American Economic Review 2008, 98:3, 897–923
http://www.aeaweb.org/articles.php?doi=10.1257/aer.98.3.897

Stability in Supply Chain Networks

By Michael Ostrovsky*

This paper studies matching in vertical networks, generalizing the theory of matching in two-sided markets. It gives sufficient conditions for the existence of stable networks and presents an algorithm for finding two of them. One is the best stable network for the agents on the "upstream" end of an industry. The other is best for the agents on the "downstream" end. The paper describes several properties of the set of stable networks and discusses applications of the theory to the design of matching markets with more than two types of agents and to the empirical analysis of supply chains. (*JEL* C78, D85, L14)

> *The woollen coat, for example, which covers the day-labourer, as coarse and rough as it may appear, is the produce of the joint labour of a great multitude of workmen. The shepherd, the sorter of the wool, the wool-comber or carder, the dyer, the scribbler, the spinner, the weaver, the fuller, the dresser, with many others, must all join their different arts in order to complete even this homely production.*
>
> — Adam Smith (1776)

Following the work of David Gale and Lloyd S. Shapley (1962), matching in two-sided markets has been an active area of research in economics. It has produced a number of successful practical applications, as well as a variety of important theoretical results on the structure of matching markets and on the close links between matching and other areas of economics, such as auction theory and competitive equilibrium theory. This stream of research focuses on markets with two sides: e.g., the marriage market between men and women, the admissions market between colleges and students, or the labor market between firms and workers. Such a market is typically viewed in isolation and interactions with other markets are ignored.

In many situations, however, markets can be closely interconnected, and an agent's behavior in one of them may be directly linked to his behavior and options in another. The preferences of a consulting firm hiring college graduates depend on the number and types of its clients. The menu that a restaurant offers to its customers depends on the availability and prices of inputs from its suppliers. The amount of iron ore that a steel manufacturer wants to buy depends on the amount of steel it plans to sell.

Interconnected markets can often be handled by the standard competitive equilibrium approach. That approach, however, is ill-suited for incorporating the discreteness and the high degree of heterogeneity of qualities, preferences, and contracts inherent in many settings. For example, on a macro level, we can talk about the average wage of management consultants and derive predictions about its movements in response to various shocks in the economy. Once, however, we look at that labor market more closely, it becomes clear that contracts are unique and personalized, assigning a specific worker to a specific firm, location, and position at a specific salary. These

* Graduate School of Business, Stanford University, Stanford, CA 94305 (e-mail: ostrovsky@gsb.stanford.edu). I am especially grateful to Al Roth and Ariel Pakes for advice and encouragement throughout this project and to Drew Fudenberg, Parag Pathak, and Michael Schwarz for detailed comments on an early draft of the paper. I also thank Attila Ambrus, John Asker, Jeremy Bulow, Jerry Green, Kate Ho, Paul Milgrom, Markus Mobius, John Roberts, Tayfun Sönmez, Bob Wilson, Pai-Ling Yin, the coeditor, and anonymous referees for helpful comments and suggestions.

features are easier to incorporate in a matching framework. The same is true of the restaurant market, the markets for iron ore and steel, and other differentiated markets. Section V talks about several such markets in more detail.

This paper generalizes the results and techniques of two-sided matching literature to a particular kind of setting with interconnected markets: supply chain networks. The basic structure of this setting is as follows. An industry includes a number of agents: workers, producers, distributors, retailers, and so on. Some agents supply basic inputs for the industry and do not consume any of the outputs (e.g., wheat farmers are the suppliers of basic inputs in a "farmer → miller → baker → retailer" supply chain). Some agents purchase the final outputs of the industry (e.g., car manufacturers are the consumers of final goods in an "iron ore supplier → steel producer → steel consumer" supply chain). The rest are intermediaries, who get their inputs from some agents in the industry, convert them into outputs at a cost, and sell the outputs to some other agents (e.g., millers, bakers, and steel producers are intermediate agents in the examples above).

A key assumption in the two-sided matching literature is that the side to which a particular agent belongs does not depend on market conditions: a man cannot become a woman and a student cannot become a college. Likewise, I assume that the positions of agents in a supply chain are exogenously fixed. Specifically, there is a predetermined upstream-downstream partial ordering on the set of agents: for a pair of agents A and B, either A is a potential direct or indirect supplier for B, or B is a potential direct or indirect supplier for A, but not both; it may also be the case that neither is a potential supplier for the other. This ordering can be very simple, with several "tiers" of firms, each of which buys its inputs from firms in the previous tier and sells its outputs to firms in the next tier. It can also be more complicated, with several alternative technologies, several paths of different lengths connecting the suppliers of basic inputs to the consumers of final outputs, and with firms being able to trade both directly and through intermediaries. However, it cannot have cycles: an agent cannot be a direct or an indirect supplier for one of its upstream nodes. Note that if there are no intermediate agents, this setting reduces to a two-sided market.

Agents can trade discrete quantities of goods, with the smallest tradeable quantity (the *unit* of quantity) defined ex ante. For example, one unit may correspond to one million tons of steel, one hour of work, or one loaf of bread. In the Gale-Shapley two-sided marriage market, one unit corresponds to marriage, and each person can "trade" at most one unit. Following the literature on two-sided matching with wages and endogenous job characteristics (Vincent P. Crawford and Elsie Marie Knoer 1981; Alexander S. Kelso, Jr., and Crawford 1982; Alvin E. Roth 1984; John W. Hatfield and Paul Milgrom 2005), units traded in the market are represented by *contracts*. In my setting, each contract specifies the buyer, the seller, the price (if monetary transfers are involved), and the serial number of the traded unit (if multiple units can be traded). A *network* is a set of contracts: it specifies who sells what to whom and at what price. Each agent has preferences over sets of contracts in which it is involved as a buyer or a seller: e.g., an intermediate agent's payoff from such a set depends on the payments it makes for its inputs (specified in its *upstream* contracts) and receives for its outputs (specified in its *downstream* contracts), as well as on the cost of converting the inputs into the outputs. For a consumer of final goods, the payoff depends on the utility from the goods it purchases and the payments it makes for these goods.

I say that a network is *chain stable* if there is no upstream–downstream sequence of agents (not necessarily going all the way to the suppliers of basic inputs and the consumers of final outputs) who could become better off by forming new contracts among themselves and possibly dropping some of their current contracts. This condition is parallel to pairwise stability in two-sided markets, and is tautologically equivalent to it if there are no intermediate agents in the industry. Note that the concept of stability in networks is not strategic—I do not study the dynamics of network formation or "what-if" scenarios analyzed by agents who may be considering temporarily dropping or adding contracts in the hopes of affecting the entire network in a way beneficial to them, although these considerations are undoubtedly important in many settings. The concept is closer in spirit to general equilibrium models, where agents perceive conditions surrounding them as given, and optimize given those conditions. Under chain stability, agents also perceive conditions surrounding them as given (i.e., which other agents are willing to form contracts with them, and what those contracts are), and optimize given these conditions.

Without restrictions on preferences, the set of stable matchings may be empty even in the two-sided one-to-many setting. A commonly used restriction that is sufficient to guarantee the existence of stable matchings in that setting is the gross substitutes condition of Kelso and Crawford (1982). In the supply chain setting, I place an analogous pair of restrictions on preferences; these restrictions are equivalent to the substitutes condition if there are no intermediate agents in the industry. The restrictions are *same-side substitutability* and *cross-side complementarity*.

Same-side substitutability is a direct generalization of the gross substitutes condition. It says that when the set of available downstream contracts of a firm expands (i.e., there are more potential customers, or the potential customers' willingness to pay goes up), while the set of available upstream contracts remains unchanged, the set of downstream contracts that the firm rejects also (weakly) expands. Symmetrically, when the set of available upstream contracts expands and the set of available downstream contracts remains unchanged, the set of rejected upstream contracts also expands.

Cross-side complementarity is a new restriction, which specifies how a firm's purchasing and selling decisions are interrelated, and thus links the markets along the supply chain. This restriction can be viewed as a mirror image of same-side substitutability. It says that when the set of available downstream contracts of a firm expands, while the set of available upstream contracts remains unchanged, the set of upstream contracts that the firm forms also weakly expands. Symmetrically, when the set of available upstream contracts expands and the set of available downstream contracts remains unchanged, the set of downstream contracts that the firm forms also weakly expands.

The restrictions of same-side substitutability and cross-side complementarity are natural and flexible; however, there are several important types of preferences and production technologies that they rule out. I discuss these restrictions in more detail in Section IA.

The main result of the paper states that under same-side substitutability and cross-side complementarity, there exists at least one chain-stable network. The proof is constructive: it presents an algorithm for finding such a network. This algorithm is a generalization of the fixed-point algorithms of Hiroyuki Adachi (2000), Federico Echenique and Jorge Oviedo (2004, 2006), and Hatfield and Milgrom (2005), which are in turn generalizations of the Deferred Acceptance Algorithm of Gale and Shapley (1962) and the Salary Adjustment Process of Crawford and Knoer (1981) and are applicable only to two-sided markets.

The set of chain-stable networks has some interesting properties. The chain-stable network formed in the constructive proof of the existence theorem is upstream-optimal: it is the *most* preferred chain-stable network for *all* suppliers of basic inputs and the *least* preferred chain-stable network for *all* consumers of final outputs. A slightly modified algorithm produces the downstream-optimal chain-stable network. The existence of these side-optimal networks is a generalization of the classic result in the theory of two-sided matching, which says that the Deferred Acceptance Algorithm converges to the man-optimal stable matching (Gale and Shapley 1962). Another property of the set of stable matchings in two-sided markets is that adding agents on one side of the market makes other agents on that side of the market weakly worse off and makes the agents on the other side weakly better off, in the two side-optimal stable matchings (Kelso and Crawford 1982; Gale and Marilda Sotomayor 1985). This result can also be extended to the supply chain setting: adding a new supplier of basic inputs to the industry makes other such suppliers weakly worse off and makes the consumers of final outputs weakly better off at both upstream- and downstream-optimal chain-stable networks. Symmetrically, adding a new consumer of final outputs makes other such consumers worse off and makes the suppliers of basic inputs better off.

These results are very different from the conclusions of several other papers that discuss generalizations and extensions of two-sided matching, but do not impose the supply chain structure of this paper. Gale and Shapley (1962) show by example that the "problem of the roommates," whose only difference from the "marriage problem" is the absence of two sides in the market, may fail to have a stable pairing. Ahmet Alkan (1988) shows that the "man–woman–child marriage problem," in which each match consists of agents of three different types, may also fail to have a stable matching. Hernan Abeledo and Garth Isaak (1991) prove that to guarantee the existence of stable pairings under arbitrary preferences, it has to be the case that each agent belongs to one of two classes, and an agent in one class can match only with agents in the other class. Note that there is no contradiction between my results and the result of Abeledo and Isaak: in their setting, matching is one-to-one, and so for a particular agent, all acceptable matches are substitutes for one another. In contrast, in the setting of this paper, some agents are intermediaries whose preferences satisfy cross-side complementarity and who therefore do not view all acceptable matches as substitutes. Finally, Hideo Konishi and M. Utku Ünver (2006, Theorem 3) show that in general multi-partner matching problems with bilateral links and responsive preferences, several different solution concepts are equivalent to pairwise stability. In their setting, however, the set of stable matchings may be empty.

To guarantee the existence of chain-stable networks, I assume that agents' preferences satisfy the cross-side complementarity and same-side substitutability conditions. A similar restriction on preferences was independently introduced by Ning Sun and Zaifu Yang (2006) in the context of exchange economies with two types of heterogeneous indivisible objects, quasi-linear preferences, and continuous prices. From the point of view of each agent, objects of one type are substitutes, while objects of different types are complements. For example, from the point of view of a manufacturing firm, workers and machines are complements, while all workers are substitutes for each other and all machines are also substitutes. Sun and Yang (2006) show that in this model a competitive equilibrium always exists. I discuss the connection between this result and the results of the current paper in Section VC.

The rest of this paper is organized as follows. Section I formally introduces the model of matching in supply chains. Section II presents the proof of the main result of the paper: under same-side substitutability and cross-side complementarity, chain-stable networks are guaranteed to exist. Section III studies the properties of the set of chain-stable networks. Section IV discusses chain stability and alternative solution concepts. Section V talks about applications and extensions of the model. Section VI concludes.

I. The Model of Matching in Supply Chains

Consider an industry consisting of a finite set A of nodes (firms, countries, agents, workers, and so on). An exogenously given "upstream–downstream" partial ordering "$>$" on this set determines possible trading relationships: $a > b$ stands for b being a downstream node for a and means that a can potentially sell something to b. If $a \not> b$ and $b \not> a$, then there can be no relationship between a and b. By transitivity, there are no loops in the market.

Relationships between pairs of nodes are represented by bilateral "contracts." Each contract $\mathbf{c}$ represents one unit of a good sold by one node to another. It is a vector, $\mathbf{c} = (s, b, l, p)$, where $s \in A$ and $b \in A$ are the "seller" and the "buyer" involved in the contract, $s > b$; $l \in \mathbb{N}$ is the "serial number" of the unit of the good represented by the contract; and $p \in \mathbb{R}$ is the price that the buyer pays to the seller for that unit. The seller involved in contract $\mathbf{c}$ is denoted by $s_{\mathbf{c}}$, the buyer is denoted by $b_{\mathbf{c}}$, and so on.

Multiple contracts between a seller and a buyer can represent multiple units of the same good or service, units of different types of goods or services, or both. For instance, if the unit is one ton, and a farmer sells 5 tons of wheat and 10 tons of rye to a miller, then this relationship will be represented by 15 contracts with 15 different serial numbers. The example in Appendix B illustrates the use of serial numbers.

