payoffs forms a lattice with the partial order under which a payoff profile dominates another if all buyers (weakly) prefer the former profile, and all sellers prefer the latter.

(4) The core contains a payoff profile that all buyers prefer to any other point in the core—the buyer-optimal core payoff—which corresponds to the minimum competitive equilibrium price vector. Similarly, there exists a seller-optimal core payoff corresponding to the maximum competitive equilibrium prices.

This result lends an interesting interpretation to the equilibrium constructed by Corominas-Bosch (2004), which we discussed in Section 3.1. A bipartite network G naturally defines an assignment game in which $a_{ij} = 1$ if $ij \in G$ and $a_{ij} = 0$ otherwise. In this assignment game, every coalition M creates a value equal to the maximum total surplus in the subnetwork induced by M in G. Corominas-Bosch (1999) points out that at every core allocation of the associated assignment game, under-demanded players in G must receive a zero payoff, which then implies that over-demanded players in G obtain a payoff of 1. The buyer-optimal core allocation yields payoffs of 1 and 0 for perfectly matched buyers and sellers, respectively. The payoffs of perfectly matched buyers and sellers are reversed in the seller-optimal core

[15]For an exhaustive survey of the research on the assignment game, refer to Section 8 of Roth and Sotomayor (1990). The reader is also encouraged to peruse the original paper of Shapley and Shubik, which is superbly written. Koopmans and Beckman (1957) provide an early analysis of the connection between efficient allocations and competitive outcomes in the assignment game.

outcome. Consequently, Corominas-Bosch (1999) concludes that the limit payoffs in the equilibrium described by Theorem 3 represent the average of the buyer- and the seller-optimal core allocations.

5.1. **A Decentralized Price Adjustment Mechanism.** The existence of extremal competitive prices for the assignment game facilitates the design of tatonnement processes that describe some intuitive price adjustment mechanisms. Crawford and Knoer (1981) develop a *decentralized* price formation process that converges to the lowest competitive prices. Their approach resembles the deferred acceptance algorithm of Gale and Shapley (1962). Let ε be a small positive number, which serves as the minimum unit of price increments. At every stage in the price adjustment mechanism, buyer i faces a personalized price p_{ij} for good j; all prices are initialized at 0. Each buyer i demands a good j that maximizes his payoff $a_{ij} - p_{ij}$ given the current price levels. A buyer does not submit any demand if no good yields positive payoffs at the prevailing prices. Each seller j tentatively assigns good j to one of the buyers i who demand it at the highest price p_{ij}. Then the personal price p_{ij} for every buyer i who demanded, but was not assigned, good j is increased by ε. The other prices do not change. The price adjustment process ends when no new demand is submitted, at which point the existing assignments become definitive and each buyer i assigned a good j pays his personal price p_{ij} to seller j. Crawford and Knoer prove that the final prices converge to the minimum competitive prices as $\varepsilon \to 0$.[16]

Note that the "decisions" at every step in the mechanism of Crawford and Knoer rely only on "local" information available to the relevant traders. At each step, buyer i's demand depends solely on his personal valuations $(a_{ij})_{j\in S}$ and prices $(p_{ij})_{j\in S}$, and similarly seller j's tentative assignment is determined by the set of buyers B_j who demand good j and the prevailing price levels $(p_{ij})_{i\in B_j}$. The steady escalation in personal prices can be interpreted as voluntary bidding by buyers whose price offers are rejected. However, it is not clear whether buyers have incentives to participate in this type of gradual Bertrand competition, particularly in situations in which their valuations are private information. A working paper by Gautier and Holzner (2015) explores this issue in the context of a labor market in which firms bid against one another in order to attract workers and there is asymmetric information regarding the network formed by job applications. Each worker observes and interacts directly only with the set of firms to which he applied and vice versa. In a departure from the mechanism of Crawford and Knoer, the wage tentatively accepted by a worker is revealed to all of the firms to which the worker submitted applications. It is assumed that all worker-firm pairs are equally productive, which is captured by an assignment game with $a_{ij} = 1$ if

[16]Kelso and Crawford (1982) generalize the result to a setting in which buyers have multi-unit demand and substitutable preferences over goods.

worker j applied to firm i and $a_{ij} = 0$ otherwise. Gautier and Holzner construct a noncooperative game of offers and counteroffers that (approximately) implements the firm-optimal core outcome for the associated assignment game.

