Trading Networks with Price-Setting Agents☆

Lawrence E. Blume[a], David Easley[b], Jon Kleinberg[c,1], Éva Tardos[d]

[a] *Economics Department, Cornell University, Ithaca, NY, and The Santa Fe Institute*

[b] *Economics Department, Cornell University, Ithaca, NY*

[c] *Dept. of Computer Science, Cornell University, Ithaca, NY*

[d] *Dept. of Computer Science, Cornell University, Ithaca, NY,*

Abstract

In a wide range of markets, individual buyers and sellers trade through intermediaries, who determine prices via strategic considerations. Typically, not all buyers and sellers have access to the same intermediaries, and they trade at correspondingly different prices that reflect their relative amounts of power in the market. We model this phenomenon using a game in which buyers, sellers, and traders engage in trade on a graph that represents the access each buyer and seller has to the traders. We show that the resulting game always has a subgame perfect Nash equilibrium, and that all equilibria lead to an efficient allocation of goods. Finally, we analyze how the profits obtained by the traders depend on the underlying graph — roughly, a trader can command a positive profit if and only if it has an "essential" connection in the network structure, thus providing a graph-theoretic basis for quantifying the amount of competition among traders.

☆A short version of this paper appeared in the EC'07 Conference Proceedings. This work has been supported in part by NSF grants CCF-0325453, IIS-0329064, CNS-0403340, and BCS-0537606, ONR grant N00014-98-1-0589, a Google Research Grant, a Yahoo! Research Alliance Grant, the Institute for the Social Sciences at Cornell, and the John D. and Catherine T. MacArthur Foundation. We would like to thank GEB Guest Editors David Parkes and Moshe Tennenholtz and three referees for helpful comments and suggestions.

Email addresses: `lb19@cornell.edu` (Lawrence E. Blume), `dae3@cornell.edu` (David Easley), `kleinber@cs.cornell.edu` (Jon Kleinberg), `eva@cs.cornell.edu` (Éva Tardos)

[1] corresponding author

Key words: Networks, trade, price formation
JEL classification: C72, L14

1. Introduction

The baseline analysis of buyers and sellers interacting in markets is the Walrasian model, wherein trade between anonymous buyers and sellers takes place at a single market clearing price. This reduced form view of trade is a powerful model which has led to many insights. However, it does not model how a single market price emerges from the behavior of market participants. The Walrasian model ignores the variety of market institutions and customs, and how traders act given the rules of the game, to agree upon prices at which to trade.

Three facts are common to a variety of market institutions: Individual buyers and sellers often trade through intermediaries, not all buyers and sellers have access to the same intermediaries, and not all buyers and sellers trade at the same price. Two important, and quite different, examples are trade of agricultural goods in developing countries and trade of financial assets.

Consider, for instance, the petty trade of agricultural goods in developing countries. Given inadequate transportation networks, and poor farmers' limited access to capital, many farmers have no alternative to trading with middlemen in inefficient local markets. A developing country may have many such partially overlapping markets existing alongside modern efficient markets (Barrett and Mutambatsere, 2008). Goods flow through a network, from the producer to a trader in one market, through that trader to a second market (or farther) where they finally meet consumers. Some buyers and sellers have links to several traders, while others may be forced to deal with only one trader. In general, most traders do much repeated business with a small clientele, very few traders are "wholesale only", and goods moving from one market to another will move through only a few hands (Fafchamps and Gabre-Madhin, 2007).

In financial markets, much of the trade between buyers and sellers is also conducted through intermediaries, such as brokers, market makers and electronic trading systems. Markets for actively-traded assets are global; there is no one market. Trade in a given asset may occur simultaneously on the floor of an exchange, on crossing networks, on electronic exchanges, and in markets in other countries. Some buyers and sellers have access to many or all of these trading venues; others have access to only one or a few of them. Each individual market, such as the NYSE, NASDAQ or London Stock Exchange, consists of many, densely connected traders. These highly connected markets are themselves linked, though less densely, through traders who arbitrage across markets. One of the most striking examples of this phenomenon occurs in the market for foreign exchange, where there is an interbank market with restricted access and a retail market with much more open access.

In this paper, we develop a framework in which market trade is mediated through middlemen, and in which the flow of goods is constrained by a preexisting network of market relationships. The market is described by a network whose edges represent the direct access that different market participants have to one another. We model trading networks as tripartite graphs, in which distinct types of vertices represent buyers, sellers, and traders. Edges connect buyers and sellers to traders. They represent the direct access market participants have to one another. In principles, such a network model can also contain edges that connect traders to other traders, although we do not consider this here. Networks for different kinds of commodities can be quite different. Networks for petty trade in vegetables, root crops and the like are, as we observed above, fairly sparse, and dominated by low-degree nodes. In a network of financial markets, the discrete markets comprising it are densely connected, with fewer links between them. The degree distribution is highly dispersed. Some trading firms are active on many markets, and may represent many clients, and so their degree is high. Other firms may specialize in only a few assets and trade only their own accounts, and so their degrees are low.

Prices in the markets we study come from the interaction between buyer or seller and the intermediary. In petty trading markets, the price is often the outcome of bargaining between buyer or seller and trader. Financial markets contain a variety of intermediation schemes. In some markets, market makers post bid and ask prices for sellers and buyers, respectively. In US financial markets, new issues are often introduced through auctions. In some markets the intermediary is a software agent that sets prices to clear the market. In most cases, the intermediary makes his profit off the spread, the difference between bid and ask, the buy and the sell price, or (and it amounts to the same thing) fixed per-transaction fees. Spreads for a given asset can differ significantly across markets, depending upon their thickness, characteristics of participants, and upon the rules of trade.

We model trade as a two-stage, complete-information game.[2] Traders strategically choose bid and ask prices to offer to the sellers and buyers to whom they are connected. The sellers and buyers respond by choosing with whom to trade, or not to trade at all. The network encodes the relative power in the structural positions of the market participants, including the implicit levels of competition among traders. We show that this game always has a subgame perfect Nash equilibrium, and that all equilibria lead to an efficient (i.e. socially optimal) allocation of goods. In particular, the market enables the "right" set of people to get the good, subject to the network constraints. We also analyze the division of the surplus from trade — how, in particular, trader profits depend on the network structure.

Our work here is connected to several lines of research in economics, finance, and algorithmic game theory. At a general level, our approach can be viewed as synthesizing two important strands of work: one that considers price-setting intermediaries, but without network-type constraints on who can trade with whom; and another strand that treats buyer-seller interaction using network structures, but without attempting to model the processes by which prices are actually formed.

The study of brokers, intermediaries and middlemen is common to many areas of economics, including finance (O'Hara, 1995), industrial organization (Hall and Rust, 2000; Lamoreaux and Sokoloff, 2002; Rubinstein and Wolinsky, 1987; Spulber, 1999) international finance (Krishna, Imai, Mukhopadhyay, and Tan, 2004) labor economics (Edid, 1994; Yavas, 1994) and macroeconomics (Hellwig, 2003; Li, 1998). Most of this literature is concerned with the role of intermediaries in facilitating or blocking information flow, their role in transactions cost reduction, and the rents they capture through management of the trading process. None of this research is concerned with the participation constraints created by the network structure of markets.