The set of possible contracts, C, is finite and is also given exogenously. In the simplest case, it can include all possible contracts between nodes in A, with all possible serial numbers from some finite set, and all possible prices from some finite set. It can also be more complicated: e.g., the US trade embargo on Cuba can be incorporated simply by removing all contracts between the nodes in these countries from set C.

Note that this framework, restricted to one "tier" of sellers and one "tier" of buyers, encompasses various two-sided matching settings considered in the literature. Setting $l \equiv constant$ and $p \equiv 0$ turns this model into the marriage model of Gale and Shapley (1962), assuming that each agent is allowed to have at most one partner. Setting $l \equiv constant$ turns it into the setup of Kelso and Crawford (1982), assuming that agents on one side are restricted to having at most one link, and into the many-to-many matching model of Roth (1984, 1985) and Charles Blair (1988), assuming that agents on both sides are allowed to have multiple links. Setting $p \equiv 0$ and assuming that all links connecting two nodes are identical turns the model into a discrete version of the schedule matching problem of Mourad Baïou and Michel Balinski (2002) and Alkan and Gale (2003).

A. *Preferences*

Each node can be involved in several contracts, some as a seller, some as a buyer, but it cannot be involved in two contracts that differ only in price p, i.e., it cannot buy or sell the same unit twice. Nodes have preferences over sets of contracts that involve them as the buyer or the seller. For example, in the simplest case of quasi-linear utilities and profits, the utility of node a involved in a set of contracts X is

$$V_a(X) = W_a(\{(s_{\mathbf{c}}, b_{\mathbf{c}}, l_{\mathbf{c}}) | \mathbf{c} \in X\}) + \sum_{\mathbf{c} \in D} p_{\mathbf{c}} - \sum_{\mathbf{c} \in U} p_{\mathbf{c}},$$

where $D = \{\mathbf{c} \in X | a = s_{\mathbf{c}}\}$ and $U = \{\mathbf{c} \in X | a = b_{\mathbf{c}}\}$, i.e., D is the set of contracts in X in which a is involved as a seller, and U is the set of contracts in which a is involved as a buyer. $W_a(\cdot)$ represents the utility from the purchased contracts for the consumers at the downstream end of the chain, the cost of producing the sold contracts for the suppliers at the upstream end of the chain, and the cost of converting inputs into outputs for the intermediate nodes.

For an agent $a \in A$ and a set of contracts X, let $Ch_a(X)$ be a's most preferred (possibly empty) subset of X, let $U_a(X)$ be the set of contracts in X in which a is the buyer (i.e., upstream contracts), and let $D_a(X)$ be the set of contracts in X in which a is the seller (i.e., downstream contracts). Subscript a will be omitted when it is clear from the context which agent's preferences are being considered. Preferences are strict, i.e., function $Ch_a(X)$ is single-valued. In the settings in which it is natural to assume that several different sets of contracts should result in identical payoffs (e.g., when two nodes can trade several identical units of a good), I assume that ties are broken in a consistent manner, e.g., lexicographically: in the case of several identical units of a good, that would imply that seller a prefers contract $(a, b, 1, p)$ to contract $(a, b, 2, p)$, but would prefer $(a, b, 2, p')$ to $(a, b, 1, p)$ for any $p' > p$.

Preferences of agent a are *same-side substitutable* if for any two sets of contracts X and Y such that $D(X) = D(Y)$ and $U(X) \subset U(Y)$, $U(X) \setminus U(Ch(X)) \subset U(Y) \setminus U(Ch(Y))$, and for any two sets X and Y such that $U(X) = U(Y)$ and $D(X) \subset D(Y)$, $D(X) \setminus D(Ch(X)) \subset D(Y) \setminus D(Ch(Y))$. That is, preferences are same-side substitutable if, choosing from a bigger set of contracts on one side, the agent does not accept any contracts *on that side* that he rejected when he was choosing from the smaller set.

Preferences of agent a are *cross-side complementary* if for any two sets of contracts X and Y such that $D(X) = D(Y)$ and $U(X) \subset U(Y)$, $D(Ch(X)) \subset D(Ch(Y))$, and for any two sets X and Y such that $D(X) \subset D(Y)$ and $U(X) = U(Y)$, $U(Ch(X)) \subset U(Ch(Y))$. That is, preferences are cross-side complementary if, when presented with a bigger set of contracts on one side, an agent does not reject any contract *on the other side* that he accepted before.

Same-side substitutability is a generalization of the gross substitutes condition introduced by Kelso and Crawford (1982) and used widely in the matching literature. If there are only two sides in the supply chain market, then these two conditions become tautologically equivalent. Cross-side complementarity is a new restriction, which is automatically satisfied in any two-sided market. In a supply chain setting, it is this restriction that "ties together" the purchasing and selling decisions of a node and thus links the markets along the supply chain. It can be viewed as a mirror image of same-side substitutability: when the set of potential contracts on the node's one side expands, same-side substitutability says that the set of *rejected* contracts on *that* side also expands, while cross-side complementarity says that the set of *accepted* contracts on *the other* side expands.

It is important to highlight what is allowed and what is not allowed by this pair of assumptions. One possibility that they rule out is scale economies and production functions with fixed costs, because in those cases a firm may decide not to produce one unit of a good at a certain price, while being willing to produce ten units at the same price, violating same-side substitutability. Complementary inputs (or outputs) are also generally ruled out.[1] Another possibility that is ruled out is an intermediary with fixed capacity (say, one unit) who can transform an input of type 1 into an output of type 1 or an input of type 2 into an output of type 2, but not both (due to the capacity constraint). An addition of a cheap type-1 input to this intermediary's set of options may cause him to shift from buying one type-2 input and selling one type-2 output to buying one type-1 input and selling one type-1 output, thus violating cross-side complementarity.

In contrast, with substitutable inputs and outputs and decreasing returns to scale, many production and utility functions can be accommodated. The simplest example is a firm that can take one kind of input and produce one kind of output at a cost, with the marginal cost of production increasing or staying constant in quantity. The input good can come from several different nodes, and the output good may go to several different nodes, with different transportation costs. Much more general cases are possible as well: preferences and production functions with quotas and tariffs, several different inputs and outputs with discrete choice demands and production functions, capacity constraints and increasing transportation costs, etc.

The interdependencies between different inputs or outputs can be rather complex as well. Consider the following example. A firm has two plants in the same location. Each plant's capacity is equal to one unit. The first plant can convert one unit of iron ore into one unit of steel for q_o^1 or it can convert one unit of steel scrap into one unit of steel for q_s^1. The second plant can convert one unit of iron ore into one unit of steel for q_o^2 or it can convert one unit of steel scrap into one unit of steel for q_s^2. Then, for a generic choice of costs and prices, the preferences of this firm will be same-side substitutable and cross-side complementary, even though the firm's

[1] See, however, the discussion in Section VC.

preferences over iron ore and scrap are not trivial. (They cannot be expressed by simply saying that two alternative inputs are perfect substitutes, with one being better than the other by a certain amount x.)

This example is an analogue of "endowed assignment valuations" in two-sided matching markets (Shapley 1962; Hatfield and Milgrom 2005), in which each firm has several unit-capacity jobs, each worker has a certain productivity at each job, and each firm has an initial endowment of workers. Even in the two-sided setting, it is an open question whether endowed assignment valuations exhaust the set of utility functions with substitutable preferences, and so it is an open question in the supply chain setting as well.

For the remainder of this paper, all preferences are assumed to be same-side substitutable and cross-side complementary, and these restrictions will often be omitted from the statements of results to avoid repetition.

B. *Stable Networks*

In the model, relationships between the nodes are represented by contracts. I call a collection of relationships between the nodes in a market a *network*, i.e., a network is simply a set of contracts. Let $\mu(a)$ denote the set of contracts involving a in network μ. Network μ is *individually rational* if, for any agent a, $Ch_a(\mu(a)) = \mu(a)$, i.e., no agent would like to unilaterally drop any of his contracts.

The most widely used solution concept in the two-sided matching literature is pairwise stability. Its analogue in the supply chain setting is chain stability, defined as follows. A *chain* is a sequence of contracts, $\{\mathbf{c}_1, \ldots, \mathbf{c}_n\}$, $n \geq 1$, such that for any $i < n$, $b_{\mathbf{c}_i} = s_{\mathbf{c}_{i+1}}$, i.e., the buyer in contract $\mathbf{c}_i$ is the same node as the seller in contract $\mathbf{c}_{i+1}$. Note that the chain does not have to go all the way from one of the most upstream nodes in the market to one of the most downstream nodes; it can connect several nodes in the middle of the market. For notational convenience, let $b_i \equiv b_{\mathbf{c}_i}$ and $s_i \equiv s_{\mathbf{c}_i}$. For a network μ, a *chain block* is a chain $\{\mathbf{c}_1, \ldots, \mathbf{c}_n\}$ such that

- $\forall i \leq n$, $\mathbf{c}_i \notin \mu$,
- $\mathbf{c}_1 \in Ch_{s_1}(\mu(s_1) \cup \mathbf{c}_1)$,
- $\mathbf{c}_n \in Ch_{b_n}(\mu(b_n) \cup \mathbf{c}_n)$, and
- $\forall i < n$, $\{\mathbf{c}_i, \mathbf{c}_{i+1}\} \subset Ch_{b_i = s_{i+1}}(\mu(b_i) \cup \mathbf{c}_i \cup \mathbf{c}_{i+1})$.

In other words, a chain block of network μ is a downstream sequence of contracts not belonging to μ, in which the buyer in one contract is the seller in the next one, such that each node involved in these contracts is willing to add all of its contracts in the sequence to its contracts in μ, possibly dropping some of its contracts in μ. A network is *chain stable* if it is individually rational and has no chain blocks.

Note that chain stability is not a strategic concept—each node views the set of contracts available to it as exogenously given and maximizes its payoff given that set, analogously to how consumers in a competitive equilibrium choose quantities taking prices as given. Hence, each node treats its contracts independently of one another, ignoring the effect of forming one contract on other nodes' willingness to pay for other contracts. For example, if a market consists of one seller (whose marginal cost is increasing in quantity) and one buyer (whose marginal benefit is decreasing in quantity), then chain stability implies that the number of units traded between these two agents is determined by the intersection of their marginal cost and marginal benefit curves. In more general networks, nodes also ignore various sorts of externalities they may impose on others (e.g., limiting the supply of inputs available to competitors by buying too much and thus reducing the competition in the market for outputs). Hence, the model is not directly applicable to cases in which there are several large players manipulating the market; it is better suited to describing competitive markets with many small players, or markets in which nodes represent countries or regions rather than firms.

C. *An Example*

I now turn to an example that illustrates the definitions introduced above. It consists of two simple two-sided matching markets "stacked" on top of each other. The markets are tied together by the restriction on the preferences of intermediate agents, who need to have a supplier in order to produce output for a customer. For simplicity, there are no prices or unit identifiers in this example. Another example, presented in Appendix B, incorporates prices, unit identifiers, and more complicated trade patterns, allowing the suppliers of basic inputs and the consumers of final outputs to trade both directly and through intermediaries.

Example 1: There are six agents in the market: two suppliers of basic inputs (a_1, a_2), two intermediaries (b_1, b_2), and two consumers of final outputs (c_1, c_2). For all i, j, and k, $a_i \succ b_j \succ c_k$.

Suppliers cannot trade directly with consumers: trade flows have to go through intermediaries. All agents have unit capacities: each supplier can supply one unit of the good; each consumer needs one unit; each intermediary can process one unit. There are no prices or quantities in this market. The set of available contracts consists of eight elements: $C = \{(a_1, b_1), (a_1, b_2), (a_2, b_1), (a_2, b_2), (b_1, c_1), (b_1, c_2), (b_2, c_1), (b_2, c_2)\}$. Prices and unit identifiers were dropped from the contracts since they play no role in this example.

Each supplier is willing to sell to any intermediary and prefers to sell to the intermediary with the same index as his own rather than to the intermediary with the other index. Hence, the most preferred set of contracts for supplier a_1 is $\{(a_1, b_1)\}$, the second most preferred is $\{(a_1, b_2)\}$, and the third is the empty set. Similarly, a_2 prefers $\{(a_2, b_2)\}$ to $\{(a_2, b_1)\}$, which he in turn prefers to the empty set.

Each consumer is willing to buy from any intermediary and prefers to buy from the intermediary with a different index rather than from the intermediary with the same index as her own. Hence, the most preferred set of contracts for consumer c_1 is $\{(b_2, c_1)\}$, the second most preferred is $\{(b_1, c_1)\}$, and the third is the empty set. Similarly, c_2 prefers $\{(b_1, c_2)\}$ to $\{(b_2, c_2)\}$, which she in turn prefers to the empty set.

Finally, an intermediary wants to trade with a consumer if, and only if, he also trades with a supplier, and vice versa. He prefers to sell to the consumer with the same index as his own, but prefers to buy from the supplier with a different index. Hence, the most preferred set of contracts for intermediary b_1 is $\{(a_2, b_1), (b_1, c_1)\}$, the next two most preferred sets (in any order) are $\{(a_1, b_1), (b_1, c_1)\}$ and $\{(a_2, b_1), (b_1, c_2)\}$, then $\{(a_1, b_1), (b_1, c_2)\}$, and, finally, the empty set. The sets containing only one element are worse for b_1 than the empty set. The preference ordering for intermediary b_2 is defined analogously. It is easy to check that for the preferences of all six nodes, same-side substitutability and cross-side complementarity hold.

Consider now the networks in Figure 1. Contracts that are in the networks are represented by solid lines. Network μ_a in Figure 1A is empty and is not chain stable: it is blocked, for instance, by the chain of contracts $\{(a_1, b_1), (b_1, c_1)\}$, shown in the figure by dashed lines. Network μ_b, consisting of contracts (a_2, b_2) and (b_2, c_1), is also unstable: it is blocked by the single-contract chain $\{(b_2, c_2)\}$. Network μ_c consists of one contract, (a_1, b_1), and is also unstable: among other things, it is not individually rational, because node b_1 would be better off dropping contract (a_1, b_1). Finally, network μ_d is chain stable: it is individually rational and is not blocked by any chain.