5.2. **Centralized Simultaneous Auctions.** Demange, Gale, and Sotomayor (1986) propose a *centralized* mechanism based on *simultaneous ascending-price auctions* that converges to the buyer-optimal core payoff. Again, $\varepsilon > 0$ denotes a small unit for price increments. At every stage $t = 0, 1, \ldots$, the auctioneer announces a price p_{jt} for each good j. Initial prices are set at $p_{j0} = 0$ for all j. At stage t, every buyer i submits his optimal demand set, which consists of the goods j that maximize i's payoff $a_{ij} - p_{jt}$ at stage t prices; i submits an empty demand set if $a_{ij} - p_{jt} < 0$ for all j. Submitted demands induce a network G_t in which i is linked to j if buyer i demands good j. If there exist matchings of G_t that cover all buyers with nonempty demand sets, then the auction ends and the market is cleared according to one of these matchings at stage t prices. In the absence of such matchings, Hall's Theorem implies the existence of a subset of buyers M with the property that $|M| > |\mathcal{L}^{G_t}(M)|$. We say that any set of goods $\mathcal{L}^{G_t}(M)$ demanded by such a group of buyers M is *over-demanded*. The auctioneer selects an over-demanded set that is minimal with respect to inclusion and increases the price of every good j in the set by ε for stage $t + 1$, i.e. $p_{j(t+1)} = p_{jt} + \varepsilon$. The prices of all other goods j' are left unchanged, $p_{j'(t+1)} = p_{j't}$.

Theorem 8 (Demange, Gale, and Sotomayor, 1986)**.** *The simultaneous auction ends in a finite number of steps and the final prices it generates converge as $\varepsilon \to 0$ to the minimum competitive prices in the associated economy.*

In an interpretation of the assignment game in the context of a bipartite network of buyers and sellers, Kranton and Minehart (2001) consider a noncooperative game in which sellers supply identical goods and each buyer has the same value for all the goods offered by his neighbors. Buyers first decide simultaneously with which sellers to form links; the cost $c > 0$ of every linked is paid by the buyer who forms it. The ensuing pattern of links defines a network G, which is observed by all buyers. Then nature draws the private value v_i of each buyer i independently from an identical distribution F.[17] Only buyer i learns the value v_i. The realized network and values generate an assignment game a in which $a_{ij} = v_i$ if $ij \in G$ and $a_{ij} = 0$ otherwise. Buyers participate in a simultaneous ascending-price auction for this assignment game. The auction proposed by Kranton and Minehart is similar to the one

[17]The conclusions of this model extend immediately to heterogeneous linking costs and value distributions.

developed by Demange, Gale, and Sotomayor (1986) except for a few design differences.[18] Kranton and Minehart prove a strong efficiency result.

Theorem 9 (Kranton and Minehart, 2001)**.** *There exists a perfect Bayesian equilibrium with the following properties:*

(1) in any subgame following the formation of a network, each buyer i competes in the auction of every seller he is linked to until the price reaches his valuation v_i, at which point he drops out of the auction, and the resulting allocation for every realization of buyer values constitutes a minimum price competitive equilibrium outcome in the associated economy;

(2) at the ex-ante network formation stage, buyers form a network that maximizes total expected welfare net of linking costs.

This efficiency result is related to a key finding of Demange (1982) and Leonard (1983), which reveals that at the buyer-optimal core allocation of any assignment game $(a_{ij})_{i\in B, j\in S}$, each buyer receives the marginal value he contributes to the grand coalition.[19] Demange and Leonard independently considered the situation in which the vector of valuations $(a_{ij})_{j\in S}$ constitute buyer i's private information and analyzed incentives for truthful revelation of preferences in a mechanism in which every buyer reports a vector of valuations and then the mechanism designer implements a minimum price competitive equilibrium of the economy with reported valuations. Using the characterization of the buyer-optimal core allocation in terms of welfare externalities, they proved that this direct revelation mechanism is *strategy-proof*: it is a weakly dominant strategy for every player to report his true valuations. The first part of the result of Kranton and Minehart refines this conclusion by addressing the issue of interim incentives in the *dynamic* environment of the simultaneous auction mechanism.[20] The intuition for the result on efficient network formation is that the positive welfare externalities

[18]Prices increase continuously and each buyer decides at every point whether to continue competing in or drop out of the auction of any linked seller. Maximal submarkets induced by the network of buyer-seller pairs (i, j) with the property that buyer i has not dropped out of seller j's auction are cleared as prices increase.