The computer science literature has taken seriously the effects of network structure on market prices. Kakade et al (2004) have characterized competitive

[2]It is important for our analysis that traders know the values that buyers and sellers place on the goods. So our model is best thought of as applying to settings in which the traders have experience in trading with these buyers and sellers. Buyers and sellers, on the other hand, do not need to know each others values. It is enough for them to be able to observe the prices quoted by traders.

equilibrium prices when buyer-seller interaction is mediated by a network. Even-Dar et al (2007) study the strategic aspects of network formation when prices arise from competitive equilibrium. Babaioff et al (2005), and Chu and Shen (2006) examine mechanism design issues for effecting trade when market participants are connected through a network. Sarma, Chakrabarty, and Gollapudi (2007) provide an algorithm that computes a Nash equilibrium for a related game of pricing and trade on a network of buyers and sellers.

The classic results of Shapley and Shubik (1972) on the assignment game can be viewed as studying the result of trade on a bipartite graph in terms of the core. They study the dual of a linear program based on the matching problem, similar to what we use for a reduced version of our model in the next section, but their concern is not with the mechanisms of price formation in markets. Most importantly, we have prices set strategically by traders. Leonard (1983) studies VCG prices for the assignment problem, Crawford and Knoer (1981) and Kelso and Crawford (1982) study dynamic auction mechanisms in this setting, Demange, Gale, and Sotomayor (1986) provides a general analysis of dynamic procedures for the setting, and Bikhchandani and Ostroy (2002) generalizes this analysis to the assignment of packages of goods.

Our work is also related to two further recent papers. Kranton and Minehart (2001) analyze the prices at which trade occurs in a network. They use a bipartite graph with direct links between buyers and sellers, and then use an ascending auction mechanism, rather than strategic intermediaries, to determine the prices. Their auction has desirable equilibrium properties, but as Kranton and Minehart note it is an abstraction of how goods are allocated and prices are determined that is similar in spirit to the Walrasian auctioneer abstraction. In fact, as we discuss in the next section, there is a network structure which implements the Kranton-Minehart outcome in our complete-information game; that is, traders produce prices at equilibrium matching the prices produced by their auction mechanism.[3]

Gale and Kariv (2007) analyze a network model of exchange in which trade is intermediated. They construct an intertemporal bargaining model in which

[3]Kranton and Minehart, however, can also analyze a more general setting in which buyers' values are private and thus buyers and sellers play a game of incomplete information. We deal only with complete information.

in each period a trader who owns a unit of the good is chosen at random to propose a trade to someone with whom he shares a connection and who does not own a unit of the good. Gale and Kariv show that as the period length converges to zero the equilibrium of their trading game becomes efficient. This approach differs from ours in that our buyers and sellers are not directly connected; instead, in our model, all trade goes through intermediaries who actually set bid and ask prices.

2. The Basic Model

Our goal in formulating the model is to express the process of price-setting in markets such as those discussed above, where the participants do not all have uniform access to one another. We are given a set B of buyers, a set S of sellers, and a set T of traders. There is an undirected graph G that indicates who is able to trade with whom. All edges have one end in $B \cup S$ and the other in T; that is, each edge has the form (i, t) for $i \in S$ and $t \in T$, or (j, t) for $j \in B$ and $t \in T$. This reflects the constraints that all buyer-seller transactions go through traders as intermediaries.

In the basic version of the model, we consider identical goods, one copy of which is initially held by each seller. We will subsequently generalize this to a setting with many distinct commodities. Buyers can value different commodities differently, and potentially sellers can value transactions with different buyers differently as well. Having different buyer valuations captures settings like house purchases; adding different seller valuations as well captures matching markets — for example, sellers as job applicants and buyers as employers, with both caring about who ends up with which "good" (and with traders acting as services that broker the job search).

The basic model has a single type of *good*, which comes in individual units. Each seller initially holds one unit of the good. All three agent types have utility functions which are linear in money. (Without loss of generality, we can take the marginal utility of money to be 1 for all agents.) Each agent $i \in B \cup S$ additionally receives θ_i units of utility from one copy of the good. No agent wants more than one copy of the good, so additional copies are valued at 0. All commodity endowments and valuations are common knowledge. Each agent has an initial endowment of money that is larger than its own individual valuation θ_i; the effect of this is to guarantee that any buyer who ends up without a copy of the good has been priced out of the market due to its valuation and network position, not a lack of funds.

We picture each good that is sold flowing along a sequence of two edges; from a seller to a trader, and then from the trader to a buyer. The particular way in which goods flow is determined by the following game: First, each trader offers a bid price to each seller it is connected to, and an ask price to each buyer it is connected to. Sellers and buyers then choose from among the offers presented to them by traders. Finally, each trader buys a copy of the good from each seller that accepts its offer, and it sells a copy of the good to each buyer that accepts its offer. If a particular trader t finds that more buyers than sellers accept its offers, then it has committed to provide more copies of the good than it has received, and we will say that this results in a large penalty to the trader for defaulting; the effect of this is that in equilibrium, no trader will choose bid and ask prices that result in a default.

More precisely, a strategy for each trader t is a specification of a *bid price* β_{ti} for each seller i to which t is connected, and an *ask price* α_{tj} for each buyer j to which t is connected.[4] Each seller or buyer then chooses at most one incident edge, indicating the trader with whom they will transact, at the indicated price. (The choice of a single edge reflects the facts that (a) sellers each initially have only one copy of the good, and (b) buyers each only want one copy of the good.) The payoffs are as follows:

1. For each seller i, the payoff from selecting trader t is β_{ti}, while the payoff from selecting no trader is θ_i. (In the former case, the seller receives β_{ti} units of money, while in the latter it keeps its copy of the good, which it values at θ_i.)
2. For each buyer j, the payoff from selecting trader t is $\theta_j - \alpha_{tj}$, while the payoff from selecting no trader is 0. (In the former case, the buyer receives the good but gives up α_{tj} units of money.)
3. For each trader t, with accepted offers from sellers $i_1, \ldots, i_s$ and buyers $j_1, \ldots, j_b$, the payoff is $\sum_r \alpha_{tj_r} - \sum_r \beta_{ti_r}$, minus a penalty π if $b > s$. The penalty is chosen to be large enough that a trader will never incur it in equilibrium, and hence we will generally not be concerned with the penalty.

[4]We can also handle a model in which a trader may choose not to make an offer to certain of its adjacent sellers or buyers.

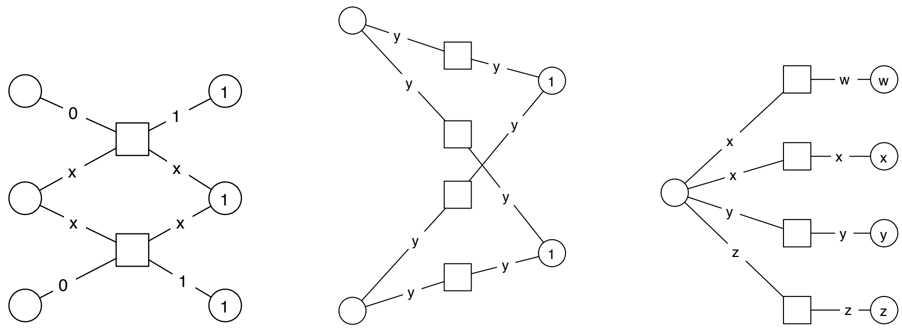

(a) *Heterogeneous outcomes* (b) *Implicit perfect competition* (c) *Auction*

Figure 1: (a) A network in which the middle seller and buyer benefit from perfect competition between the traders, while the other sellers and buyers have no power due to their position in the network. (b) A form of *implicit perfect competition*: all bid/ask spreads will be zero in equilibrium, even though no trader directly "competes" with any other trader for the same buyer-seller pair. (c) An auction, mediated by traders, in which the buyer with the highest valuation for the good ends up with it.