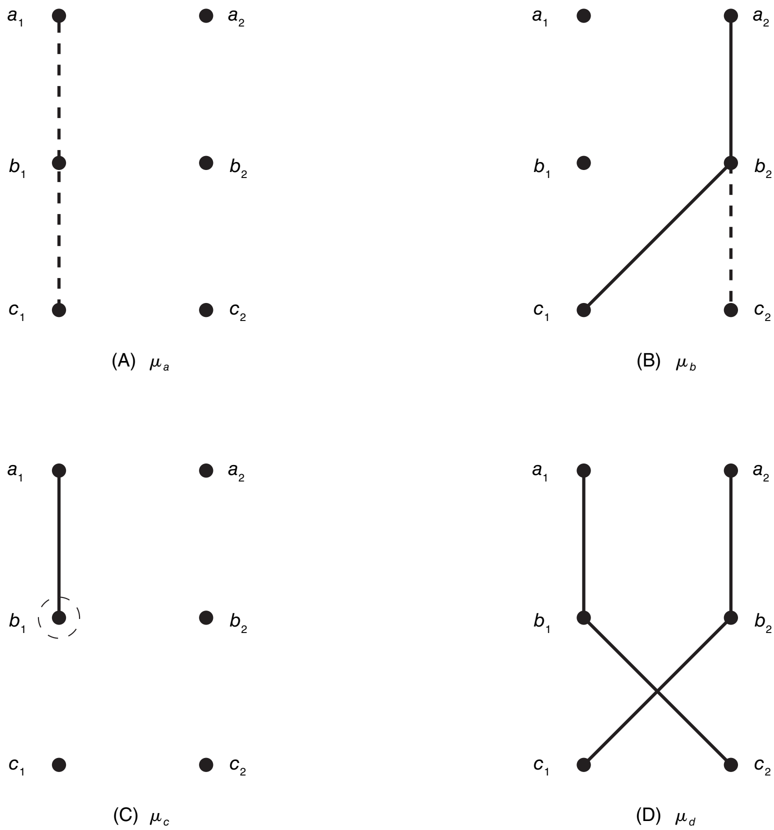

FIGURE 1. UNSTABLE AND STABLE NETWORKS

II. Existence of Stable Networks

This section presents the main result of the paper: under the assumptions of Section I, chain-stable networks are guaranteed to exist. The proof is constructive. It provides an algorithm for finding a special chain-stable network, whose properties I will discuss in Section III. The algorithm generalizes the fixed-point algorithms developed by Adachi (2000), Echenique and Oviedo (2004, 2006), and Hatfield and Milgrom (2005) in the context of two-sided markets. These algorithms, in turn, are descendants of the Deferred Acceptance Algorithm of Gale and Shapley (1962) and the Salary Adjustment Process of Crawford and Knoer (1981).

The proof relies on objects called pre-networks and on a special mapping, T, defined on the set of pre-networks. The definitions are as follows.

A *pre-network* is a set of arrows from nodes in A to other nodes in A, with the following properties. Each arrow $\mathbf{r}$ is a vector $(o_{\mathbf{r}}, d_{\mathbf{r}}, \mathbf{c}_{\mathbf{r}})$, where $o_{\mathbf{r}}$ ("origin of arrow $\mathbf{r}$") and $d_{\mathbf{r}}$ ("destination of arrow $\mathbf{r}$") are two different nodes and $\mathbf{c}_{\mathbf{r}}$ ("contract attached to arrow $\mathbf{r}$") is a contract involving both $o_{\mathbf{r}}$ and $d_{\mathbf{r}}$. If $o_{\mathbf{r}}$ is the seller and $d_{\mathbf{r}}$ is the buyer of contract $\mathbf{c}_{\mathbf{r}}$, the arrow is "downstream." Otherwise, $o_{\mathbf{r}}$ is the buyer and $d_{\mathbf{r}}$ is the seller of contract $\mathbf{c}_{\mathbf{r}}$, and the arrow is "upstream." For a pre-network ν and a node a, $\nu(a)$ denotes the set of contracts attached to arrows pointing <u>to</u> a, i.e., $\nu(a) = \{\mathbf{c} | \mathbf{r} = (o_{\mathbf{r}}, a, \mathbf{c}) \in \nu\}$.

There can be multiple arrows going from $o_{\mathbf{r}}$ to $d_{\mathbf{r}}$, but any two arrows going from $o_{\mathbf{r}}$ to $d_{\mathbf{r}}$ must have different contracts attached to them (these contracts may differ in serial numbers, prices, or both). Arrows going in opposite directions (from node a to node b and from node b to node a) can have identical contracts attached to them. Hence, there are two arrows corresponding to each contract: one from the seller of the contract to the buyer, and one from the buyer to the seller. Let R denote the set of all arrows.

Mapping T on the set of pre-networks is defined as follows. Consider any pre-network ν. Then pre-network $T(\nu)$ consists of such arrows $\mathbf{r}$ that the contract attached to the arrow, $\mathbf{c}_{\mathbf{r}}$, belongs to the most preferred subset chosen by the arrow's origin, $o_{\mathbf{r}}$, from set $(\nu(o_{\mathbf{r}}) \cup \mathbf{c}_{\mathbf{r}})$, i.e.,

$$T(\nu) = \{\mathbf{r} \in R | \mathbf{c}_{\mathbf{r}} \in Ch_{o_{\mathbf{r}}}(\nu(o_{\mathbf{r}}) \cup \mathbf{c}_{\mathbf{r}})\}.$$

Pre-network ν^* is a *fixed point* of mapping T if $T(\nu^*) = \nu^*$.

The first step of the proof is Lemma 1 below. It shows that fixed points of mapping T play a very special role: there is a one-to-one correspondence between the set of chain-stable networks and the set of fixed points of mapping T. Thus, to prove the main result, it will be sufficient to construct a fixed-point pre-network.

Formally, define mapping F from the set of pre-networks to the set of networks as follows. Take any pre-network ν, and consider any contract $\mathbf{c}$. Contract $\mathbf{c}$ belongs to $\mu = F(\nu)$ if, and only if, both the arrow from the seller of $\mathbf{c}$ to the buyer of $\mathbf{c}$ with $\mathbf{c}$ attached and the arrow from the buyer of $\mathbf{c}$ to the seller of $\mathbf{c}$ with $\mathbf{c}$ attached are contained in ν, i.e.,

$$F(\nu) = \{\mathbf{c} \in C | (s_{\mathbf{c}}, b_{\mathbf{c}}, \mathbf{c}) \in \nu \text{ and } (b_{\mathbf{c}}, s_{\mathbf{c}}, \mathbf{c}) \in \nu\}.$$

LEMMA 1: *For any pre-network ν^* such that $T(\nu^*) = \nu^*$, network $\mu^* = F(\nu^*)$ is chain stable. Moreover, for any chain-stable network μ^*, there exists exactly one fixed-point pre-network ν^* such that $\mu^* = F(\nu^*)$.*

Only the first statement of Lemma 1 is used to prove the existence of chain-stable networks. The second statement comes into play in Section III, where I use the one-to-one correspondence between the set of fixed points of mapping T and the set of chain-stable networks to establish several properties of the latter.

To see the intuition behind Lemma 1, note that for any pre-network ν, in the pre-network $T(\nu)$, each node a "points" to the nodes with contracts that it would like to form, if it assumes that it can also choose from all contracts attached to the arrows pointing to it in ν. Hence, in a fixed-point pre-network $\nu^* = T(\nu^*)$, each node a "points" to the contracts that it would like to form, if it assumes that it can also choose from all the contracts pointing to it. Mapping F simply takes "mutually acceptable" contracts (i.e., attached to arrows pointing from both the seller to the buyer and from the buyer to the seller), puts them in the network $\mu^* = F(\nu^*)$, and throws out the contracts that are acceptable to only one or none of the nodes. Network μ^* is individually rational, because only "acceptable" contracts were kept in the network. Also, it is not blocked by any chain of contracts: in any such chain, the buyers and sellers of all contracts would have to be willing to form these contracts; this would imply that ν^* contains all arrows going along this chain, in both directions; and this in turn would violate the assumption that these contracts formed a chain block and thus did not belong to network $\mu^* = F(\nu^*)$. To prove that for any chain-stable network μ^*, there exists exactly one fixed-point pre-network ν^* such that $\mu^* = F(\nu^*)$, I construct an iterative algorithm that for any chain-stable network μ^* finds a fixed-point pre-network ν^* such that $F(\nu^*) = \mu^*$. All formal proofs are in Appendix A.

To prove that chain-stable networks exist, it is now sufficient to show that mapping T has a fixed point. To find a fixed point, I first introduce a partial ordering on the set of pre-networks, as follows. Let ν_1 and ν_2 be two pre-networks. Then, ν_1 is said to be less than or equal to ν_2 ($\nu_1 \leq \nu_2$) if the set of downstream arrows in ν_1 is a *subset* of the set of downstream arrows in ν_2, while the set of upstream arrows in ν_1 is a *superset* of the set of upstream arrows in ν_2. Let $\nu_{\min}$ be the pre-network that includes all possible upstream arrows and no downstream arrows, and let $\nu_{\max}$ be the pre-network that includes no upstream arrows and all possible downstream arrows. By construction, for any pre-network ν, $\nu_{\min} \leq \nu \leq \nu_{\max}$.

The following key lemma shows that mapping T is isotone, i.e., order-preserving.

LEMMA 2: *For any pair of pre-networks ν_1 and ν_2 such that $\nu_1 \leq \nu_2$, we have $T(\nu_1) \leq T(\nu_2)$.*

Intuitively, if $\nu_1 \leq \nu_2$, then in ν_1 each node a has fewer "options" (i.e., contracts attached to arrows pointing to it) from potential sellers and more "options" from potential buyers than in ν_2. Then, by same-side substitutability and cross-side complementarity, given its options in ν_1, node a will be less "picky" in forming contracts with potential sellers and more "picky" in forming contracts with potential buyers than it would be given its options in ν_2. Hence, for any node a there will be more arrows pointing upstream from a (and fewer arrows pointing downstream) in $T(\nu_1)$ than in $T(\nu_2)$.

It is now easy to construct an algorithm for finding a chain-stable network. Namely, take the smallest pre-network, $\nu_{\min}$. Apply mapping T to it, getting $T(\nu_{\min})$. Since $\nu_{\min}$ is the smallest pre-network, by definition, $\nu_{\min} \leq T(\nu_{\min})$. Now, apply mapping T to $T(\nu_{\min})$. By Lemma 2, $T(\nu_{\min}) \leq T(T(\nu_{\min})) = T^2(\nu_{\min})$. Applying mapping T repeatedly, we get an increasing sequence of pre-networks: $\{\nu_{\min}, T(\nu_{\min}), T^2(\nu_{\min}), T^3(\nu_{\min}), \ldots\}$. Since the set of all pre-networks is finite, after a finite number of steps this sequence has to converge to a fixed point $\nu^*_{\min}$. We can now apply mapping F to this fixed-point pre-network to get a chain-stable network $\mu^*_{\min}$. This completes the description of the algorithm, and also proves Theorem 1, stated below.

THEOREM 1: *There exists a chain-stable network.*

One can also use pre-network $\nu_{\max}$ instead of $\nu_{\min}$ as the starting point of the algorithm above. Then, the algorithm would converge to a possibly different fixed point, $\nu^*_{\max}$, resulting in a possibly different chain-stable network, $\mu^*_{\max}$. In Section III, I describe some special properties of fixed-point pre-networks $\nu^*_{\min}$ and $\nu^*_{\max}$ and the corresponding chain-stable networks $\mu^*_{\min}$ and $\mu^*_{\max}$. Before moving on to the properties of these networks, however, I illustrate the algorithm presented above using the setting of Example 1.

A. *T-Algorithm: An Example*

Recall the setting of Example 1, with two suppliers, two intermediaries, and two consumers. In this example, for any two nodes that can trade, there can be only one contract (since there are no prices or unit identifiers), and so in discussing pre-networks, we can simply talk about the arrow $\overrightarrow{xy}$ going from node x to node y. This market has 16 different arrows: four downstream arrows from suppliers to intermediaries, four downstream arrows from intermediaries to consumers, four upstream arrows from intermediaries to suppliers, and four upstream arrows from consumers to intermediaries. Pre-network $\nu_{\min}$, which contains no downstream arrows and all eight upstream arrows, is shown in Figure 2A.