[19]Elliott (2015) develops a characterization of the buyer-optimal core payoffs which provides further economic insight. A buyer i and a seller j are said to be an *efficient match* for each other if i receives good j at some efficient assignment. Buyer i constitutes an *outside option* for seller j if i is an efficient match for j in the assignment game obtained by eliminating a buyer who is an efficient match for j in the original game. An *opportunity path* originating at buyer i is a sequence $i_0, j_0, i_1, j_1, \ldots$ with $i = i_0$ such that for $k \geq 0$ seller j_k is an efficient match for buyer i_k and buyer i_{k+1} is an outside option for seller j_k (the last player in the sequence is either a buyer with no efficient match or a seller with no outside option). Elliott proves that the buyer-optimal core payoff of each buyer i can be computed by alternatively adding and subtracting the values along any opportunity path $i_0, j_0, i_1, j_1, \ldots$ (i.e., $a_{i_0 j_0} - a_{i_1 j_0} + a_{i_1 j_1} - \ldots$) originating at i. This finding generalizes a result of Kranton and Minehart (2000).

[20]While Kranton and Minehart restrict attention to assignment games derived from networks with homogeneous goods, Gul and Stacchetti (2000) obtain a general version of this result for settings with multi-unit demand for heterogeneous goods and utility functions that satisfy the gross substitutes condition.

created by adding a link to any network are fully captured in the auction by the buyer who forms and bears the cost of that link.

To test the limits of Kranton and Minehart's efficiency result, Elliott (2015) considers a network formation game in which (1) payoffs are determined by a fixed convex combination of the buyer- and the seller-optimal core outcomes and (2) establishing a link requires investments from both buyer and seller. When linking costs are shared according to an exogenous rule, hold-up problems may lead to underinvestment in link formation, which in some cases eliminates all the potential gains from trade; overinvestment in links that do not affect the maximum total surplus (but change its division) is also possible and can dissipate half of the gains from trade. The alternative assumption that buyers and sellers negotiate the split of linking costs endogenously eliminates underinvestment but may exacerbate overinvestment.

5.3. **A Noncooperative Bargaining Game.** All models discussed thus far in this section predict that an efficient allocation will arise. Elliott and Nava (2014) challenge this conclusion by analyzing a noncooperative bargaining model derived from an assignment game $(a_{ij})_{i\in B, j\in S}$. Let $(p_k > 0)$ denote a probability distribution over the set of players. At every date $t = 0, 1, \ldots$, a single player k who has not traded by date t is recognized as the *proposer* with probability p_k.[21] Player k then chooses a bargaining partner h from the other side of the market and proposes a price (which determines how the surplus a_{kh} is divided between the two parties). If player h accepts the offer, then k and h trade at the agreed price and exit the market. Players have a common discount factor $\delta \in (0, 1)$.[22]

To gain some intuition for the competitive forces induced by this bargaining protocol, consider the monopoly scenario illustrated in Figure 8 in which a single seller, player 0, bargains with two buyers, players 1 and 2. Suppose that buyer 1 has a higher value than buyer 2 for the good supplied by the seller, i.e., $a_{10} > a_{20}$, and that each of the three players is recognized as the proposer with probability 1/3. If $a_{10} > 2a_{20}$, then for δ above a certain threshold, there exists a unique MPE, in which the seller trades exclusively with buyer 1 and both players 0 and 1 obtain limit payoffs of $a_{10}/2$ as $\delta \to 1$. If instead $a_{10} \leq 2a_{20}$, then there exists a family of MPEs for high δ that yields limit payoffs of $a_{20}, a_{10} - a_{20}$, and 0 as $\delta \to 1$ for players 0, 1, and 2, respectively. We conclude that in both cases the limit MPE payoffs of players 0 and 1 are summarized by the formulae $\max(a_{10}/2, a_{20})$ and $\min(a_{10}/2, a_{10} - a_{20})$, respectively, whereas buyer 2 receives a limit MPE payoff of 0.

[21] Probabilities are not renormalized as players trade and exit the market, so there may be stages where no proposer is selected.

[22] This bargaining protocol is similar to the second protocol analyzed by Abreu and Manea (2012b), except that in the present setting a trader selected by nature automatically becomes the proposer in the bargaining round with his chosen partner, while in the model of Abreu and Manea either party serves as the proposer with probability 1/2 after the selected trader activates a link.