This defines the basic elements of the game. The equilibrium concept we use is subgame perfect Nash equilibrium. An equilibrium specifies a set of bids and asks for each trader, and also a choice by each buyer and seller of whom (if anyone) to trade with. Thus implicit in equilibrium is a set of paths along which goods flow.

2.1. Some Examples

To help with thinking about the model, we now describe three illustrative examples, depicted in Figure 1. To keep the figures from getting too cluttered, we adopt the following conventions: sellers are drawn as circles in the leftmost column and will be named $i_1, i_2, \ldots$ from top to bottom; traders are drawn as squares in the middle column and will be named $t_1, t_2, \ldots$ from top to bottom; and buyers are drawn as circles in the rightmost column and will be named $j_1, j_2, \ldots$ from top to bottom. All sellers in the examples will have valuations for the good equal to 0; the valuation of each buyer is drawn inside its circle; and the bid or ask price on each edge is drawn on top of the edge.

In Figure 1(a), we show how nodes with different positions in the network topology can achieve different payoffs, even when all buyer valuations are the same numerically. Specifically, seller i_2 and buyer j_2 occupy powerful positions, because the two traders are competing for their business; on the other hand, the other sellers and buyers are in weak positions, because they each have only one option. And indeed, in every equilibrium, there is a real number $x \in [0,1]$ such that both traders offer bid and ask prices of x to i_2 and j_2 respectively, while they offer bids of 0 and asks of 1 to the other sellers and buyers. Thus, this example illustrates a few crucial ingredients that we will identify at a more general level shortly. Specifically, i_2 and j_2 experience the benefits of *perfect competition*, in that the Bertrand-style competition between the two traders competing for their business drives the bid-ask spreads to 0. On the other hand, the other sellers and buyers experience the downsides of *monopoly* — they receive 0 payoff since they have only a single option for trade, and the corresponding trader makes all the profit. Note further how this natural behavior emerges from the fact that traders are able to offer different prices to different agents — capturing the fact that there is no one fixed "price" in the kinds of markets that motivate the model, but rather different prices reflecting the relative power of the different agents involved.

This first example shows perhaps the simplest way in which a trader's profit on a particular transaction can drop to 0: when there is another trader who can replicate its function precisely. (In this example, two traders each had the ability to move a copy of the good from i_2 to j_2.) But as our subsequent results will show, traders make zero profit more generally due to global, graph-theoretic reasons. The example in Figure 1(b) gives an initial indication of this: one can show that for every equilibrium, there is a $y \in [0,1]$ such that every bid and every ask price is equal to y (and any y in the interval describes one such equilibrium). In other words, all traders make zero profit, whether or not a copy of the good passes through them — and yet, no two traders have any seller-buyer paths in common. The price spreads have been driven to zero by a global constraint imposed by the long cycle through all the agents; this is an example of *implicit perfect competition* determined by the network topology.[5]

In Figure 1(c), we provide an example of a network which implements the standard second-price auction outcome.[6] Suppose the buyer valuations from top

[5]Notice that in this equilibrium buyers and sellers are indifferent about which offers to accept. Not all best response flow-paths of goods will be equilibria.

[6]We restrict attention to complete information, so this is not the usual incomplete information auction setting.

to bottom are $w > x > y > z$. The bid and ask prices shown are consistent with an equilibrium in which i_1 and j_1 accept the offers of trader t_1, and no other buyer accepts the offer of its adjacent trader: thus, trader t_1 receives the good with a bid price of x, and makes $w - x$ by selling the good to buyer j_1 for w. In this way, the network works like an auction for a single good in which the traders act as "proxies" for their adjacent buyers. The buyer with the highest valuation for the good ends up with it, and the surplus is divided between the seller and the associated trader. Note that one can construct a k-unit auction with $\ell > k$ buyers just as easily, by building a complete bipartite graph on k sellers and ℓ traders, and then attaching each trader to a single distinct buyer. Buyers are each connected to a single trader, so asks will be at buyer valuation. It is not hard to see the bids of the agents that trade will be at the $k+1$st buyer's valuation, as the trader associated with this buyer can take over a lower bid. Notice that the second-price outcome is not robust to the network structure. Suppose, for instance, that the network in Figure 1(c) is modified so that the topmost trader sells to two buyers, and that those buyers have the two highest valuations. The equilibrium asks remain unchanged, so the seller receives the third highest price.

2.2. Extending the Model to Distinguishable Goods

We extend the basic model to a setting with distinguishable goods, as follows. Instead of having each agent $i \in B \cup S$ have a single numerical valuation θ_i, we index valuations by pairs of buyers and sellers: if buyer j obtains the good initially held by seller i, it gets a utility of θ_{ji}, and if seller i sells its good to buyer j, it experiences a loss of utility of θ_{ij}. This generalizes the case of indistinguishable goods, since we can always have these pairwise valuations depend only on one of the indices. A strategy for a trader now consists of offering bids (respectively, asks) in the form of vectors, essentially specifying a "menu" with a price attached to each buyer (resp. seller).[7] Specifically, a trader t can offer a buyer j a menu of asks α_{tji}, a vector of values for all the products that she is connected to, where α_{tji} is the ask for the product of seller i. Symmetrically, a trader t can offer to each seller i a menu of bids β_{tij} for selling to different buyers j. Each buyer and seller selects an offer from an adjacent trader, and the payoffs to all agents are determined as before.

[7] We can also handle a model in which a trader offers a bid to each seller that specifies both a price *and* a buyer, and offers an ask to each buyer that specifies both a price and a seller

This general framework captures matching markets as considered by Shapley and Shubik (1972) and Leonard (1983): for example, a job market that is mediated by agents or employment search services (as in hiring for corporate executives, or sports or entertainment figures). Here the sellers are job applicants, buyers are employers, and traders are the agents that mediate the job market. Of course, if one specifies pairwise valuations on buyers but just single valuations for sellers, we model a setting where buyers can distinguish among the goods, but sellers don't care whom they sell to – this (roughly) captures settings like housing markets.

2.3. Results

In this section we briefly describe our results. We will defer details to the next two sections. Our results are essentially that there is always an equilibrium; that in equilibrium the trade is socially optimal given the constraints on trade imposed by the network structure; and that a trader can make a profit if and only if he has an incident edge in the trading graph which is essential for social welfare.

First we must define what we mean by an equilibrium. A *no-crossing* strategy for trader t is a vector of bids and asks such that for each ij pair, $\beta_{tij} \leq \alpha_{tij}$. An *equilibrium* is a subgame-perfect Nash equilibrium in which each trader uses a no-crossing strategy. The no-crossing restriction is natural because to employ a crossing strategy is to propose to execute a trade which will lose money. All of our results, however, hold for all subgame perfect Nash equilibria.

Second, we define what it means for trade to be socially optimal. An optimal trade is the solution of the transportation problem: sending goods from sellers to buyers connected to them via a traders, in order to maximize the social value of trade. Our first result is that in any equilibrium of our trading game the trade is efficient. We provide a proof of this result in Section 4.1.