To find a chain-stable network, we first need to iterate mapping T, starting with $\nu_{\min}$, until we converge to a fixed point. Figure 2B shows the first iteration, $T(\nu_{\min})$. Compared to $\nu_{\min}$, it has

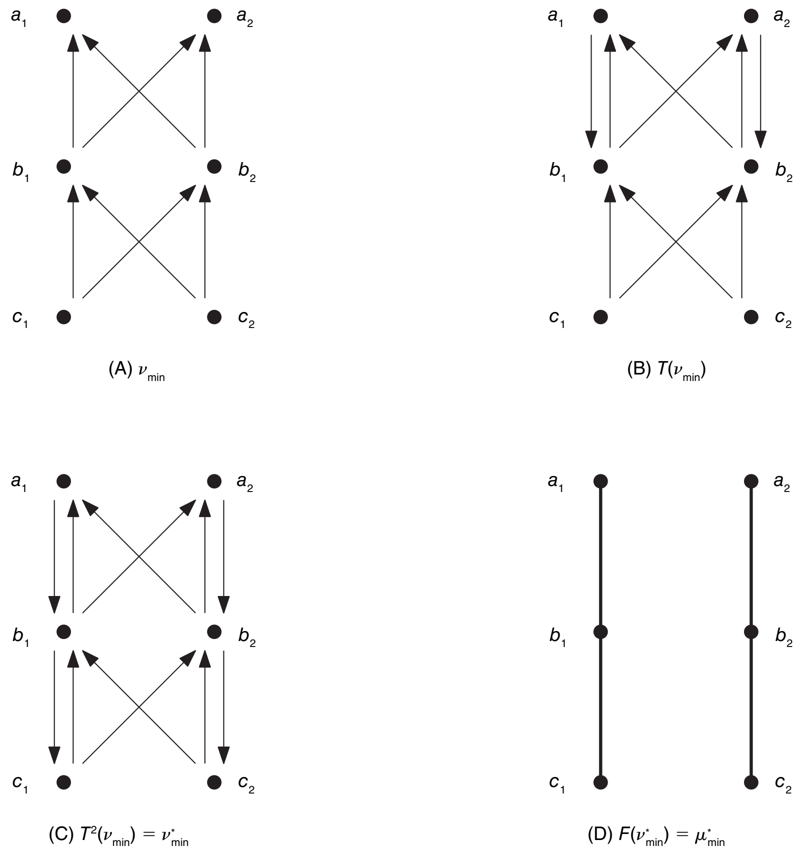

FIGURE 2. T-ALGORITHM

two new arrows: $\overrightarrow{a_1b_1}$ and $\overrightarrow{a_2b_2}$. To check that, say, arrow $\overrightarrow{a_1b_1}$ is in $T(\nu_{\min})$, note that $\nu_{\min}(a_1) = \{(a_1, b_1), (a_1, b_2)\}$ and $(a_1, b_1) \in Ch_{a_1}(\nu_{\min}(a_1) \cup (a_1, b_1)) = \{(a_1, b_1)\}$. To check that, say, arrow $\overrightarrow{b_1c_2}$ is not in $T(\nu_{\min})$, note that $\nu_{\min}(b_1) = \{(b_1, c_1), (b_1, c_2)\}$ and $(b_1, c_2) \notin Ch_{b_1}(\nu_{\min}(b_1) \cup (b_1, c_2)) = \emptyset$. The remaining 14 arrows can be checked analogously.

Figure 2C shows the next iteration, $T^2(\nu_{\min})$, which contains two new arrows. This iteration turns out to be a fixed point: one can check that $T^3(\nu_{\min}) = T^2(\nu_{\min})$. Hence, in this example, $\nu^*_{\min} = T^2(\nu_{\min})$. Finally, we need to apply mapping F to $\nu^*_{\min}$, erasing one-directional arrows and replacing reciprocal arrows with the corresponding contracts. The resulting chain-stable network, $\mu^*_{\min}$, is shown in Figure 2D.

III. Properties of Stable Networks

Theorem 1 states that the set of chain-stable networks is not empty. In this section, I present several results concerning the properties of this set, generalizing similar results concerning the properties of the set of pairwise-stable matchings in various two-sided settings.

By Lemma 1, there is a one-to-one correspondence between the set of fixed points of mapping T and the set of chain-stable networks. I will now rely on this result to establish several properties of the latter, using an auxiliary lemma. This lemma describes the special structure of the set of fixed points of mapping T. Pre-networks $\nu^*_{\min}$ and $\nu^*_{\max}$ are, respectively, its lowest and highest fixed points, and the set of fixed points of mapping T is a lattice contained between them. This result is a direct corollary of Tarski's fixed point theorem for isotone functions on lattices, but I also give a short independent proof for completeness.

LEMMA 3: *The set of fixed points of mapping T is a lattice. Its lowest element is $\nu^*_{\min}$ and its highest element is $\nu^*_{\max}$.*

Now, let $\overline{A} = \{a \in A : U_a(C) = \emptyset\}$ and $\underline{A} = \{a \in A : D_a(C) = \emptyset\}$, i.e., $\overline{A}$ and $\underline{A}$ are the sets of suppliers of basic inputs for the market ("suppliers") and consumers of final outputs ("consumers"), respectively. In the two-sided matching setup, one side of the market can be viewed as $\overline{A}$ and the other as $\underline{A}$; in more general networks, there is also a set of "intermediate" nodes, $A \backslash (\overline{A} \cup \underline{A})$.

The following theorem shows that *all* suppliers of basic inputs are *at least* as well off in network $\mu^*_{\min}$ as they are in any other chain-stable network, and *all* consumers of final outputs are *at most* as well off in network $\mu^*_{\min}$ as they are in any other chain-stable network. Symmetrically, all suppliers of basic inputs are at most as well off in network $\mu^*_{\max}$ as they are in any other chain-stable network, and all consumers of final outputs are at least as well off in network $\mu^*_{\max}$ as they are in any other chain-stable network.

THEOREM 2: *Let $\mu^*_{\min} = F(\nu^*_{\min})$, $\mu^*_{\max} = F(\nu^*_{\max})$, and let μ^* be a chain-stable network. Then any $a \in \overline{A}$ (weakly) prefers $\mu^*_{\min}$ to μ^* and μ^* to $\mu^*_{\max}$, and any $a \in \underline{A}$ (weakly) prefers $\mu^*_{\max}$ to μ^* and μ^* to $\mu^*_{\min}$.*

Recall the setting of Example 1. The set of chain-stable networks in this example is shown in Figure 3. It is easy to check that networks $\mu^*_{\min}$ and $\mu^*_{\max}$ are side-optimal. Note also that in this particular example, the most preferred chain-stable network for both intermediate nodes is μ^*_4, and their least preferred chain-stable network is μ^*_3. Generally, however, stable networks most or least preferred by all of the intermediate nodes need not exist.

The final result of this section shows that when a new supplier of basic inputs is added to the market, the set of chain-stable networks (or, more precisely, the boundaries of this set—networks $\mu^*_{\min}$ and $\mu^*_{\max}$) moves in the direction favorable to the consumers of final outputs and unfavorable to the suppliers of basic inputs. Symmetrically, when a new consumer of final outputs is added, the set of chain-stable networks moves in the opposite direction. In other words, more competition on one end of an industry is bad for the agents on that end and good for the agents on the other end. Of course, removing a supplier of basic inputs or a consumer of final outputs has the opposite effect on the remaining nodes.

More formally, let $A' = A \cup a'$ and let $\mu'_{\min}$ and $\mu'_{\max}$ be the smallest and the largest chain-stable matchings in A'.

THEOREM 3: *If $U_{a'}(A) = \emptyset$, i.e., a' is a supplier of basic inputs, then each $a \in \overline{A}$ is* at least *as well off in $\mu^*_{\max}$ as in $\mu'_{\max}$ and* at least *as well off in $\mu^*_{\min}$ as in $\mu'_{\min}$; each $a \in \underline{A}$ is* at most *as well off in $\mu^*_{\max}$ as in $\mu'_{\max}$ and* at most *as well off in $\mu^*_{\min}$ as in $\mu'_{\min}$.*

Symmetrically, if $D_{a'}(A) = \emptyset$, i.e., a' is a consumer of final outputs, then each $a \in \overline{A}$ is at most *as well off in $\mu^*_{\max}$ as in $\mu'_{\max}$ and* at most *as well off in $\mu^*_{\min}$ as in $\mu'_{\min}$; each $a \in \underline{A}$ is* at least *as well off in $\mu^*_{\max}$ as in $\mu'_{\max}$ and* at least *as well off in $\mu^*_{\min}$ as in $\mu'_{\min}$.*

Recall, again, the setting of Example 1, and suppose supplier of basic inputs a_1 is removed from the market. According to Theorem 3, the other supplier of basic inputs (a_2) should benefit

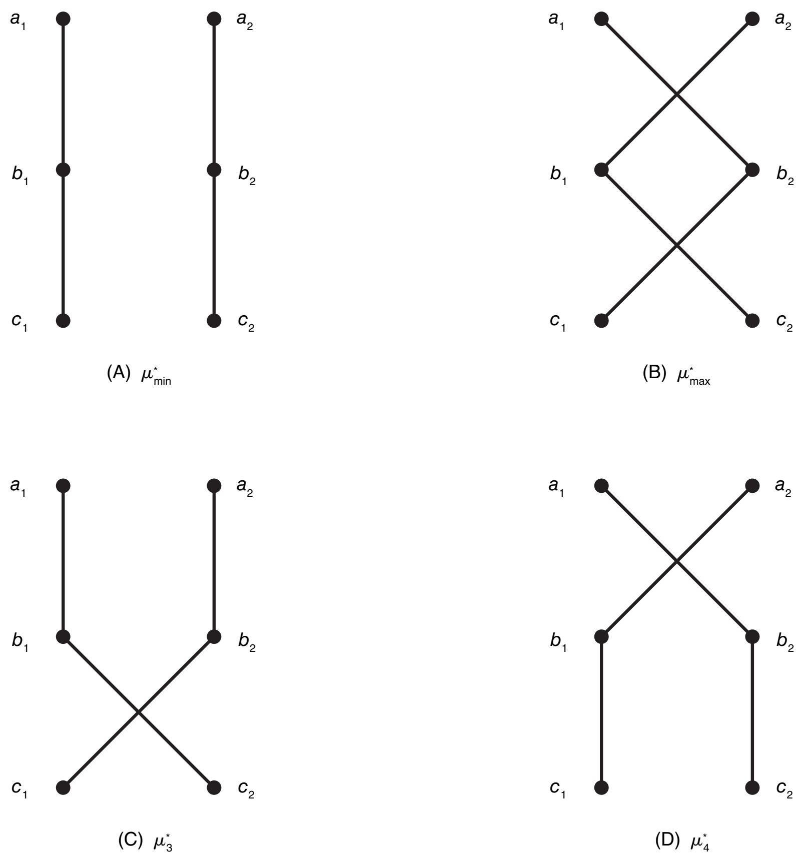

FIGURE 3. ALL CHAIN-STABLE NETWORKS

from the reduced competition, while the consumers of final outputs $(c_1$ and $c_2)$ should suffer. Indeed, this smaller market has only one stable network, which consists of contracts (a_2, b_2) and (b_2, c_2) —this can be verified by running the *T*-algorithm twice, starting from the two extreme pre-networks. In this network, node a_2 is connected with its preferred intermediary, and is therefore at least as well off as in both networks $\mu^*_{\min}$ and $\mu^*_{\max}$ in the larger market. Node c_1 is left without an intermediary, and is therefore strictly worse off than it was in networks $\mu^*_{\min}$ and $\mu^*_{\max}$. Finally, node c_2 is connected with its least preferred intermediary, and is therefore at most as well off as in $\mu^*_{\min}$ and $\mu^*_{\max}$.

As before, the change in the welfare of intermediate agents is ambiguous; it can go either way. Also, adding new intermediate nodes can have opposite effects on different extreme nodes (e.g., some suppliers of basic inputs may become better off, while others may become worse off), as well as on other intermediate nodes.

IV. Chain Stability and Alternative Solution Concepts

The model introduced and studied in this paper involves a new solution concept: chain stability. In this section, I discuss some alternative solution concepts, argue that chain stability is a natural one to use in the present setting, and establish results on the connections between chain stability, tree stability, and the weak core.

In two-sided one-to-one matching markets, the set of pairwise-stable matchings coincides with the core and with other solution concepts proposed in the matching literature. In more general models, this is no longer true, even when preferences are substitutable. In one-to-many matching markets, the set of pairwise-stable matchings is equal to the weak core but not to the strict core. In two-sided many-to-many matching markets, even that result no longer holds: Example 2.6 in Blair (1988) shows that the intersection of the set of pairwise-stable matchings and the core may be empty. Several other solution concepts for the many-to-many matching model have also been studied, including setwise stability (Roth 1984; Sotomayor 1999), bargaining set (Echenique and Oviedo 2006), and credible group stability (Konishi and Ünver 2006). Nevertheless, pairwise stability remains the preferred solution concept for studying matching even in two-sided many-to many markets, due to its simplicity, natural interpretation, and apparent empirical success. As Roth and Sotomayor (1990, 156) argue, "identifying and organizing large coalitions may be more difficult than making private arrangements between two parties, and the experience of those regional [many-to-many matching] markets in the United Kingdom that are built around stable mechanisms suggests that pairwise stability is still of primary importance in these markets."

Since the model of this paper includes two-sided many-to-many matching models as special cases, different solution concepts can result in different predictions; e.g., the core and the set of chain-stable networks may be disjoint. However, just like pairwise stability in the two-sided setting, chain stability is a natural solution concept in the supply chain environment. The reason is that chain blocks are particularly easy to identify and organize: a customer need only pick up the phone and call a potential supplier asking him whether he would like to form a contract; the potential supplier, after receiving that phone call, calls one of his potential suppliers, and so on. If there is a chain block, it can be easily identified in this way, and subsequently the contracts can be formed. In fact, firms are often organized in a way that helps them identify chain blocks (or, more precisely, the parts of chain blocks that concern them). Sales and marketing departments identify available downstream contracts and keep track of which ones may potentially be available in the future. Procurement departments identify upstream options. Top management identifies ways of forming profitable combinations of upstream and downstream contracts. It is common for companies to view themselves as members of supply chains and to mention their suppliers to potential customers,[2] and vice versa.[3]

In contrast, coalitions involving several *competing* firms require much more coordination and information exchange between the agents. Many firms do invest in "competitive intelligence," trying to learn what their competitors are doing or planning to do. The goal, however, is to figure out how to respond to the behavior of competitors, or to learn the latest technologies or innovations from them, but not to identify large coalitional deviations that would involve these competitors. Moreover, such deviations would often be illegal, as violations of antitrust laws. Hence, while deviations involving competing agents are conceivable in principle, they must be rare in

[2] For example, a steel processor in the United Kingdom writes on its Web site: "Sourcing steel from all over the world, primarily Western Europe, Steel & Alloy currently process over 320,000 tonnes of steel per annum. We are approved processors for both Arcelor and Corus and are the preferred processor for Salzgitter within the U.K. Automotive Sector. These partnerships give us the stability and technical support needed to provide a total supply solution for our customers" (http://www.steelalloy.co.uk, accessed July 20, 2006).