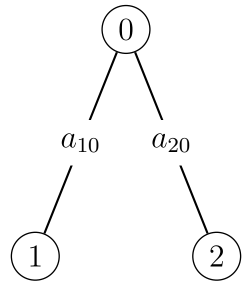

FIGURE 8. An example with one seller and two buyers

A different bargaining protocol proposed by Manea (2014c) in the context of an intermediation model generates the same predictions for the limit MPE payoffs. In the setting of the example above, the protocol stipulates that the seller gets a new opportunity to choose one of the buyers as a bargaining partner in every period, and both the seller and the chosen partner assume the role of the proposer with probability $1/2$ after the choice is made. If $a_{10} < 2a_{20}$, then in either model the seller takes advantage of his outside option of trading with the low-value buyer with positive probability, but the probability of such an inefficient agreement converges to 0 as $\delta \to 1$.[23] As discussed in Manea (2014c), the common limit MPE payoffs can be explained as follows. When players are patient, the seller is able to pick his preferred outcome between two scenarios: (1) a bilateral monopoly situation corresponding to a two-player bargaining game between the seller and buyer 1 (where buyer 2 is absent); and (2) a second-price auction in which the seller extracts all the surplus created by the low-value buyer. Thus the outside option of trading with buyer 2 is "credible" in equilibrium only if $a_{10} \leq 2a_{20}$.[24]

Elliott and Nava explore the extent to which the conclusion of asymptotic efficiency of MPEs generalizes beyond the simple example above. To approach this issue, they restrict attention to assignment games that have a unique efficient assignment, denoted η. An agreement between buyer i and seller j is said to be *efficient* if $\eta(i) = j$. In this setting, Elliott and Nava call a family of MPEs for $\delta \in (0, 1)$ *limiting efficient* if for every date t when a player k with $k \neq \eta(k)$ is selected as the proposer following a series of efficient agreements, k trades with $\eta(k)$ at date t with a conditional probability that converges to 1 as $\delta \to 1$.

[23]While the limit MPE payoffs coincide under the two bargaining protocols, the structure of equilibrium agreements for high δ is not exactly analogous when $a_{10} < 2a_{20}$. In the MPE of Manea's model for this case, the seller chooses to bargain with the low-value buyer 2 with positive probability and trades with buyer 2 with probability 1 conditional on this choice. In the model of Elliott and Nava, the seller makes an offer to buyer 2 with probability 0 conditional on being recognized as the proposer, whereas in the event that buyer 2 is the proposer, he forges an agreement with the seller with positive probability.

[24]See Shaked and Sutton (1984) for a classic study of outside options.

Motivated by the example with two buyers and a single seller discussed above, define the *core outside option* o_k of player k as the maximum surplus that k can create by trading with a player on the opposite side of the market who is left unassigned by η,

$$o_k = \max \{a_{kh} \mid \eta(h) = h \ \& \ h \text{ is on the opposite side of the market from } k\}.$$

In a limiting efficient MPE, all traders unassigned under η remain in the market indefinitely with limit probability 1 and receive limit payoffs of 0, so every player should obtain a limit payoff at least as high as his core outside option. Then an agreement between buyer i and seller $\eta(i)$ departs from their "Rubinstein shares" $(p_i a_{i\eta(i)}/(p_i + p_{\eta(i)}), p_{\eta(i)} a_{i\eta(i)}/(p_i + p_{\eta(i)}))$ in a two-player game in which they bargain in isolation only if the outside option of either player "binds." Based on this intuition, Elliott and Nava define the *shifted Rubinstein payoff* u_k of a player k assigned under η as follows:[25]

$$u_k = \begin{cases} o_k & \text{if } o_k \geq \frac{p_k}{p_k + p_{\eta(k)}} a_{k\eta(k)} \\ a_{k\eta(k)} - o_{\eta(k)} & \text{if } o_{\eta(k)} \geq \frac{p_{\eta(k)}}{p_k + p_{\eta(k)}} a_{k\eta(k)} \\ \frac{p_k}{p_k + p_{\eta(k)}} a_{k\eta(k)} & \text{otherwise.} \end{cases}$$

For any player left unassigned by η, the shifted Rubinstein payoff is simply defined to be 0. Elliott and Nava prove that limiting efficient MPEs must yield the shifted Rubinstein payoffs and provide a necessary condition for the existence of such equilibria.