Theorem 1. *Every equilibrium results in an efficient allocation of the goods.*

In fact this conclusion holds for any subgame perfect Nash equilibrium, and not just those using no-crossing strategies.

Next we claim that an equilibrium exists. In fact, we give two constructions for finding an equilibrium. In the first construction, traders may make crossing bid-ask offers. We then give a second construction showing that crossing

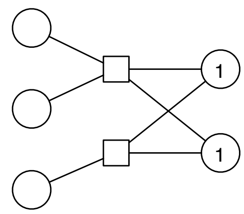

Figure 2: The top trader is essential for social welfare. Yet the only equilibrium is to have bid and ask values equal to 0, and the trader makes no profit.

offers are not essential to the existence of equilibrium. We prove these existence results in Section 4.1 as well.

Theorem 2. *There exists a pure equilibrium supporting any socially optimal trade.*

Finally, we turn to the question of how a trader's profit is affected by its position in the network. As Figures 1(a) and 1(b) suggest, a trader cannot make a profit unless it has some type of monopoly. We will show that in the special case of pair-traders, when each trader is associated only with a single seller-buyer pair (thus, traders in this case essentially serve as a "trade route" between the two), a trader can make profit if and only if he is essential for social welfare. Interestingly, in the general case, not all traders with monopoly power can make profit, as shown by the example in Figure 2. We claim that the only equilibrium in this example is to have all asks and all bids equal to 0. Clearly all bids will be 0 as the sellers are each connected to only one trader. To see why the asks must also be 0, first note that both traders must offer the same ask for both products, as otherwise the one with higher ask would get undercut by the other. Further, in any arrangements of trades, there is a trader who has access to an unsold good. If the common ask is $\alpha > 0$, the trader who has access to the unsold good can undercut the ask α.[8]

[8]If any edge connecting trader 1 to any buyer or seller, is deleted, then both traders can make a profit of 1. This example also shows that if we extend the game to have the traders buy a network first, then the resulting network will not be efficient even if the edges are very cheap: If we modify values so that keeping that edge (t_1, b_2) improves social welfare a bit (for example, if seller i_3 has a tiny non-zero value for the good), trader t_1 will want to drop the edge to improve his profit, but this now hurts social welfare. This shows some of the fundamental contrasts between our model

We can characterize when there exists an equilibrium in which trader i makes a positive profit. This is the subject of the next theorem, which we prove in Section 4.2.

Theorem 3. *There is an equilibrium in which trader t makes a positive profit if and only if there is a seller or buyer i adjacent to t such that the connection of trader t to agent i is essential for social welfare — that is, if deleting the edge connecting t and i decreases the value of the optimal allocation.*

Overview of the Analysis.. To prove these theorems we use a description of optimal trade in terms of linear programming. We consider the bipartite graph $G = (S \cup B, E)$ connecting sellers and buyers where an edge $e = (i, j)$ connects a seller i and a buyer j if there is a trader adjacent to both. For each buyer or seller i, we use $adj(i) \subset T$ to denote the set of traders adjacent to i. Then the set of edges in the bipartite graph connecting buyers and sellers is

$$E = \{(i, j) : adj(i) \cap adj(j) \neq \emptyset\}.$$

On this graph, we then solve the following transshipment-style linear programming problem in which the value of edge (i, j) is equal to $\theta_{ji} - \theta_{ij}$ (since the value of trading between i and j is independent of which trader conducted the trade). An integral solution to this program corresponds to a socially optimal trade, where for an edge $e = (i, j)$ with $x_e = 1$, trade between seller i and buyer j can be conducted by any trader t adjacent to both.

$$\begin{aligned} \max SV(x) &= \sum_{e \in E : e=(i,j)} x_e(\theta_{ji} - \theta_{ij}) \\ x_e &\geq 0 \quad \forall e \in E \\ \sum_{e \in adj(i)} x_e &\leq 1 \quad \forall i \in S \\ \sum_{e \in adj(j)} x_e &\leq 1 \quad \forall j \in B \end{aligned} \tag{1}$$

and that of Kranton and Minehart (2001), who show that if all edges are sufficiently cheap then the network-formation version of their game will result in an efficient network.

The idea of the proof is to show that equilibria form optimal primal and dual solutions of this linear program. To do this, we need to think of the equilibria in terms of the profits the players obtain, rather than the bid and ask values. This is in fact necessary, as the equilibrium bids and asks may not form convex sets. To see an example of such a non-convex bid set, consider the auction example in Figure 1(c) with valuations $w > x > y > z$. For the set of bids, the values in the figure give only one possible equilibrium; in all equilibria the asks will be w, x', y', z' for x', y', z' at least x, y, z respectively, since the buyers are monopolized. The asks of w, x', y', z' together with any bid values $x'', y'', z'' \leq w$ for the bottom three traders results in an equilibrium with trader t_1 offering the bid $\max(x'', y'', z'')$, assuming this maximum is at least x. This set of bids does not form a convex set. In this equilibrium any trader other than t_1 offering a bid above x is making crossing offers; non-crossing equilibria can also create non-convex sets using a slightly more complex example.

We will think of the dual of the linear program (1) in terms of the profits of the players. The constraints above are associated with buyers and sellers, so dual variables can be naturally associated with their profits. However, in this linear program there are no dual variables associated with traders, while our main interest is exactly in the price-setting behavior and profit of the traders. In Section 3 we will consider the special case of pair-traders, where edges e in the linear program directly correspond to traders, and we can add an additional constraint $x_e \leq 1$, and think of the corresponding dual variable as the trader's profit. However, when traders are connected to many buyers and sellers, then each trader is solving a more complicated optimization problem. To see the connection between the traders' profit and the linear programming dual of (1) in this case, we have to understand the trader's optimization problem better, as we'll see in Section 4.

3. Markets with Pair-Traders

The ideas behind the analysis of the general model are easier to explain if we restrict our attention to networks with traders that we refer to as *pair-traders*. In this case, each trader is connected to just one buyer and one seller. The techniques we develop to handle this case form a useful basis for reasoning about the case of traders that may be connected arbitrarily to the sellers and buyers.

We will relate profits in a subgame perfect Nash equilibrium to optimal solutions of a linear program, use this relation to show that all equilibria result in efficient allocation of the goods, and show that an equilibrium in pure strategies always exists. First, we consider the simplest model where sellers have indistinguishable items, and each buyer is interested in getting one item. Then we extend the results to the more general case of a matching market, as discussed in the previous section, where valuations depend on the identity of the seller and buyer.

Given that we are working with pair-traders in this section, we can represent the problem using a bipartite graph G whose node set is $B \cup S$, and where each trader t, connecting seller i and buyer j, appears as an edge $t = (i, j)$ in G. Note, however, that we allow multiple traders to connect the same pair of agents.