[3] For example, a trading company in India looking for a timber supplier writes in its online ad, "Our client is a supplier of packing and packaging wooden materials to large scale factories in and around Chennai. […] We would like to know, what type of wood you will be able to supply us, in break bulk/in containers? […]" (http://www.alibaba.com/manufacturer/13917627/Buy_Logs_And_Timber.html, accessed July 20, 2006).

practice, at least in industries that comply with antitrust laws. Thus, while chain stability may not be an attractive solution concept for modeling industries with monopolies and cartels, it is a natural one in competitive markets.

Solution concepts based on deviations by large coalitions involving competing agents also suffer from an additional, theoretical problem: the set of such solutions may be empty, even under very strong restrictions on preferences. Example 3 in Sotomayor (1999) shows that the set of setwise-stable matchings in two-sided markets may be empty.[4] Example 2 in Konishi and Ünver (2006) shows that the core in a two-sided many-to-many matching problem may also be empty, even under responsive preferences. In contrast, as the results of this paper show, a chain-stable network is guaranteed to exist, at least under same-side substitutability and cross-side complementarity, just like a pairwise-stable matching is guaranteed to exist in any two-sided many-to-many matching market under substitutability. Of course, the similarity between the empirical and theoretical arguments behind pairwise stability in the two-sided case and chain stability in the more general case is not a coincidence: chain stability reduces to pairwise stability if there are no intermediate agents.

Still, it is important to understand the differences and similarities between various solution concepts in matching markets. Several papers address these issues in two-sided markets (see, e.g., the recent papers by Echenique and Oviedo (2006) and Konishi and Ünver (2006), and references in those papers). The following two results provide a starting point for the analysis of the relationship between chain stability and other solution concepts in supply chain networks.

The first result shows that under same-side substitutability and cross-side complementarity, blocking by "trees" is equivalent to blocking by chains. More formally, a sequence of contracts $\mathbf{c}_1, \dots, \mathbf{c}_i$ is a *path* from node a to node b if: (i) node a is involved in contract $\mathbf{c}_1$ and not involved in any contract $\mathbf{c}_j$ for $j > 1$; (ii) node b is involved in contract $\mathbf{c}_i$ and not involved in any contract $\mathbf{c}_j$ for $j < i$; and (iii) any other node x involved in one of the contracts $\mathbf{c}_j$ for $1 \leq j \leq i$ is involved in exactly two such contracts, and these two contracts are adjacent in the sequence (i.e., if one of the contracts is $\mathbf{c}_k$, then the other is either $\mathbf{c}_{k-1}$ or $\mathbf{c}_{k+1}$). Note that while each chain is a path, there are paths that are not chains: e.g., a pair of contracts with the same buyer and two different sellers is a path connecting the two sellers, but is not a chain. A *tree* is a set of contracts such that for any two nodes involved in these contracts, there exists exactly one path in this set connecting the two nodes. Note that every chain is a tree. A network, μ, is blocked by a tree, τ, if $\tau \cap \mu = \emptyset$ and for every node a involved in τ, $\tau(a) \subset Ch_a(\mu(a) \cup \tau(a))$. A network is tree stable if it is not blocked by any tree.

THEOREM 4: *Under same-side substitutability and cross-side complementarity, the set of tree-stable networks is equal to the set of chain-stable networks.*

The final result of this section shows that in a special case, in which each node is restricted to having at most one upstream contract and at most one downstream contract, the set of chain-stable networks coincides with the weak core of the matching game. Network μ is in the *weak core* of the matching game if, and only if, there is no other network μ' and set M of nodes such that (i) for every node $a \in M$, for every contract $\mathbf{c}$ involving a, the other node involved in $\mathbf{c}$ is also in set M; (ii) every node $a \in M$ weakly prefers the set of contracts in which it is involved in μ' to the set of contracts in which it is involved in μ; and (iii) at least one node $a \in M$ strictly prefers the set of contracts in which it is involved in μ' to the set of contracts in which it is involved in μ.

[4] A matching is setwise stable if, roughly, there are no coalitions of agents who can form additional links or remove existing ones among themselves and possibly remove some of the links to other agents in such a way that they all become better off.

THEOREM 5: *If each node $a \in A$ can have at most one upstream contract and at most one downstream contract and has same-side substitutable and cross-side complementary preferences, then the set of chain-stable networks is equal to the weak core of the matching game.*

V. Applications and Extensions

Applications of matching theory have traditionally been found in market design literature, which views stability as a desirable property of a mechanism and uses the deferred acceptance algorithm or its extensions to match market participants (see Roth 2007 for a historical overview). Recently, however, researchers have also started using it as a positive theory, imposing stability as a restriction on market outcomes and using this restriction to estimate various parameters of interest: e.g., Donald Boyd et al. (2003) use it to study the labor market for public school teachers, Patrick Bajari and Jeremy T. Fox (2005) impose it in their analysis of the outcomes of FCC spectrum auctions, and Morten Sørensen (2007) relies on it to model the matching of startups to venture capitalists.

The theory of matching in supply chains can also be used in both kinds of applications. First, *T*-algorithm or its extensions could potentially be used in market design. Second, the model can be viewed as a positive theory of matching in vertical networks, allowing researchers to impose chain stability restriction to estimate various parameters of interest and to run counterfactual policy simulations. Not every supply chain can be modeled in this framework: e.g., the presence of economies of scale or multiple complementary inputs would not fit the model's assumptions. At the same time, there is a wide range of markets to which the theory can be applied; some examples of such markets are discussed below.

A. *Market Design*

A natural area for design applications of the results of this paper are professional services markets. Such markets often consist of three groups of agents: workers providing these services who sign contracts with firms (and view these firms as substitutes), clients who also sign contracts with the firms (and view them as substitutes), and the firms that sign contracts with the workers and the clients. From the point of view of a firm, workers are substitutes for one another, clients are substitutes as well, and workers and clients are complements: the more clients the firm has, the more workers it is willing to hire, and vice versa.

For example, high school students in New York City are currently matched to public schools via a mechanism based on an extension of the deferred acceptance algorithm (Atila Abdulkadiroğlu, Parag A. Pathak, and Roth 2005). The schools also participate in another matching market—the labor market for teachers (see Boyd et al. (2003) for the analysis of this market in a matching framework). At present, the organization of the latter market is similar to that of most other labor markets: schools interview candidates and make them offers and then the candidates choose whether to accept or reject them (the terms of these offers, however, are substantially restricted by the agreement between the New York City Department of Education (DOE) and the United Federation of Teachers (UFT)). In the future, given the success of the student matching mechanism, the NYC DOE and the UFT may want to run a similar program for matching teachers to schools. In that case, *T*-algorithm and the framework of this paper could be used to coordinate these two matching markets and ensure that the number of teachers matched to each school is in line with student enrollment there.

The experience of two-sided market design suggests that the mechanism would require some modifications (e.g., the algorithm for matching medical residents to hospitals had to deal with the issue of married couples, and the algorithm for matching students to public schools had to deal with the issue of breaking indifferences; see Roth 2007). Nevertheless, the deferred acceptance algorithm and the theory of stability in two-sided markets are at the heart of the mechanisms that were eventually implemented, and similarly the core ideas behind *T*-algorithm and the theory of matching in supply chains could form the basis for the mechanisms in more complicated markets.

Another caveat is incentive compatibility. In theory, even with only two sides, at least for some agents it is not a dominant strategy to report their preferences truthfully (Roth 1982). This lack of incentive compatibility automatically carries over to the more general setting of supply chain networks. In practice, however, the lack of incentive compatibility does not seem to cause problems in two-sided markets (Roth and Elliott Peranson 1999), most likely due to the fact that the number of agents in these markets is large (Nicole Immorlica and Mohammad Mahdian 2005; Fuhito Kojima and Pathak 2007). It is an open question whether these results generalize to the supply chain setting.

B. *Positive Theory of Stability in Supply Chain Networks*

In markets that lack coordinating bodies like the DOE and the UFT, implementing a centralized matching mechanism may be infeasible. The theory of matching in supply chains, however, may still be applicable to such markets: it can be viewed in a descriptive light, like, for example, various equilibrium concepts in the empirical industrial organization literature.

One class of markets to which the theory is particularly well suited is industries with the following feature. There is one key good (or type of good) that moves along the supply chain. It may be transformed, repackaged, processed, transported, etc., by intermediaries, but, crucially, has to remain the key input and output for firms at each stage of the chain. Of course, in all industries there will be some other complementary inputs used by the firms in the chain (e.g., firms in all industries use computers, printers, office supplies, etc.), but as long as the cost of these complementary inputs is small or these inputs are simple and homogeneous, they can be modeled simply as an expense, included in the cost of converting inputs into outputs. In some industries (e.g., construction), firms along supply chains combine several complementary inputs to produce final goods, with inputs themselves consisting of multiple complementary parts, many of them heterogeneous, complex, and an important part of the final cost of the outputs. For such industries, the theory developed in the current paper is not applicable. In many others, however, this "key good" can be easily identified. Wheat, which is turned into flour, is such a good in the farmer–miller–baker supply chain. Many other agricultural supply chains have similar key goods (milk, fish, coffee, grapes, and so on). Oil and natural gas industries have long and complicated supply chains involving production, transportation, refining, and distribution among other steps, but the key good moving along the chain is easily identifiable. Various manufacturing supply chains also have such goods: e.g., wood in the supply chain for wooden furniture; iron ore and scrap metal in the steel supply chain; and cotton, wool, and linen in the textile supply chain. Some less typical industries also have this feature. For example, supply chains for illegal drugs, such as heroin and cocaine, involve, in essence, moving just one good from the source of production to destination, with only a small amount of processing along the way and extreme trade barriers and transportation costs, resulting in a high degree of price heterogeneity.

C. *Two-Sided Markets with Complementarities*

The results of this paper rely on same-side substitutability. There is, however, one special case in which this restriction can be partially relaxed: two-sided matching markets with two types of agents or objects on one of the sides. The agents of one type are substitutes for one another, but the agents of different types are complements, from the point of view of agents on the other side of the market. For example, workers and machines may be complements from the point of view of a manufacturing firm, while workers are substitutes for one another and machines are substitutes for one another as well. A similar model is analyzed by Sun and Yang (2006), who assume that firms' preferences are quasi-linear and prices are continuous. They show that in such a market, a competitive equilibrium always exists.

The model of the current paper contains the discrete price version of Sun and Yang's model as a special case. Indeed, suppose prices can take only a finite number of values (e.g., have to be multiples of one cent and cannot exceed some arbitrarily large amount), and consider a "supply chain market" in which machines are viewed as suppliers of basic inputs, firms are viewed as intermediaries, and workers are viewed as consumers of final outputs. Then the complementarity of the two types of agents on one side of the two-sided market becomes equivalent to the cross-side complementarity of firms' preferences in the "supply chain market," and all the assumptions of the current paper are satisfied. Hence, all the conclusions remain valid. Stable networks exist. Such networks form a lattice with the one most preferred by the machines and the one most preferred by the workers as the extreme elements. At these extreme networks, the payoffs of workers will weakly decrease and the payoffs of machines will weakly increase if a worker is added to the market (and vice versa). Sun and Yang (2006) do not present any results on the lattice structure, the opposition of interests, or the comparative statics. Note also that these results do not require quasi-linear preferences and continue to hold even if prices are fixed or not fully flexible, and workers and machines can contract with multiple firms and may care not just about the payments they receive but also about the identities of the firms they contract with.

Moreover, the results can be extended to a setting in which the suppliers of these complementary goods are themselves intermediaries in two distinct vertical networks. For instance, if the producers of machines in the example above require steel to make them, we can consider the steel producers → machine producers → manufacturing firms → manufacturing firm workers "supply chain market" and apply to it the results of Sections II, III, and IV.

VI. Conclusion

This paper introduces and studies matching in supply chain networks. The model allows for matching both with and without prices and can incorporate endogenously determined quantities and types of traded goods or more general contracts. The results rely on two important assumptions. First, the roles of agents in an industry are predetermined. There is an "upstream–downstream" partial ordering, which specifies who can be a supplier and who can be a customer for each node. This ensures that trade proceeds only in one direction, from the suppliers of basic inputs to the consumers of final outputs. Second, the preferences of all agents satisfy same-side substitutability and cross-side complementarity. Each agent views its inputs and outputs as complements, views its inputs as substitutes for one another, and also views its outputs as substitutes for one another.

In this setting, the paper studies chain-stable networks. A chain-stable network is a set of bilateral contracts such that no upstream–downstream sequence of agents can add a chain of contracts (and drop, if necessary, some other contracts) that would make them all better off. The main result of the paper is that under same-side substitutability and cross-side complementarity, chain-stable networks are guaranteed to exist. There is an iterative algorithm for finding a chain-stable network. This algorithm converges to a special chain-stable network: in all other chain-stable networks, all suppliers of basic inputs are weakly worse off and all consumers of final outputs are weakly better off. A slightly modified algorithm finds a "symmetric" (and generally different) chain-stable network, preferred by all consumers of final outputs. At these side-optimal chain-stable networks, an intuitive result on comparative statics holds: adding a consumer of final outputs makes other consumers of final outputs worse off and makes the suppliers of basic inputs better off, while adding a supplier of basic inputs has the opposite effect.

Appendix A: Proofs

PROOF OF LEMMA 1:

The proof consists of two independent steps. The first step shows that for any pre-network ν such that $T(\nu) = \nu$, $\mu = F(\nu)$ is a chain-stable network. The second step shows that for any chain-stable network μ, there exists a unique fixed-point pre-network ν such that $F(\nu) = \mu$.