Theorem 10 (Elliott and Nava, 2014)**.** *The payoffs for every family of limiting efficient MPEs converge to the shifted Rubinstein payoffs as players become patient. Limiting efficient MPEs exist only if the profile of shifted Rubinstein payoffs belongs to the core of the underlying assignment game.*

The first part of this result follows from an inductive argument on the number of buyers assigned under η. Consider a subgame for a family of limiting efficient MPEs where at least two buyers assigned under η are left. Then for each remaining buyer-seller pair $(i, \eta(i))$, there is positive limit probability that another efficient trade takes place. By the inductive hypothesis, $(i, \eta(i))$ obtain the shifted Rubinstein payoffs following that trade, and this property extends to the earlier stage. The induction base case, in which a single buyer is assigned under η, has the flavor of the example discussed above. The second part of the result follows from the first part, which establishes that any limiting efficient MPE must yield the shifted Rubinstein payoffs $(u_k)_{k \in B \cup S}$. If this payoff profile does not belong to the core of the assignment game, it must be that $u_i + u_j < a_{ij}$ for a buyer-seller pair (i, j). Then buyer i can obtain a limit payoff strictly greater than u_i by deviating from the underlying equilibria to make an offer slightly higher than u_j to j when he is selected as the proposer.

[25]The first two cases correspond to the outside options of k and $\eta(k)$ binding. The efficiency of the assignment η implies that it is not possible for both outside options to bind simultaneously.

Elliott and Nava argue that Theorem 10 has negative welfare implications because the shifted Rubinstein payoffs and the core of the assignment game are not systematically related. However, limiting efficiency is a strong welfare criterion. A more natural efficiency notion in this setting would be a version of Abreu and Manea's (2012b) concept of asymptotic efficiency. Specifically, a family of MPEs for $\delta \in (0,1)$ with corresponding payoffs $(v_k^\delta)_{k \in B \cup S}$ is *asymptotically efficient* if $\lim_{\delta \to 1} \sum_{k \in B \cup S} v_k^\delta$ equals the maximum total surplus. Note that, under the assumption of a unique efficient assignment, a family of asymptotically efficient MPEs must eventually lead to the efficient assignment with a probability converging to 1 as $\delta \to 1$. Limiting efficiency requires that efficient agreements arise without any endogenous delay (in addition to the exogenous delay inherent to the bargaining protocol) as players become patient. Clearly, every family of limiting efficient MPEs is necessarily asymptotically efficient. At this time, it is not known whether a reverse relationship holds and whether Theorem 10 extends to also characterize asymptotically efficient MPEs.

6. Conclusion

This chapter surveys the growing body of research on bilateral trade in networks.[26] The theoretical models explored here advance our understanding of how local competitive forces shape trading outcomes and the balance of bargaining power in different parts of a network. The predictions of these models are qualitatively similar for some networks yet diverge sharply for others. We attempted to explain the modeling choices that account for discrepancies, but in many cases a complete comparison between models is intractable. We found that the nature of the matching process and the selection of the solution concept have substantially different implications for the dynamics of trade, the balance of bargaining power, welfare properties of market outcomes, and equilibrium multiplicity. In some of the models, the question of what types of networks and equilibrium concepts lead to inefficient trade or multiple equilibria requires further investigation. It would also be desirable to get a better grasp of which results are robust to the specification of the matching and bargaining process. Relatedly, it would be interesting to know what bargaining protocols offer more realistic predictions for markets in which participants meet and bargain "freely." Empirical analysis of real markets where network data is available and laboratory experiments can shed light on this issue.

We finally comment on a couple of restrictive assumptions that are prevalent in the existing literature. The work surveyed here deals exclusively with the case in which buyers have unit demand and sellers have unit supply. It would be useful to develop tractable models in which traders have multi-unit supply and demand. Another strong assumption maintained through

[26]We restricted attention to markets in which intermediation is not possible; a review of the literature on intermediation in networks can be found in Chapter 27.

most of the literature is that the underlying network is common knowledge among traders. Relaxing this assumption in plausible ways to allow for incomplete information about local network structure constitutes another important topic for future research.