3.1. Indistinguishable Goods

The socially optimal trade for the case of indistinguishable goods is the solution of the transportation problem: sending goods along the edges representing the traders. The edges along which trade occurs correspond to a matching in this bipartite graph, and the optimal trade is described by the linear program (1) where we can index variables directly with the traders, and the objective function simplifies to the following.

$$\max SV(x) \quad = \sum_{t \in T : t=(i,j)} x_t(\theta_j - \theta_i) \tag{2}$$

Next we consider an equilibrium. Each trader $t = (i, j)$ must offer a bid β_t and an ask α_t. (We omit the subscript denoting the seller and buyer since we are dealing with pair-traders.) Given the bid and ask price, the agents react to these prices, as described earlier. Instead of focusing on prices, we will focus on profits. If a seller i sells to a trader $t \in adj(i)$ with bid β_t then his profit is $p_i = \beta_t - \theta_i$. Similarly, if a buyer j buys from a trader $t \in adj(j)$ with ask α_t, then his profit is $p_j = \theta_j - \alpha_t$. Finally, if a trader t trades with ask α_t and bid β_t then his profit is $y_t = \alpha_t - \beta_t$. All agents not involved in trade make 0 profit. We will show that the profits at equilibrium are an optimal solution to the following linear program.

$$\begin{aligned}
\min sum(p,y) &= \sum_{i\in B\cup S} p_i + \sum_{t\in T} y_t \\
y_t &\geq 0 \;\; \forall t \in T: \\
p_i &\geq 0 \;\; \forall i \in S \cup B: \\
y_t &\geq (\theta_j - p_j) - (\theta_i + p_i) \;\; \forall t = (i,j) \in T
\end{aligned} \tag{3}$$

Lemma 1. *At equilibrium the profits must satisfy the above inequalities.*

Proof. Clearly all profits are nonnegative, as trading is optional for all agents.

To see why the last set of inequalities holds, consider two cases separately. For a trader t who conducted trade, we get equality by definition. For other traders $t = (i,j)$, note that offering a bid $\beta_t > p_i + \theta_i$ would get the seller to sell to trader t, as $p_i + \theta_i$ is the price that seller i sold for (or $p_i = 0$ if seller i decided to keep the good). Similarly, for any ask $\alpha_t < \theta_j - p_j$, the buyer will buy from trader t. So unless $\theta_j - p_j \leq \theta_i + p_i$ the trader has a profitable deviation. ■

Our first claim, that in any equilibrium of the trading game the trade is socially optimal, follows from the duality of the program describing socially optimal assignments and the program describing equilibria. Note that in the equilibrium problem we have replaced price setting by the traders with "surplus setting". These two approaches are equivalent by Lemma 1. Our duality approach is a direct generalization of the approach of Shapley and Shubik (1972) and Leonard (1983) to our model. The remarkable fact is that this approach works even when the price-setting agents (traders in our case) are part of the strategic model.

Proposition 1. *In any equilibrium the trade is efficient.*

Proof. Let x be a flow of goods resulting in an equilibrium, and let variables p and y be the profits.

Consider the linear program describing the socially optimal trade. We will also add a set of additional constraints $x_t \leq 1$ for all traders $t \in T$; this can be added to the description, as it is implied by the other constraints. Now we claim that the two linear programs are duals of each other. The variables p_i for agents

$B \cup S$ correspond to the equations $\sum_{t \in adj(i)} x_t \leq 1$. The additional dual variable y_t corresponds to an additional inequality $x_t \leq 1$.

The optimality of the social value of the trade will follow from the claim that the solution of these two linear programs derived from an equilibrium satisfy the complementary slackness conditions for this pair of linear programs, and hence both x and (p, y) are optimal solutions to the corresponding linear programs.

There are three different complementary slackness conditions we need to consider, corresponding to the three sets of variables x, y and p. Any agent can only make profit if he transacts, so $p_i > 0$ implies $\sum_{t \in adj(i)} x_t = 1$, and similarly, $y_t > 0$ implies that $x_t = 1$ also. Finally, consider a trader t with $x_t > 0$ that trades between seller i and buyer j, and recall that we have seen above that the inequality $y_t \geq (\theta_j - p_j) - (\theta_i + p_i)$ is satisfied with equality for those who trade. ■

Next we argue that equilibria exist. We do this by constructing equilibria from the solution to the dual of the social optimality program.

Proposition 2. *For any efficient trade between buyers and sellers there is a pure equilibrium of bid-ask values that supports this trade.*

Proof. First we prove the existence of a subgame perfect Nash equilibrium. Then we show how to modify the construction to produce an equilibrium; a non-crossing subgame perfect Nash equilibrium. Consider an efficient trade; let $x_t = 1$ if t trades and 0 otherwise; and consider an optimal solution (p, y) to the dual linear program.

We would like to claim that all dual solutions correspond to equilibrium profits, but unfortunately this is not exactly true. Before we can convert a dual solution to equilibrium prices, we may need to modify the solution slightly as follows. Consider any agent i that is only connected to a single trader t. Because the agent is only connected to a single trader, the variables y_t and p_i are dual variables corresponding to the same primal inequality $x_t \leq 1$, and they always appear together as $y_t + p_i$ in all inequalities, and also in the objective function. Thus there is an optimal solution in which $p_i = 0$ for all agents i connected only to a single trader.

Assume (p, y) is a dual solution where agents connected only to one trader have $p_i = 0$. For a seller i, let $\beta_t = \theta_i + p_i$ be the bid for all traders t adjacent to i. Similarly, for each buyer j, let $\alpha_t = \theta_j - p_j$ be the ask for all traders t adjacent to j. We claim that this set of bids and asks, together with the trade x, are an equilibrium. To see why, note that all traders t adjacent to a seller or buyer i offer the same ask or bid, and so trading with any trader is equally good for agent i. Also, if i is not trading in the solution x then by complementary slackness $p_i = 0$, and hence not trading is also equally good for i. This shows that sellers and buyers don't have an incentive to deviate.

We need to show that traders have no incentive to deviate either. When a trader t is trading with seller i and buyer j, then profitable deviations would involve increasing α_t or decreasing β_t. But by our construction (and assumption about monopolized agents) all sellers and buyers have multiple identical ask/bid offers, or trade is occurring at valuation. In either case such a deviation cannot be successful.

Finally, consider a trader $t = (i, j)$ who doesn't trade. A deviation for t would involve offering a lower ask to seller i and a higher bid to seller j than their current trade. However, $y_t = 0$ by complementary slackness, and hence $p_i + \theta_i \geq \theta_j - p_j$, so i sells for a price at least as high as the price at which j buys, so trader t cannot create profitable trade.

Next we show that the equilibrium can be constructed so that it is non-crossing. Consider an optimal solution to the dual linear program. To get an equilibrium without crossing bids, we need to do a more general modification than just assuming that $p_i = 0$ for all sellers and buyers connected to only a single trader. Let the set $\overline{E}$ be the set of edges $t = (i, j)$ that are *tight*, in the sense that we have the equality $y_t = (\theta_j - p_j) - (\theta_i + p_i)$. This set $\overline{E}$ contains all the edges where trade occurs, and some more edges. We want to make sure that $p_i = 0$ for all sellers and buyers that have degree at most 1 in $\overline{E}$. Consider a seller i that has $p_i > 0$. We must have i involved in a trade, and the edge $t = (i, j)$ along which the trade occurs must be tight. Suppose this is the only tight edge adjacent to agent i; then we can decrease p_i and increase y_t till one of the following happens: either $p_i = 0$ or the constraint of some other agent $t' \in adj(i)$ becomes tight. This change only increases the set of tight edges $\overline{E}$, keeps the solution feasible, and does not change the objective function value. So after doing this for all sellers, and analogously changing y_t and p_j for all buyers, we get an optimal solution where all sellers and buyers i either have $p_i = 0$ or have at least two adjacent tight edges.