Step 1: Let us show that for any pre-network ν such that $T(\nu) = \nu$, $\mu = F(\nu)$ is a chain-stable network.

First, we need to show that there are no contracts in μ that differ only in price. To see this, note that if there are two contracts that differ only in price in μ, that implies that each of the two agents involved in these contracts would choose both contracts when selecting from some larger set containing them. But this is impossible, because each agent, by definition, chooses only one contract with a given partner and a particular serial number—the one with the most favorable price.

Second, we need to show that network μ is individually rational. To see that, note that for any agent a, $\mu(a) = Ch_a(\nu(a))$, and so a does not want to drop any of its contracts in μ (because that would imply that $\mu(a) \neq Ch_a(\mu(a)) = Ch_a(Ch_a(\nu(a))) = Ch_a(\nu(a)) = \mu(a)$).

Finally, we need to show that there are no chain blocks. Suppose $(\mathbf{c}_1, \mathbf{c}_2, \ldots, \mathbf{c}_n)$ is a chain block of μ, and let s_i and b_i denote the seller and the buyer involved in contract i. Since $\mathbf{c}_1 \in Ch_{s_1}(\mu(s_1) \cup \mathbf{c}_1)$, it has to be the case that $\mathbf{c}_1 \in Ch_{s_1}(\nu(s_1) \cup \mathbf{c}_1)$ (otherwise, $Ch_{s_1}(\nu(s_1) \cup \mathbf{c}_1) = Ch_{s_1}(\nu(s_1)) = \mu(s_1)$, and hence no subset of $\mu(s_1) \cup \mathbf{c}_1 \subset \nu(s_1) \cup \mathbf{c}_1$ can be better for s_1 than $\mu(s_1)$), and so the arrow $\mathbf{r}_1$ from s_1 to b_1 with $\mathbf{c}_1$ attached must be in $T(\nu) = \nu$. Now, by assumption, s_2 would like to sign contracts $\mathbf{c}_1$ and $\mathbf{c}_2$, i.e., $\{\mathbf{c}_1, \mathbf{c}_2\} \subset Ch_{s_2}(\mu(s_2) \cup \mathbf{c}_1 \cup \mathbf{c}_2)$. If neither $\mathbf{c}_1$ nor $\mathbf{c}_2$ are in $Ch_{s_2}(\nu(s_2) \cup \mathbf{c}_1 \cup \mathbf{c}_2)$, then $Ch_{s_2}(\nu(s_2) \cup \mathbf{c}_1 \cup \mathbf{c}_2) = Ch_{s_2}(\nu(s_2)) = \mu(s_2)$, and so $\{\mathbf{c}_1, \mathbf{c}_2\} \not\subset Ch_{s_2}(\mu(s_2) \cup \mathbf{c}_1 \cup \mathbf{c}_2)$, which would contradict our assumptions. Suppose $\mathbf{c}_2 \notin Ch_{s_2}(\nu(s_2) \cup \mathbf{c}_1 \cup \mathbf{c}_2)$. Then $\mathbf{c}_1 \in Ch_{s_2}(\nu(s_2) \cup \mathbf{c}_1 \cup \mathbf{c}_2) = Ch_{s_2}(\nu(s_2) \cup \mathbf{c}_1)$, and so there must be an arrow from s_2 to s_1 with contract $\mathbf{c}_1$ attached in $T(\nu) = \nu$, which together with the fact that there is an arrow from s_1 to s_2 with $\mathbf{c}_1$ attached in ν would imply that $\mathbf{c}_1 \in \mu$, which would also contradict our assumptions. Hence, it must be the case that $\mathbf{c}_2 \in Ch_{s_2}(\nu(s_2) \cup \mathbf{c}_1 \cup \mathbf{c}_2)$. Proceeding by induction, there is an arrow from s_i to s_{i+1} with $\mathbf{c}_i$ attached in ν for any $i < n$. Similarly, we could have started from node b_n, and so there must be an arrow going from b_n to $b_{n-1} = s_n$ with $\mathbf{c}_n$ attached in ν, which implies that $\mathbf{c}_n \in \mu$—contradiction. Therefore, for any $\nu = T(\nu)$, $F(\nu)$ is a chain-stable network.

Step 2: Let us now show that for any chain-stable network μ, there exists a unique fixed-point pre-network ν such that $F(\nu) = \mu$. The proof is constructive.

Let M be the set of all networks and N be the set of all pre-networks. Define the following mappings:

- $G: M \to N$. $G(\mu) = \{\mathbf{r} \in R | \mathbf{c}_\mathbf{r} \in \mu\}$, i.e., arrow $\mathbf{r}$ belongs to $\nu = G(\mu)$ if and only if contract $\mathbf{c}_\mathbf{r}$ belongs to μ.
- $H_k: M \to N$. $H_k(\mu) = T^k(G(\mu))$, i.e., $H_0(\mu) = \mathrm{G}(\mu)$ and $H_k(\mu) = T(H_{k-1}(\mu))$ for $k > 0$.

We will now show that for any chain-stable network μ, for some n, $H_n(\mu) = H_{n+1}(\mu)$ and moreover, for each n, $F(H_n(\mu)) = \mu$, thus giving a constructive proof of Step 2. For convenience, let $\nu^n = H_n(\mu)$, $\nu^0 = G(\mu)$. We will show by induction on k that: (i) $F(\nu^k) = \mu$ and (ii) $\nu^k \supset \nu^{k-1}$.

For $k = 1$, (ii) follows from the individual rationality of μ and (i) follows from the absence of chain blocks of length 1 (i.e., pairwise blocks) of μ. Suppose (i) and (ii) hold up to $k - 1$. Let us show that they hold for k.

(i) Suppose arrows $\mathbf{r}$ and $\mathbf{r}'$ with contract $\mathbf{c_r}$ attached are in ν^k, but $\mathbf{c_r} \notin \mu$. We will now "grow" a chain block of μ from this contract $\mathbf{c_r}$.

Consider arrow $\mathbf{r}$ first; without loss of generality, assume it is an upstream arrow. Let $\mathbf{c}_0 = \mathbf{c_r}$ and $\mathbf{r}_0 = \mathbf{r}$. If $\mathbf{c_r} \in Ch_{o_\mathbf{r}}(\mu(o_\mathbf{r}) \cup \mathbf{c_r})$, stop. Otherwise, since

$$\mathbf{c_r} \in Ch_{o_\mathbf{r}}(\nu^{k-1}(o_\mathbf{r}) \cup \mathbf{c_r}),$$

$$\nu^{k-1}(o_\mathbf{r}) \cup \mathbf{c_r} = \big(D_{o_\mathbf{r}}(\nu^{k-1}(o_\mathbf{r}))\big) \cup \big(U_{o_\mathbf{r}}(\nu^{k-1}(o_\mathbf{r})) \cup \mathbf{c_r}\big), \text{ and}$$

$$\mu(o_\mathbf{r}) \subset \nu^{k-1}(o_\mathbf{r}),$$

it follows from same-side substitutability that

$$\mathbf{c_r} \in Ch_{o_\mathbf{r}}\big((D_{o_\mathbf{r}}(\nu^{k-1}(o_r))) \cup (U_{o_\mathbf{r}}(\mu(o_\mathbf{r})) \cup \mathbf{c_r})\big).$$

Since by assumption, $\mathbf{c_r} \notin Ch_{o_\mathbf{r}}(\mu(o_\mathbf{r}) \cup \mathbf{c_r}) = Ch_{o_\mathbf{r}}((D_{o_\mathbf{r}}(\mu(o_r))) \cup (U_{o_\mathbf{r}}(\mu(o_\mathbf{r})) \cup \mathbf{c_r}))$ and $\mu(o_\mathbf{r}) \subset \nu^{k-1}(o_\mathbf{r})$, it has to be the case that $D_{o_\mathbf{r}}(\mu(o_\mathbf{r}))$ is a strict subset of $D_{o_\mathbf{r}}(\nu^{k-1}(o_\mathbf{r}))$. Let $\mathbf{x}_1, \mathbf{x}_2, \ldots, \mathbf{x}_m$ be the contracts in $D_{o_\mathbf{r}}(\nu^{k-1}(o_\mathbf{r}))\backslash D_{o_\mathbf{r}}(\mu(o_\mathbf{r}))$. Then for some j, $\mathbf{x}_j \in Ch_{o_\mathbf{r}}(\mu(o_\mathbf{r}) \cup \mathbf{x}_j \cup \mathbf{c_r})$ (otherwise, by same-side substitutability, for any j, $\mathbf{x}_j \notin Ch_{o_\mathbf{r}}(\mu(o_\mathbf{r}) \cup \mathbf{c_r} \cup \mathbf{x}_1 \cup \mathbf{x}_2 \cup \cdots \cup \mathbf{x}_m)$, which is then equal to $Ch_{o_\mathbf{r}}(\mu(o_\mathbf{r}) \cup \mathbf{c_r})$, which contradicts our assumption that $\mathbf{c_r} \notin Ch_{o_\mathbf{r}}(\mu(o_\mathbf{r}) \cup \mathbf{c_r})$). It must also be the case that $\mathbf{c_r} \in Ch_{o_\mathbf{r}}(\mu(o_\mathbf{r}) \cup \mathbf{x}_j \cup \mathbf{c_r})$, because otherwise $\mathbf{x}_j \in Ch_{o_\mathbf{r}}(\mu(o_\mathbf{r}) \cup \mathbf{x}_j)$ and so the downstream arrow with $\mathbf{x}_j$ attached is in ν^1, and is therefore in ν^{k-1} (by statement (ii) in the step of induction, $\nu^1 \subset \nu^{k-1}$). But then both arrows with $\mathbf{x}_j$ attached are in ν^{k-1}, which contradicts assumption (i) of induction for $k - 1$. Hence, $\{\mathbf{c_r}, \mathbf{x}_j\} \subset Ch_{o_\mathbf{r}}(\mu(o_\mathbf{r}) \cup \mathbf{x}_j \cup \mathbf{c_r})$.

Let $\mathbf{c}_1 = \mathbf{x}_j$. By construction, the upstream arrow $\mathbf{r}_1$ with $\mathbf{c}_1$ attached is in ν^{k-1}, but $\mathbf{c}_1 \notin \mu$. Let o_1 denote the origin of arrow $\mathbf{r}_1$. If $\mathbf{c}_1 \in Ch_{o_1}(\mu(o_1) \cup \mathbf{c}_1)$, stop; otherwise, following the procedure above, generate $\mathbf{c}_2 \in D_{o_1}(\nu^{k-2}(o_1))\backslash D_{o_1}(\mu(o_1))$, and so on. At some point, this procedure will have to stop (since we keep going downstream). Now, "grow" $\mathbf{c_r}$ in the other direction, starting with arrow $\mathbf{r}'$. We end up with a chain $\mathbf{c}_x, \mathbf{c}_{x+1}, \ldots, \mathbf{c}_0, \ldots, \mathbf{c}_{y-1}, \mathbf{c}_y$, which, by construction, is a chain block of μ—contradiction.

(ii) Suppose some upstream arrow $\mathbf{r}$ is in ν^{k-1}, but not in ν^k, i.e., $\mathbf{c_r} \in Ch_{o_\mathbf{r}}(\nu^{k-2}(o_\mathbf{r}) \cup \mathbf{c_r})$, but $\mathbf{c_r} \notin Ch_{o_\mathbf{r}}(\nu^{k-1}(o_\mathbf{r}) \cup \mathbf{c_r})$. Then by (i), $Ch_{o_\mathbf{r}}(\nu^{k-1}(o_\mathbf{r}) \cup \mathbf{c_r}) = \mu(o_\mathbf{r}) = Ch_{o_\mathbf{r}}(D_{o_\mathbf{r}}(\nu^{k-1}(o_\mathbf{r})) \cup U_{o_\mathbf{r}}(\mu(o_\mathbf{r})) \cup \mathbf{c_r})$. From $\mathbf{c_r} \notin Ch_{o_\mathbf{r}}(D_{o_\mathbf{r}}(\nu^{k-1}(o_\mathbf{r})) \cup U_{o_\mathbf{r}}(\mu(o_\mathbf{r})) \cup \mathbf{c_r})$, by same-side substitutability and by assumptions of induction for k−2, we get $\mathbf{c_r} \notin Ch_{o_\mathbf{r}}(D_{o_\mathbf{r}}(\nu^{k-1}(o_\mathbf{r})) \cup U_{o_\mathbf{r}}(\nu^{k-2}(o_\mathbf{r})) \cup \mathbf{c_r})$, and from that, by cross-side complementarity and assumption (ii) of induction for k−1 (i.e., $\nu^{k-2} \subset \nu^{k-1}$), we get $\mathbf{c_r} \notin Ch_{o_\mathbf{r}}(D_{o_\mathbf{r}}(\nu^{k-2}(o_\mathbf{r})) \cup U_{o_\mathbf{r}}(\nu^{k-2}(o_\mathbf{r})) \cup \mathbf{c_r}) = Ch_{o_\mathbf{r}}(\nu^{k-2}(o_\mathbf{r}) \cup \mathbf{c_r})$, and so $\mathbf{r}$ is not in ν^{k-1}—contradiction. The proof for a downstream arrow $\mathbf{r}'$ is completely analogous. This completes the proof of statements (i) and (ii) of induction.

Now, since $G(\mu) \subset H_1(\mu) \subset H_2(\mu) \subset \ldots$ is an increasing sequence and the set of possible arrows is finite, this sequence has to converge, i.e., for some n, $H_n(\mu) = H_{n+1}(\mu) \equiv H(\mu)$. By (ii), all arrows in $G(\mu)$ are also present in $H(\mu)$, and by (i), any pair of arrows with the same contract attached in $H(\mu)$ is also present in $G(\mu)$. Therefore, $F(H(\mu)) = \mu$.