Appendix A. Proof of the First Part of Theorem 4

We start by showing that all matches result in first-period agreement in any MPE for every discount factor. Fix an MPE for the network G and the discount factor δ. Consider a link ij that is part of a maximum matching of G. Let u_i^{ij} and u_j^{ij} denote the expected payoffs of players i and j, respectively, in the MPE conditional on the event that they are matched with each other and do not reach an agreement in the first period. These payoffs are evaluated under the assumption that the other players conform to their equilibrium strategies, with discounting applied from the perspective of the first period. When i is selected to make an offer to j in the first period, his conditional expected payoff is at least $1 - u_j^{ij}$. Indeed, player j must accept any offer greater than u_j^{ij} from i since j's expected payoff in the MPE conditional on not reaching an agreement with i in the first period is u_j^{ij} (regardless of i's rejected offer). Furthermore, player j can secure a payoff of u_j^{ij} in the event that i is the proposer by rejecting any offer that i makes. Hence the sum of equilibrium continuation payoffs of i and j conditional on i being the proposer in the match (i, j)—an event which is commonly known by i and j—is at least 1. The same conclusion holds if instead j is selected to make an offer to i. If v_i^{ij} and v_j^{ij} denote the expected payoffs of players i and j, respectively, in the MPE conditional on i being matched to j in the first period (before a proposer is recognized), it follows that $v_i^{ij} + v_j^{ij} \geq 1$. Then for any maximum matching H of G, summing the previous inequality over the links of H, we obtain $\sum_{ij \in H}(v_i^{ij} + v_j^{ij}) \geq |H| = m$, where m is the maximum total surplus in G. Since the sum of expected payoffs of all players in the MPE is at least as large as the average of the expression $\sum_{ij \in H}(v_i^{ij} + v_j^{ij})$ over all maximum matchings H, it must be that the total sum of MPE payoffs is greater than or equal to m. Clearly, the total sum of MPE payoffs cannot exceed m, so it must be exactly equal to m. As $\delta < 1$, this is possible only if all matches result in agreement.

We can now prove that there exists exactly one MPE. When players i and j are matched in the original network, they correctly anticipate that all other matches result in agreement. If i and j fail to reach an agreement then they are left in a subnetwork in which the maximum matching has only one link. Hence the remaining subnetwork would consist of either the single link ij, a star with multiple spokes, or a 3-player cycle (possibly with some isolated nodes). It is routine to verify that there exists a unique MPE in every such subnetwork. Players i and j share the same beliefs about the distribution of subnetworks that may arise in the second period. Since the sum of expected payoffs of i and j in any possible subnetwork cannot exceed 1, the sum of their discounted expected payoffs conditional on not reaching an agreement when matched with each other must be smaller than 1. Using the notation defined above, we have that $u_i^{ij} + u_j^{ij} < 1$. Then in any MPE, when i is selected to propose to j in the original network, i offers j a payoff of u_j^{ij} and j accepts the offer with probability 1. Therefore, there exists a unique MPE, in which every match materializes in an agreement in any subgame, as claimed.[27]

References

1. Abreu, D. and M. Manea (2012a) Markov equilibria in a model of bargaining in networks, *Games Econ. Behav.*, **75**, 1-16.
2. Abreu, D. and M. Manea (2012b) Bargaining and efficiency in networks, *J. Econ. Theory*, **147**, 43-70.
3. Binmore, K.G., A. Shaked and J. Sutton (1985) Testing noncooperative bargaining theory: A preliminary study, *Amer. Econ. Rev.*, **75**, 1178-1180.
4. Charness, G., M. Corominas-Bosch and G.R. Frechette (2007) Bargaining and network structure: An experiment, *J. Econ. Theory*, **136**, 28-65.
5. Corominas-Bosch, M. (1999) On two-sided network markets, Ph.D. Thesis, Universitat Pompeu Fabra, Barcelona.
6. Corominas-Bosch, M. (2004) Bargaining in a network of buyers and sellers, *J. Econ. Theory*, **115**, 35-77.
7. Crawford, V. and E. Knoer (1981) Job matching with heterogeneous firms and workers, *Econometrica*, **49**, 437-450.
8. Démange, G. (1982) Strategyproofness in the assignment market game, mimeo.
9. Démange, G., D. Gale and M. Sotomayor (1986) Multi-item auctions, *J. Pol. Econ.*, **94**, 863-872.
10. Edmonds, J. (1965) Paths, trees, and flowers, *Canadian J. Math.*, **17**, 449-467.
11. Elliott, M. (2015) Inefficiencies in networked markets, *Am. Econ. J: Micro.*, forthcoming.
12. Elliott, M. and F. Nava (2014) Decentralized bargaining: Efficiency and the core, mimeo.
13. Gale, D. (1986a) Bargaining and competition, part I: Characterization, *Econometrica*, **54**, 785-806.
14. Gale, D. (1986b) Bargaining and competition, part II: Existence, *Econometrica*, **54**, 807-818.
15. Gale, D. (1987) Limit theorems for markets with sequential bargaining, *J. Econ. Theory*, **43**, 20-54.
16. Gale, D. and L. Shapley (1962) College admissions and the stability of marriage, *Am. Math. Monthly*, **69**, 9-15.
17. Gallai, T. (1964) Maximale systeme unabhängiger kanten, *Magyar Tud. Akad. Mat. Kutató Int. Közl*, **9**, 401-413.
18. Gauer, F. (2014) Strategic formation of homogeneous bargaining networks, mimeo.
19. Gautier, P. and C. Holzner (2015) Maximum matching in the labor market under incomplete information, mimeo.
20. Hall, P. (1935) On representatives of subsets, *J. London Math. Soc.*, **10**, 26-30.
21. Jackson, M.O. and A. Wolinsky (1996) A strategic model of social and economic networks, *J. Econ. Theory*, **71**, 44-74.