Now we can set asks and bids to form a non-crossing equilibrium. For all traders $t = (i, j)$ associated with an edge $t \in \overline{E}$ we set α_t and β_t as before: we set the bid $\beta_t = p_i + \theta_i$ and the ask $\alpha_t = \theta_j - p_j$. For a trader $t = (i, j) \notin \overline{E}$ we have that $p_i+\theta_i > \theta_j-p_j$ and we set $\alpha_t = \beta_t$ to be any value in the range $[\theta_j-p_j, p_i+\theta_i]$. This guarantees that for each seller or buyer the best sell or buy offer is along the edge where trade occurs in the solution. The ask-bid values along the tight edges guarantee that traders who trade cannot increase their spread. Traders $t = (i, j)$ who do not trade cannot make profit due to the constraint $p_i + \theta_i \geq \theta_j - p_j$. ■

3.2. Distinguishable Goods

We now consider the case of distinguishable goods. As in the previous section, we can write a transshipment linear program for the socially optimal trade (1) with traders directly corresponding to edges. We can show that the dual of this linear program corresponds to trader profits. Recall that we needed to add the constraints $x_t \leq 1$ for all traders. The dual is then the program (3) with the only change being in the last constraint, which becomes:

$$y_t \quad \geq \quad (\theta_{ji} - p_j) - (\theta_{ij} + p_i) \;\; \forall t = (i, j) \in T \tag{4}$$

It is not hard to extend the proofs of Propositions 1 and 2 to this case. Profits in an equilibrium satisfy the dual constraints, and profits and trade satisfy complementary slackness. This shows that trade is socially optimal. Taking an optimal dual solution where $p_i = 0$ for all agents that are monopolized, we can convert it to a subgame perfect Nash equilibrium, and with a bit more care, we can also create an equilibrium.

Proposition 3. *All equilibria for the case of pair-traders with distinguishable goods result in socially optimal trade.*

3.3. Trader Profits

We have seen that all equilibria are efficient. However, it turns out that equilibria may differ in how the value of the allocation is spread between the sellers, buyers and traders. Figure 3 depicts a simple example of this phenomenon.

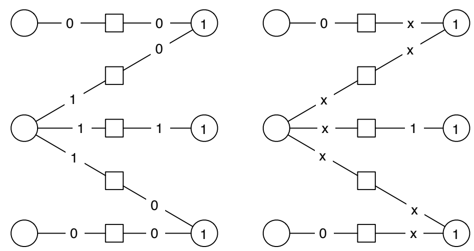

(a) *No trader profit* (b) *Trader profit*

Figure 3: Left: an equilibrium with crossing bids where traders make no money. Right: an equilibrium without crossing bids for any value $x \in [0, 1]$. Total trader profit ranges between 1 and 2.

Our goal is to understand how a trader's profit is affected by its position in the network. For the pair-trader case, this is straightforward. We can use the characterization we obtained to work out the range of profits a trader can make. To maximize the profit of a trader t (or a subset of traders T') all we need to do is to find an optimal solution to the dual linear program maximizing the value of y_t (or the sum $\sum_{t \in T'} y_t$).

Remark 1. *For any trader t or subset of traders T' the maximum total profit they can make in any equilibrium can be computed in polynomial time.*

One way to think about the profit of a trader $t = (i, j)$ is as a subtraction from the value of the corresponding edge (i, j). The value of the edge is the social value $\theta_{ji} - \theta_{ij}$ if the trader makes no profit, and decreases to $\theta_{ji} - \theta_{ij} - y_t$ if the trader t insists on making y_t profit. Trader t gets y_t profit in equilibrium, if after this decrease in the value of the edge, the edge is still included in the optimal transshipment.

Proposition 4. *In the pair-trader model a trader t can make profit in an equilibrium if and only if t is essential for the social welfare, that is, if deleting agent t decreases social welfare. The maximum profit he can make is exactly his marginal value to society, that is, the increase his presence causes in the social welfare.*

We can also find the minimum possible profit. Recall that in the proof of proposition 2, monopolized buyers and sellers cannot make profit. We can find the minimum profit a trader or set of traders can make, by minimizing the value

y_t (or sum $\sum_{t \in T'} y_t$) over all optimal solutions that satisfy $p_i = 0$ whenever i is connected to only a single trader.

Remark 2. *For any trader t or subset of traders T' the minimum total profit they can make in any equilibrium can be computed in polynomial time.*

4. General Model

In this section we extend the results from pair traders to general traders and prove Theorems 1 and 2. The proofs are more complex than in the pair-traders model as traders are no longer identified with edges, and hence can do more complex optimizations. We start by understanding the trader's optimization problem; we then prove that the trade is socially optimal in all equilibria (Theorem 1), and that equilibria always exist (Theorem 2). In subsection 4.2 we provide a characterization of trader profits for the general model (Theorem 3).

4.1. Trader Optimization and Equilibria

To relate equilibrium trade to socially optimal trade, we need to solve trader t's decision problem. Suppose we have bid and ask values for all other traders $t' \in T$, $t' \neq t$. What bid and ask offers are trader t's best response to the current set of bids and asks? For each seller i, let p_i be the maximum profit seller i can make using bids by other traders, and symmetrically let p_j be the maximum profit buyer j can make using asks by other traders (let $p_i = 0$ for any seller or buyer i who cannot make profit). Now consider a seller-buyer pair (i, j) that trader t can connect. Trader t will have to make a bid of at least $\beta_{tij} = \theta_{ij} + p_i$ to seller i and an ask of at most $\alpha_{tji} = \theta_{ji} - p_j$ to buyer j to get this trade, so the maximum profit she can make on this trade is $v_{tij} = \alpha_{tji} - \beta_{tij} = \theta_{ji} - p_j - (\theta_{ij} + p_i)$. The optimal trade for trader t is obtained by solving a matching problem to find the matching between the sellers $S(t)$ and buyers $B(t)$ connected to trader t that maximizes the total value v_{tij} for trader t. This matching is the optimal solution of the following linear program.

$$\begin{aligned}
\max V_t(x_t) &= \sum_{i\in S(t), j\in B(t)} x_{tij} v_{tij} \\
x_{tij} &\geq 0 \quad \forall i \in S(t); \forall j \in B(t) \\
\sum_{j\in B(t)} x_{tij} &\leq 1 \;\; \forall i \in S(t) \\
\sum_{i\in S(t)} x_{tij} &\leq 1 \;\; \forall j \in B(t)
\end{aligned} \tag{5}$$

We will also need the dual of this linear program. Let q_{ti} denote the dual variable associated with the constraint on seller or buyer i. The dual is then the following problem.

$$\begin{aligned}
\min val_t(q_t) &= \sum_{i\in B(t)\cup S(t)} q_{ti} \\
q_{ti} &\geq 0 \;\; \forall i \in S(t) \cup B(t). \\
q_{ti} + q_{tj} &\geq v_{tij} \;\; \forall i \in S(t), j \in B(t).
\end{aligned} \tag{6}$$

We view q_{ti} as the *profit made by* t *from trading with seller or buyer* i. Lemma 2 summarizes the above discussion.

Lemma 2. *For a trader t, given the lowest bids β_{tij} and highest asks α_{tji} that can be accepted for sellers $i \in S(t)$ and buyers $j \in B(t)$, the best trade t can make is the maximum value matching between $S(t)$ and $B(t)$ with value $v_{tij} = \alpha_{tji} - \beta_{tij}$ for the edge (i, j). This maximum value is equal to the minimum of the dual linear program (6).*

Proof of Theorem 1. Consider an equilibrium, with $x_e = 1$ if and only if trade occurs along edge $e = (i, j)$. Trade is a solution to the transshipment linear program (1).