Finally, we need to show that for two fixed points of mapping T, ν_1^* and ν_2^*, $F(\nu_1^*) \neq F(\nu_2^*)$. Suppose $\nu_1^* \neq \nu_2^*$ and $F(\nu_1^*) = F(\nu_2^*) = \mu$. Consider the set of agents for whom the upstream arrows originating from them are not the same in ν_1^* and ν_2^*. Take one of the "most downstream" agents in this set (i.e., such an agent o that there is nobody downstream from him in this set), and take an upstream arrow $\mathbf{r}$ originating from o such that it is in only one of the two pre-networks. Without loss of generality, $\mathbf{r} \in \nu_1^*$ and $\mathbf{r} \notin \nu_2^*$. $\mathbf{r} \notin \nu_2^* \Rightarrow \mathbf{c_r} \notin Ch_o(\nu_2^*(o) \cup \mathbf{c_r}) = Ch_o(\nu_2^*(o)) = \mu(o) = Ch_o(D_o(\nu_2^*(o)) \cup U_o(\mu(o)) \cup \mathbf{c_r})$. By the assumption that o is the "most downstream" agent whose upstream arrows differ in the two pre-networks, $D_o(\nu_2^*(o)) = D_o(\nu_1^*(o))$, and hence $\mathbf{c_r} \notin Ch_o(D_o(\nu_1^*(o)) \cup U_o(\mu(o)) \cup \mathbf{c_r})$. Now, since $F(\nu_1^*) = \mu$, $U_o(\nu_1^*(o)) \supset U_o(\mu(o))$, and so by same-side substitutability, $\mathbf{c_r} \notin Ch_o(D_o(\nu_1^*(o)) \cup U_o(\nu_1^*(o)) \cup \mathbf{c_r}) = Ch_o(\nu_1^*(o) \cup \mathbf{c_r})$, and therefore $\mathbf{r} \notin \nu_1^*$—contradiction.

PROOF OF LEMMA 2:

We need to check that all downstream arrows in $T(\nu_1)$ belong to $T(\nu_2)$ and that all upstream arrows in $T(\nu_2)$ belong to $T(\nu_1)$. Consider a downstream arrow $\mathbf{r}$ in $T(\nu_1)$. By definition of mapping T, $\mathbf{c_r} \in Ch_{o_\mathbf{r}}(\nu_1(o_\mathbf{r}) \cup \mathbf{c_r})$. Since $\nu_1 \leq \nu_2$, by the definition of the partial order on pre-networks, the set of downstream arrows in ν_1 is a subset of the set of downstream arrows in ν_2, and so $U_{o_\mathbf{r}}(\nu_1(o_\mathbf{r}))$, i.e., the set of contracts attached to downstream arrows pointing to $o_\mathbf{r}$ in ν_1 is a subset of $U_{o_\mathbf{r}}(\nu_2(o_\mathbf{r}))$, i.e., the set of contracts attached to downstream arrows pointing to $o_\mathbf{r}$ in ν_2. Analogously, we have $D_{o_\mathbf{r}}(\nu_1(o_\mathbf{r})) \supset D_{o_\mathbf{r}}(\nu_2(o_\mathbf{r}))$. By same-side substitutability, $\mathbf{c_r} \in Ch_{o_\mathbf{r}}(\nu_1(o_\mathbf{r}) \cup \mathbf{c_r}) = Ch_{o_\mathbf{r}}(U_{o_\mathbf{r}}(\nu_1(o_\mathbf{r})) \cup D_{o_\mathbf{r}}(\nu_1(o_\mathbf{r})) \cup \mathbf{c_r})$ implies $\mathbf{c_r} \in Ch_{o_\mathbf{r}}(U_{o_\mathbf{r}}(\nu_1(o_\mathbf{r})) \cup D_{o_\mathbf{r}}(\nu_2(o_\mathbf{r})) \cup \mathbf{c_r})$. This, by cross-side complementarity, implies that $\mathbf{c_r} \in Ch_{o_\mathbf{r}}(U_{o_\mathbf{r}}(\nu_2(o_\mathbf{r})) \cup D_{o_\mathbf{r}}(\nu_2(o_\mathbf{r})) \cup \mathbf{c_r})$, i.e., $\mathbf{c_r} \in Ch_{o_\mathbf{r}}(\nu_2(o_\mathbf{r}) \cup \mathbf{c_r})$, and so $\mathbf{r} \in T(\nu_2)$.

For upstream arrows, the argument is completely symmetric.

PROOF OF LEMMA 3:

Take two fixed points of mapping T, ν_1^* and ν_2^*. Let ν_{12} be the least upper bound of these two pre-networks in the original lattice. $\nu_{12} \geq \nu_1^*$, $\nu_{12} \geq \nu_2^* \Rightarrow T(\nu_{12}) \geq T(\nu_1^*) = \nu_1^*$, $T(\nu_{12}) \geq T(\nu_2^*) = \nu_2^* \Rightarrow T(\nu_{12}) \geq \nu_{12}$, and so for some n, $\nu_{12} \leq T(\nu_{12}) \leq T^2(\nu_{12}) \leq \cdots \leq T^n(\nu_{12}) = T^{n+1}(\nu_{12}) = \nu_{12}^*$. By construction, $\nu_{12}^* \geq \nu_1^*$ and $\nu_{12}^* \geq \nu_2^*$. To see that any other upper bound of ν_1^* and ν_2^* (say, ν_3^*) has to be greater than ν_{12}^*, note that $\nu_3^* \geq \nu_1^*$, $\nu_3^* \geq \nu_2^*$ implies $\nu_3^* \geq \nu_{12} \Rightarrow T(\nu_3^*) = \nu_3^* \geq T(\nu_{12}) \Rightarrow \cdots \Rightarrow \nu_3^* \geq \nu_{12}^*$. The greatest lower bound of ν_1^* and ν_2^* can be constructed in an analogous way.

To show that $\nu_{\min}^*$ is the lowest fixed point, consider another fixed point ν^*, and note that $\nu^* \geq \nu_{\min} \Rightarrow T(\nu^*) = \nu^* \geq T(\nu_{\min}) \Rightarrow \cdots \Rightarrow \nu^* \geq \nu_{\min}^*$. Analogously, $\nu_{\max}^*$ is the highest fixed point of mapping T.

PROOF OF THEOREM 2:

Let $\nu^* = F^{-1}(\mu^*)$, i.e, the pre-network corresponding to μ^*, such that $F(\nu^*) = \mu^*$. Since $\nu_{\min}^*$ and $\nu_{\max}^*$ are the extreme fixed points of mapping T, $\nu_{\min}^* \leq \nu^* \leq \nu_{\max}^*$. Take any $a \in \bar{A}$ (the proof for the symmetric case $a \in \underline{A}$ is completely analogous). By definition of $\bar{A}$, node a can be connected only with downstream nodes: it can have customers but not suppliers. Hence, in any pre-network, every single arrow pointing to node a is an upstream arrow. Therefore, $\nu_{\min}^*(a)$, which is the set of arrows pointing to a in $\nu_{\min}^*$, is a superset of $\nu^*(a)$, which in turn is a superset of $\nu_{\max}^*(a)$. But one is always (weakly) better off choosing from a larger set, and so $Ch_a(\nu_{\min}^*(a)) = \mu_{\min}^*(a)$ is at least as good for a as $Ch_a(\nu^*(a)) = \mu^*(a)$, which in turn is at least as good as $Ch_a(\nu_{\max}^*(a)) = \mu_{\max}^*(a)$.

PROOF OF THEOREM 3:

The proof consists of two independent steps—one compares $\mu_{\max}^*$ with $\mu_{\max}'$ and the other compares $\mu_{\min}^*$ with $\mu_{\max}'$.

Step 1: Consider $\nu_{\max}^*$. Add node a' to market A, so that $U_{a'}(A) = \emptyset$. Let $\nu_+ = \nu_{\max}^* \cup \{\mathbf{r} : a' = d_\mathbf{r} \;\&\; \mathbf{c_r} \in Ch_a(\nu_{\max}^*(a) \cup \mathbf{c_r})\}$; that is, ν_+ contains all arrows in $\nu_{\max}^*$ plus all arrows $\mathbf{r}$ from nodes $a \in A$ to the new node a' such that a would like to add the attached contract $\mathbf{c_r}$ to its list of contracts (and possibly drop some of its other contracts). Now, note that $T(\nu_+) \geq \nu_+$ (for any $a \in A$, $\nu_+(a) = \nu^*_{\max}(a)$, and so all arrows in $T(\nu_+)$ originating from points in A are exactly the same as in ν_+; all new arrows originate from a' and thus necessarily point downstream). But then $\nu_+ \leq T(\nu_+) \leq \cdots \leq T^n(\nu_+) = T^{n+1}(\nu_+) \leq \nu'_{\max}$, where $\nu'_{\max} = F^{-1}(\mu'_{\max})$ is the highest fixed point in market A'. This, in turn, implies that for any $a \in \bar{A}$, $\nu^*_{\max}(a) = \nu_+(a) \supset \nu'_{\max}(a)$, and so a is at least as well off in $\mu^*_{\max} = Ch_a(\nu^*_{\max}(a))$ as in $\mu'_{\max} = Ch_a(\nu'_{\max}(a))$. Similarly, for any $a \in \underline{A}$, $\nu^*_{\max}(a) = \nu_+(a) \subset \nu'_{\max}(a)$, and so a is at most as well off in $\mu^*_{\max} = Ch_a(\nu^*_{\max}(a))$ as in $\mu'_{\max} = Ch_a(\nu'_{\max}(a))$.

Step 2: Now start with the larger market A' and consider the lowest fixed point of T, $\nu'_{\min}$. Exclude node $a' \in \bar{A}'$ with all the arrows going to and from a'. Denote the resulting pre-network on A by ν_-. Note that $T(\nu_-) \leq \nu_-$ (for any node $a \in A$, $U_a(\nu_-(a)) \subset U_a(\nu'_{\min}(a))$ and $D_a(\nu_-(a)) = D_a(\nu'_{\min}(a))$; thus (i) by same-side substitutability, the set of upstream arrows originating at a in $T(\nu_-)$ is a superset of the set of upstream arrows originating at a in $T(\nu'_{\min})$; and (ii) by cross-side complementarity, the set of downstream arrows originating at a in $T(\nu_-)$ is a subset of the set of downstream arrows originating at a in $T(\nu'_{\min})$). Therefore, $\nu_- \geq T(\nu_-) \geq \ldots\ T^n(\nu_-) = T^{n+1}(\nu_-) \geq \nu^*_{\min}$. This, in turn, implies that for any $a \in \bar{A}$, $\nu'_{\min}(a) = \nu_-(a) \subset \nu^*_{\min}(s)$, and so a is at least as well off in $\mu^*_{\min} = Ch_a(\nu^*_{\min}(a))$ as in $\mu'_{\min} = Ch_a(\nu'_{\min}(a))$. Similarly, for any $a \in \underline{A}$, $\nu'_{\min}(a) \supset \nu_-(a) \supset \nu^*_{\min}(a)$, and so a is at most as well off in $\mu^*_{\min} = Ch_a(\nu^*_{\min}(a))$ as in $\mu'_{\min} = Ch_a(\nu'_{\min}(a))$.

The case where a' is added to the other end of the market is completely symmetric.

PROOF OF THEOREM 4:

Since every chain is a tree, the set of tree-stable networks is a subset of the set of chain-stable networks. Let us now show that any chain-stable network is also tree stable.

Consider a network, μ^*, that is chain stable but not tree stable. Let τ be a tree with the smallest possible number of contracts blocking μ^*. Since, by assumption, τ is not a chain, there must exist a node, a, that is involved in at least two contracts in τ as a seller or in at least two contracts in τ as a buyer. Assume that a is involved in contracts $\{\mathbf{c}_1, \ldots, \mathbf{c}_k\} \subset \tau$ as a seller, $k \geq 2$; the case in which a is involved in two or more contracts as a buyer is completely symmetric and is therefore omitted. Let $v = Ch_a[\mu^*(a) \cup \mathbf{c}_1 \cup U_a(\tau(a))] \cap U_a(\tau(a))$, that is, the set of upstream contracts in blocking tree τ that a would choose to add to μ^* if the only additional downstream contract it had was $\mathbf{c}_1$. Note that, by same-side substitutability, $\mathbf{c}_1 \in Ch_a(\mu^*(a) \cup \mathbf{c}_1 \cup U_a(\tau(a)))$, and so $(\mathbf{c}_1 \cup v) \subset Ch_a(\mu^*(a) \cup (\mathbf{c}_1 \cup v))$. Set v can, of course, be empty. Let τ' be the subset of τ which consists of contracts that involve only the nodes that have paths connecting them to a and containing either $\mathbf{c}_1$ or a contract from v. In other words, τ' is obtained by cutting off the branches of tree τ (viewing a as the root) that do not start with contracts in $\mathbf{c}_1 \cup v$. By construction, τ' is a tree, $\tau'(a) = (\mathbf{c}_1 \cup v) \subset Ch_a(\mu^*(a) \cup \tau'(a))$, and for any other node b involved in τ', $\tau'(b) = \tau(b)$, and so $\tau'(b) \subset Ch_b(\mu^*(b) \cup \tau'(b))$. Therefore, τ' is a tree block of μ^*, and contains fewer contracts than τ does, which contradicts the assumption that τ is a tree with the smallest possible number of contracts blocking μ^*.

PROOF OF THEOREM 5:

Suppose network μ is in the weak core, but has a chain block, $(\mathbf{c}_1, \ldots, \mathbf{c}_k)$.