[27]A reexamination of the argument shows that if the information structure is modified so that all players observe the entire matching activated by nature at each stage (but not the offers made in contemporaneous matches), then the resulting game also has a unique MPE, which generates the same expected payoffs as Polanski's benchmark model.

22. Kelso, A.S. and V.P. Crawford (1982) Job matching coalition formation and gross substitutes, *Econometrica*, **50**, 1483-1504.
23. Koopmans, T.C., and M. Beckman (1957) Assignment problems and the location of economic activities, *Econometrica*, **25**, 53-76.
24. Kranton, R. and D. Minehart (2000) Competition for goods in buyer-seller networks, *Rev. Econ. Design*, **5**, 301-331.
25. Kranton, R. and D. Minehart (2001) A theory of buyer-seller networks, *Amer. Econ. Rev.*, **91**, 485-508.
26. Lovász, L. and M.D. Plummer (2009) *Matching Theory*, AMS Chelsea Publishing, Providence, RI.
27. Leonard, H.B. (1983) Elicitation of honest preferences for the assignment of individuals to positions, *J. Pol. Econ.*, **91**, 461-479.
28. McLennan, A. and H. Sonnenschein (1991) Sequential bargaining as a noncooperative foundation of general equilibrium, *Econometrica*, **59**, 1395-1424.
29. Manea, M. (2011a) Bargaining in stationary networks, *Amer. Econ. Rev.*, **101**, 2042-2080.
30. Manea, M. (2011b) Bargaining in stationary networks: Online appendix, https://www.aeaweb.org/aer/data/aug2011/20090320_app.pdf.
31. Manea, M. (2014a) Bargaining in dynamic markets, mimeo.
32. Manea, M. (2014b) Steady states in matching and bargaining, mimeo.
33. Manea, M. (2014c) Intermediation in networks, mimeo.
34. Nguyen, T. (2014) Coalitional bargaining in networks, mimeo.
35. Polanski, A. (2007) Bilateral bargaining in networks, *J. Econ. Theory*, **134**, 557-565.
36. Polanski, A. and F. Vega-Redondo (2014) Markets, bargaining, and networks with heterogeneous agents, mimeo.
37. Roth, A. E. and M.A.O. Sotomayor (1990) *Two-Sided Matching: A Study in Game-Theoretic Modeling and Analysis*, Econometric Society Monographs, Cambridge University Press, Cambridge.
38. Rubinstein, A. (1982) Perfect equilibria of a bargaining model, *Econometrica*, **50**, 97-110.
39. Rubinstein, A. and A. Wolinsky (1985) Equilibrium in a market with sequential bargaining, *Econometrica*, **53**, 295-328.
40. Rubinstein, A. and A. Wolinsky (1990) Decentralized trading, strategic behaviour and the Walrasian outcome, *Rev. Econ. Stud.*, **57**, 63-78.
41. Shaked, A. and J. Sutton (1984) Involuntary unemployment as a perfect equilibrium in a bargaining model, *Econometrica*, **52**, 1351-1364.
42. Shapley, L.S. and M. Shubik (1971) The assignment game I: The core, *Int. J. Game Theory*, **1**, 111-130.
43. Stahl, I. (1972) *Bargaining Theory*, Stockholm School of Economics, Stockholm.