Let p_i denote the profit of seller or buyer i. Each trader t currently has the best solution to his own optimization problem. A trader t finds his optimal trade (given bids and asks by all other traders) by solving an assignment problem. Let q_{ti} for $i \in B(t) \cup S(t)$ denote the optimal dual solution (6).

When setting up the optimization problem for a trader t above, we used p_i to denote the maximum profit i can make without the offer of trader t. Note that this p_i is exactly the same p_i we use here, the profit of agent i. This is clearly true for all traders t' that are not trading with i in the equilibrium. To see why it is true for the trader t that i is trading with, we use that the current set of bid-ask values is an equilibrium. If for any agent i the bid or ask of trader t were the *unique* best option, then t could extract more profit by offering a bit larger ask or a bit smaller bid, a contradiction.

We show the trade x is optimal by considering the dual solution $z_i = p_i + \sum_t q_{ti}$ for all agents $i \in B \cup S$. We claim z is a dual solution, and it satisfies complementary slackness with trade x. To see this we need to show a few facts.

- We need that $z_i > 0$ implies that i trades. If $z_i > 0$ then either $p_i > 0$ or $q_{ti} > 0$ for some trader t. Agent i can only make profit $p_i > 0$ if he is involved in a trade. If $q_{ti} > 0$ for some t, then trader t must trade with i, as his solution is optimal, and by complementary slackness for the dual solution, $q_{ti} > 0$ implies that t trades with i.

- For an edge (i, j) associated with a trader t we need to show the dual solution is feasible, that is $z_i+z_j \geq \theta_{ji}-\theta_{ij}$. Recall $v_{tij} = \theta_{ji}-p_j-(\theta_{ij}+p_i)$, and the dual constraint of the trader's optimization problem requires $q_{ti} + q_{tj} \geq v_{tij}$. Putting these together, we have $z_i + z_j \geq p_i + q_{ti} + p_j + q_{tj} \geq v_{tij} + p_i + p_j = \theta_{ji} - \theta_{ij}$.

- Finally, we need to show that the trade variables x also satisfy the complementary slackness constraint: when $x_e > 0$ for an edge $e = (i, j)$ then the corresponding dual constraint is tight. Let t be the trader involved in the trade. By complementary slackness of t's optimization problem we have $q_{ti} + q_{tj} = v_{tij}$. To see that z satisfies complementary slackness we need to argue that for all other traders $t' \neq t$ we have both $q_{t'i} = 0$ and $q_{t'j} = 0$. This is true as $q_{t'i} > 0$ implies by complementary slackness of t''s optimization problem that t' must trade with i at optimum, and $t \neq t'$ is trading. ∎

Proof of Theorem 2. As we did for Proposition 2, we first prove the existence of a subgame perfect Nash equilibrium, and then modify the construction to get an equilibrium. Let x be a socially optimal trade, and let z be an optimal dual solution. We will think of z_i as the profit associated with agent i, and will divide z_i into a profit p_i for agent i, and profits associated with traders q_{ti} for $t \in adj(i)$ such that $p_i + \sum_{t \in adj(i)} q_{ti} = z_i$ with $q_{ti} \neq 0$ only for the one trader t that agent i is trading with.

For each seller or buyer i if i is associated with only a single trader t then let $q_{ti} = z_i$ and set $p_i = 0$ (as monopolized sellers and buyers cannot have profit). For all other sellers and buyers i let $p_i = z_i$ and $q_{it} = 0$ for all traders $t \in adj(i)$. Now for all traders t, sellers $i \in S(t)$ and all buyers $j \in B(t)$, the trader will make the bid $\beta_{tij} = \theta_{ij} + p_i$, and the ask $\alpha_{tji} = \theta_{ji} - p_j$, offering p_i profit to the seller i and p_j profit to buyer j.

We claim that the trade x and these bid and ask values are at equilibrium. Each seller and buyer i has multiple options with p_i profit, so they don't have a profitable way to deviate.

To see that a trader t does not want to deviate, consider the trader's optimization problem. All sellers and buyers have multiple best options, so the value p_i is equal to the profit agent i can make also without any given trader t, and hence trader t will use $v_{tij} = \theta_{ji} - p_j - (\theta_{ij} + p_i)$ in her own optimization problem.

We claim that the values q_{ti} form a dual solution to t's optimization problem. For all agents i adjacent to a trader t we have that $z_i = p_i + q_{ti}$ as $q_{t'i} = 0$ for all these agents. This implies that q_t is feasible, as $z_i = p_i + q_{ti}$ is feasible for the overall optimization problem. To show that it is also optimal for t's optimization problem, we need to argue that the complementary slackness constraint for q_{ti} is satisfied, that is $q_{ti} > 0$ implies that i is trading with t. This is true as $q_{ti} > 0$ implies that t is the only trader adjacent to i, and further $z_i > 0$ and so i is trading in the solution x, so must be trading with t.

Now we construct an equilibrium. Consider an optimal trade x and a dual solution z. As before, we will divide the profit z_i between i and the trader t trading with i. We will use q_{ti} as the trader t's profit associated with agent i for any $i \in S(t) \cup B(t)$.

We will need to guarantee the following properties:

- Trader t trades with agent i whenever $q_{ti} > 0$. This is one of the complementary slackness conditions to make sure the current trade is optimal for trader t.

- For all seller-buyer pairs (i, j) that a trader t can trade with, we have

$$p_i + q_{ti} + p_j + q_{tj} \geq \theta_{ji} - \theta_{ij}, \qquad (7)$$

which will make sure that q_t is a feasible dual solution for the optimization problem faced by trader t.

- We need to have equality in (7) when trader t is trading between i and j. This is one of the complementary slackness conditions for trader t, and will ensure that the trade of t is optimal for the trader.

- Finally, we want to arrange that each agent i with $p_i > 0$ has multiple offers for making profit p_i, and the trade occurs at one of his best offers. To guarantee this in the corresponding bids and asks we need to make sure that whenever $p_i > 0$ there are multiple $t \in adj(i)$ that have equation in the above constraint (7).

We start by setting $p_i = z_i$ for all $i \in S \cup B$ and $q_{ti} = 0$ for all $i \in S \cup B$ and traders $t \in adj(i)$. This guarantees all invariants except the last property about multiple $t \in adj(i)$ having equality in (7). We will modify p and q to gradually enforce the last condition, while maintaining the others.

Consider a seller with $p_i > 0$. By optimality of the trade and dual solution z, seller i must trade with some trader t, and that trader will have equality in (7) for the buyer j that he matches with i. If this is the only trader t that has a tight constraint in (7) involving seller i then we increase q_{ti} and decrease p_i till either $p_i = 0$ or another trader $t' \neq t$ will be achieve equality in (7) for some buyer edge adjacent to i (possibly a different buyer j'). This change maintains all invariants, and increases the set of sellers that also satisfy the last constraint. We can do a similar change for a buyer j that has $p_j > 0$ and has only one trader t with a tight constraint (7) adjacent to j. After possibly repeating this for all sellers and buyers, we get profits satisfying all constraints.