Let $(\mathbf{x}_1, \ldots, \mathbf{x}_m) \subset \mu$ be the longest chain in μ such that the seller in contract $\mathbf{c}_1$ is the buyer in contract $\mathbf{x}_m$, and let $(\mathbf{y}_1, \ldots, \mathbf{y}_n) \subset \mu$ be the longest chain in μ such that the buyer in contract $\mathbf{c}_k$ is the seller in contract $\mathbf{y}_1$. Let $\mu' = \{\mathbf{x}_1, \ldots, \mathbf{x}_m, \mathbf{c}_1, \ldots, \mathbf{c}_k, \mathbf{y}_1, \ldots, \mathbf{y}_n\}$, and let M be the set of nodes involved in μ'. Then μ' weakly dominates μ via coalition M, and hence μ could not be in the weak core. The proof of the fact that any network in the weak core is individually rational is very similar, and is therefore omitted.

Now consider any chain-stable network μ^* that is not in the weak core, and consider a network μ' that weakly dominates it via some coalition M and has the smallest possible number of contracts in $(\mu' \backslash \mu^*)$ among such networks. Take a node $a \in M$ that strictly prefers its set of contracts in μ' to its set of contracts in μ^* and doesn't have any upstream nodes that strictly prefer their sets of contracts in μ' to their sets of contracts in μ^*. $Ch_a(\mu^*(a) \cup \mu'(a)) \neq \mu^*(a)$. If $Ch_a(\mu^*(a) \cup \mu'(a)) \subset \mu^*(a)$, then μ^* is not individually rational, contradicting its chain stability. Otherwise, take contract $\mathbf{c}_1 \in Ch_a(\mu^*(a) \cup \mu'(a)) \backslash \mu^*(a)$. Contract $\mathbf{c}_1$ must be downstream for a, because a was chosen as one of the most upstream nodes that strictly benefit from a switch from μ^* to μ', and because all preferences are strict.

Let b be the buyer in contract $\mathbf{c}_1$. Preferences of agent b are strict, $\mu'(b) \neq \mu^*(b)$, and therefore $Ch_b(\mu^*(b) \cup \mu'(b)) \neq \mu^*(b)$. $Ch_b(\mu^*(b) \cup \mu'(b)) \not\subset \mu^*(b)$ by individual rationality, and so set $Z = Ch_b(\mu^*(b) \cup \mu'(b)) \backslash \mu^*(b)$ is not empty. There are three possibilities: (i) Z contains only $\mathbf{c}_1$, (ii) Z contains only some downstream contract $\mathbf{c}_2$, and (iii) Z contains $\mathbf{c}_1$ and a downstream contract $\mathbf{c}_2$. Let us consider these possibilities one by one.

(i) In this case, $(\mathbf{c}_1)$ is a chain block of μ^*.

(ii) In this case, consider network μ'' that includes contract $\mathbf{c}_2$, the longest possible chain in μ' that begins with $\mathbf{c}_2$, and the longest possible chain in μ^* that ends at node b. Then μ'' weakly dominates μ^* via the coalition of all nodes involved in μ'', and $|\mu'' \backslash \mu^*| < |\mu' \backslash \mu^*|$, contradicting the assumptions.

(iii) In this case, consider the buyer of contract $\mathbf{c}_2$ and repeat the same operation with this buyer as what we did with buyer b.

Eventually, since we keep going downstream, we will have to end up at case (i) or (ii) and so will either find a chain block of μ^*, or a network μ'' that weakly dominates μ^* such that $|\mu'' \backslash \mu^*| < |\mu' \backslash \mu^*|$, both of which are impossible by assumption.

Appendix B: An Example with Prices and Unit Identifiers

This Appendix presents an example of a setting with prices, unit identifiers, and contracts that link the suppliers of basic inputs with the consumers of final outputs both directly and through intermediaries.

There are five agents in the market: two suppliers (a_1 and a_2), one intermediary (b), and two consumers (c_1 and c_2). All agents have the same capacity: three units. For each supplier, it costs 0 to supply no units, 5 to supply one unit, 20 to supply two units, and 45 to supply three units of input, i.e., the marginal cost of the first unit is 5, the marginal cost of the second unit is 15, and the marginal cost of the third unit is 25. For each intermediary, it costs 0 to remain inactive, 5 to turn one unit of input into one unit of output, 15 to turn two units of inputs into two units of output, and 30 to turn three units of input into three units of output. An intermediary can freely dispose of unused inputs, but cannot produce more units of output than the number of units of input that he buys. Finally, each consumer has utility 0 from no consumption, 55 from consuming one unit of output, 100 from consuming two units, and 135 from consuming three or more units. A consumer can also buy units of input directly from suppliers, and turn them into units of output for his own consumption at the cost of 11 per unit. For example, his utility from buying one unit of output and two units of input is equal to $135 - 22 = 113$. All utilities are quasi-linear in money.

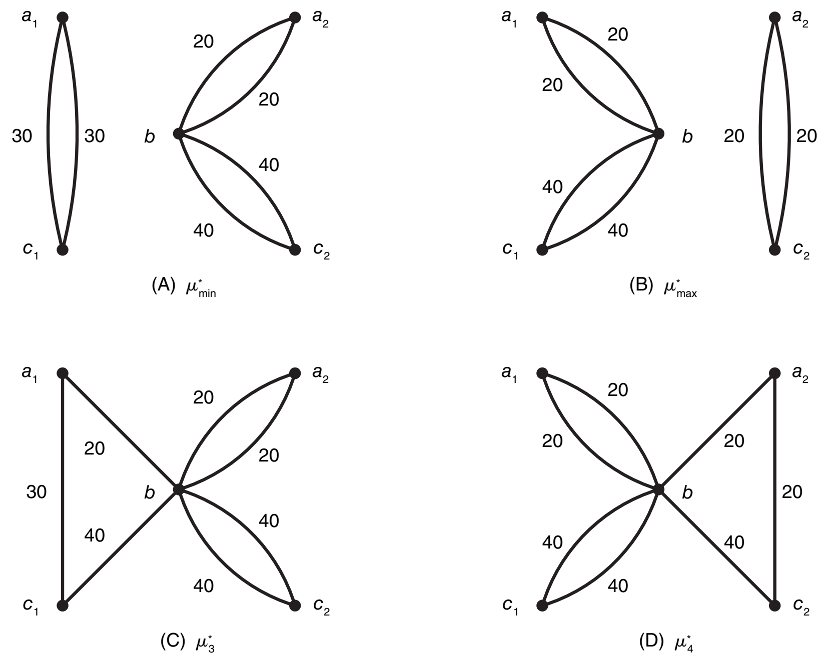

FIGURE B1. CHAIN-STABLE NETWORKS

There are also per-unit costs of transporting goods from sellers to buyers, summarized in the matrix below. Transportation costs are split equally between the seller and the buyer.

	b	c_1	c_2
a_1	1	3	2
a_2	2	4	3
b	—	2	1

There are six different prices at which trade is allowed to occur: 0, 10, 20, 30, 40, and 50. Each pair of agents who can trade, can trade up to three units of a good, with unit identifiers 1, 2, or 3. Agents' preferences over different units are lexicographic: if the numerical payoffs of a node from two different sets of contracts are the same, a node prefers the set with smaller unit identifiers.

It is easy to check that the nodes' preferences satisfy same-side substitutability and cross-side complementarity, and therefore the set of chain-stable networks has to be non-empty. Indeed, Figures B1(A) and B1(B) illustrate, respectively, the supplier- and the consumer-optimal chain-stable networks $\mu^*_{\min}$ and $\mu^*_{\max}$. In $\mu^*_{\min}$, supplier a_1 sells two units of input to consumer c_1 at the price of 30 for each unit, supplier a_2 sells two units of input to intermediary b at the price of 20 for each unit, and intermediary b sells two units of output to consumer c_2 at the price of 40 for each unit. In $\mu^*_{\max}$, supplier a_2 sells two units of input to consumer c_2 at the price of 20 for each unit, supplier a_1 sells two units of input to intermediary b at the price of 20 for each unit, and intermediary b sells two units of output to consumer c_1 at the price of 40 for each unit. More formally, the supplier-optimal chain-stable network consists of contracts $(a_1, c_1, 1, 30)$, $(a_1, c_1, 2, 30)$,

$(a_2, b, 1, 20)$, $(a_2, b, 2, 20)$, $(b, c_2, 1, 40)$, and $(b, c_2, 2, 40)$ and the consumer-optimal chain-stable network consists of contracts $(a_1, b, 1, 20)$, $(a_1, b, 2, 20)$, $(b, c_1, 1, 40)$, $(b, c_1, 2, 40)$, $(a_2, c_2, 1, 20)$, and $(a_2, c_2, 2, 20)$. These networks were obtained by running the *T*-algorithm starting from the lowest and the highest pre-networks; intermediate steps are available upon request. Some other stable networks for this market are shown in Figures B1(C) and B1(D).

REFERENCES

Abdulkadiroğlu, Atila, Parag A. Pathak, and Alvin E. Roth. 2005. "The New York City High School Match." *American Economic Review*, 95(2): 364–67.

Abeledo, Hernan, and Garth Isaak. 1991. "A Characterization of Graphs that Ensure the Existence of Stable Matchings." *Mathematical Social Sciences,* 22(1): 93–96.

Adachi, Hiroyuki. 2000. "On a Characterization of Stable Matchings." *Economics Letters*, 68(1): 43–49.

Alkan, Ahmet. 1988. "Nonexistence of Stable Threesome Matchings." *Mathematical Social Sciences,* 16(2): 207–09.

Alkan, Ahmet, and David Gale. 2003. "Stable Schedule Matching under Revealed Preference." *Journal of Economic Theory*, 112(2): 289–306.

Baïou, Mourad, and Michel Balinski. 2002. "The Stable Allocation (or Ordinal Transportation) Problem." *Mathematics of Operations Research,* 27(3): 485–503.

Bajari, Patrick, and Jeremy T. Fox. 2005. "Complementarities and Collusion in an FCC Spectrum Auction." National Bureau of Economic Research Working Paper 11671.

Blair, Charles. 1988. "The Lattice Structure of the Set of Stable Matchings with Multiple Partners." *Mathematics of Operations Research,* 13(4): 619–28.

Boyd, Donald, Hamilton Lankford, Susanna Loeb, and James Wyckoff. 2003. "Analyzing the Determinants of the Matching Public School Teachers to Jobs: Estimating Compensating Differentials in Imperfect Labor Markets." National Bureau of Economic Research Working Paper 9878.

Crawford, Vincent P., and Elsie Marie Knoer. 1981. "Job Matching with Heterogeneous Firms and Workers." *Econometrica*, 49(2): 437–50.

Echenique, Federico, and Jorge Oviedo. 2004. "Core Many-to-One Matchings by Fixed-Point Methods." *Journal of Economic Theory*, 115(2): 358–76.

Echenique, Federico, and Jorge Oviedo. 2006. "A Theory of Stability in Many-to-Many Matching Markets." *Theoretical Economics*, 1(2): 233–73.

Gale, David, and Lloyd S. Shapley. 1962. "College Admissions and the Stability of Marriage." *American Mathematical Monthly,* 69(1): 9–15.

Gale, David, and Marilda A. O. Sotomayor. 1985. "Ms. Machiavelli and the Stable Matching Problem." *American Mathematical Monthly,* 92(4): 261–68.

Hatfield, John William, and Paul R. Milgrom. 2005. "Matching with Contracts." *American Economic Review*, 95(4): 913-35.

Immorlica, Nicole, and Mohammad Mahdian. 2005. "Marriage, Honesty, and Stability." In *Proceedings of the Sixteenth Annual ACM-SIAM Symposium on Discrete Algorithms*, 53–62. Philadelphia: SIAM.

Kelso, Alexander S., Jr., and Vincent P. Crawford. 1982. "Job Matching, Coalition Formation, and Gross Substitutes." *Econometrica*, 50(6): 1483–1504.

Kojima, Fuhito, and Parag A. Pathak. 2007. "Incentives and Stability in Large Two-Sided Matching Markets." http://www.people.fas.harvard.edu/~kojima/largematchingmarket.pdf.

Konishi, Hideo, and M. Utku Ünver. 2006. "Credible Group Stability in Many-to-Many Matching Problems." *Journal of Economic Theory*, 129(1): 57–80.

Roth, Alvin E. 1982. "The Economics of Matching: Stability and Incentives." *Mathematics of Operations Research,* 7(4): 617–28.

Roth, Alvin E. 1984. "Stability and Polarization of Interests in Job Matching." *Econometrica*, 52(1): 47–57.

Roth, Alvin E. 1985. "Conflict and Coincidence of Interest in Job Matching: Some New Results and Open Questions." *Mathematics of Operations Research,* 10(3): 379–89.

Roth, Alvin E. 2007. "Deferred Acceptance Algorithms: History, Theory, Practice, and Open Questions." National Bureau of Economic Research Working Paper 13225.

Roth, Alvin E, and Elliott Peranson. 1999. "The Redesign of the Matching Market for American Physicians: Some Engineering Aspects of Economic Design." *American Economic Review*, 89(4): 748–80.

Roth, Alvin E, and Marilda A. O. Sotomayor. 1990. *Two-Sided Matching: A Study in Game-Theoretic Modeling and Analysis*. New York: Cambridge University Press.

Shapley, Lloyd S. 1962. "Complements and Substitutes in the Optimal Assignment Problem." *Naval Research Logistics Quarterly,* 9(1): 45–48.

Smith, Adam. 2003. *An Inquiry into the Nature and Causes of the Wealth of Nations*. New York, NY: Bantam Classic (Org. pub. 1776).

Sørensen, Morten. 2007. "How Smart Is Smart Money? A Two-Sided Matching Model of Venture Capital." *Journal of Finance*, 62(6): 2725–62.

Sotomayor, Marilda. 1999. "Three Remarks on the Many-to-Many Stable Matching Problem." *Mathematical Social Sciences,* 38(1): 55–70.

Sun, Ning, and Zaifu Yang. 2006. "Equilibria and Indivisibilities: Gross Substitutes and Complements." *Econometrica*, 74(5): 1385–1402.