Now we get equilibrium bid and ask values as follows. For a trader t that has equality for the seller–buyer pair (i, j) in (7) we offer $\alpha_{tji} = \theta_{ji} - p_j$ and $\beta_{tij} = \theta_{ij} + p_i$. For all other traders t and seller–buyer pairs (i, j) we have the invariant (7), and using this we know we can pick a value γ in the range

$\theta_{ij} + p_i + q_{ti} \geq \gamma \geq \theta_{ji} - (p_j + q_{tj})$. We offer bid and ask values $\beta_{tij} = \alpha_{tji} = \gamma$. Neither the bid nor the ask will be the unique best offer for the buyer, and hence the trade x remains an equilibrium. ■

4.2. Trader Profits

For the case of general traders the effect of a trader's position in the network on its profit is more complex than in the pair-traders model, as already indicated by Figure 2 discussed earlier in relation to Theorem 3.

First we show how to maximize the total profit of a set of traders. By the proof of Theorems 1 and 2 the profit of trader t in an equilibrium is $\sum_i q_{ti}$. To find the maximum possible profit for a trader t or a set of traders T', we need to do the following: Find profits $p_i \geq 0$ and $q_{ti} > 0$ so that $z_i = p_i + \sum_{t \in adj(i)} q_{ti}$ is an optimal dual solution, and also satisfies the constraints (7) for any seller i and buyer j connected through a trader $t \in T$. Now, subject to all these conditions, we maximize the sum $\sum_{t \in T'} \sum_{i \in S(t) \cup B(t)} q_{ti}$. Note that this maximization is a secondary objective function to the primary objective that z is an optimal dual solution. Then we use the proof of Theorem 2 to show how to turn this into an equilibrium.

Proposition 5. *The maximum value for $\sum_{t \in T'} \sum_i q_{ti}$ in the optimization problem above is the maximum profit the set T' of traders can make.*

Proof. We prove the Proposition by first getting the maximum value for a subgame perfect Nash equilibrium, and then showing it can be implemented by an equilibrium. By the proof of Theorem 1 the profits of trader t can be written in this form, so the set of traders T' cannot make more profit than claimed in this theorem.

To see that T' can indeed make this much profit, we use the proof of Theorem 2. We modify that proof to start with profit vectors p and q_t for $t \in T'$. For each seller or buyer i if i is associated with only a single trader $t \notin T'$ then let $q_{ti} = z_i$ and set $p_i = 0$ (as monopolized sellers and buyers cannot have profit). For all other sellers and buyers i use p_i and set q_{it}. Now for all traders t, sellers $i \in S(t)$ and all buyers $j \in B(t)$, the trader will make the bid $\beta_{tij} = \theta_{ij} + p_i$, and the ask $\alpha_{tji} = \theta_{ji} - p_j$, offering p_i profit to the seller i and p_j profit to buyer j.

We claim that the trade x and these bid and ask values are at equilibrium. First note that if a seller or buyer i is monopolized by a trader $t \in T$ than we must have $p_i = 0$, as otherwise decreasing p_i and increasing q_{ti} would improve the secondary objective function without violating any constraint. Now use the proof of Theorem 2 to see that the bids and asks form an equilibrium.

The equilibrium construction follows part of the proof of Theorem 2. We first verify that the constructed above satisfies the first three of the four required properties, and then we can follow the proof of Theorem 2 to make the fourth property true.

The first property requires that whenever $q_{ti} > 0$, trader t is trading with agent i. This obvious for $t \notin T'$ as $q_{ti} = 0$. Consider a trader $t \in T'$, and an agent i with $q_{ti} > 0$. We must have $z_i > 0$ so by the complementary slackness condition for z agent i must be involved in a trade. Let t' be the trader involved in this trade. Say agent i is a seller, and trader t' is matching i with a buyer j. By complementary slackness on edge (i, j) associated with the trader t' we have the equation $z_i + z_j = \theta_{ji} - \theta_{ij}$. The edge (i, j) is associated with a trader $t' \neq t$, so we have $p_i + q_{t'i} \leq z_i - q_{ti}$ and $p_j + q_{t'i} \leq z_i - q_{tj}$, and by the condition above for t' we need $p_i + q_{t'i} + p_j + q_{t'i} \geq \theta_{ji} - \theta_{ij}$, so we must have $q_{ti} = q_{tj} = 0$.

The second property requires $p_i + q_{ti} + p_j + q_{tj} \geq \theta_{ji} - \theta_{ij}$ for all traders t and trades (i, j) they can make. This is an explicit condition we used for optimization.

The third property required that there is equality in the above whenever t trades between i and j. This follows from the complementary slackness for z on this edge, and from the first property that $q_{t'i} = q_{t'j} = 0$ for all other traders $t' \neq t$.
■

In Section 3.3 we showed that in the case of pair traders, a trader t can make money if he is essential for efficient trade. This is not true for the type of more general traders we consider here, as was shown by the example in Figure 2. The correct characterization is given by Theorem 3, which we now prove.

Proof of Theorem 3. First we show the direction that if a trader t can make money there must be an agent i so that t's connection to i is essential to social welfare. Let p, q be the profits in an equilibrium where t makes money, as described by

Theorem 1 with $\sum_{i \in S(t) \cup B(t)} q_{ti} > 0$. So we have some agent i with $q_{ti} > 0$. We claim that the connection between agent i and trader t must be essential, in particular, we claim that social welfare must decrease by at least q_{ti} if we delete i from $adj(t)$. To see why note that decreasing the value of all edges of the form (i, j) associated with trader t by q_{ti} keeps the same trade optimum, as we get a matching dual solution by simply resetting q_{ti} to zero.

To see the opposite, assume deleting i from $adj(t)$ decreases social welfare by some value γ. Assume i is a seller (the case of buyers is symmetric), and decrease by γ the social value of each edge (i, j) for any buyer j such that t is the only agent connecting i and j. By assumption the trade is still optimal, and we let z be the dual solution for this matching. This creates a solution with $\sum_{t \in T'} \sum_i q_{ti} + \gamma$, so the maximum in Theorem 3 is at least γ. ■

Finally, we show that the minimum possible profit is found as a dual solution minimizing the dual variables associated with agents monopolized by some trader.

Remark 3. *For any trader t or subset of traders T', the minimum total profit they can make in any equilibrium can be computed in polynomial time.*

5. Extensions of the Model

In this final section we extend our results to two more general problems. The first extension concerns the case with indistinguishable goods where buyers and sellers may have multiple items. The second extension is to consider trading costs. The easiest way to do this is to treat the more general idea of trader-dependent valuations (by buyers and sellers), which may be of independent interest.

5.1. Agents with Multiple Goods

Our basic model assumes that all sellers have one good, and buyers are interested in at most one good. We can extend the results to the case when agents (both sellers are buyers) have concave utility functions for goods. To be concrete, assume that each seller i has an endowment k_i of goods, and each buyer j is possibly interested in up to k_j units of the good, and as before the goods are indistinguishable. We will assume that both sellers and buyers have concave utility functions. For a buyer j we can use θ_{jk} to denote the buyer's value for the k^{th} copy of the item, and will assume that θ_{jk} is decreasing in k. For a seller i we use θ_{ik} for the value of k^{th} item sold, and we assume that θ_{ik} is increasing in k. So both sellers and buyers value the item more when they have fewer copies.

Our results for the case with $k_i = 1$ for all agents i extend to this case by a simple reduction. We replace each agent i with k_i separate agents, all connected to the same set of traders. The k^{th} copy of agent i will have valuation θ_{ik}. Valuations are monotone, so in all socially optimal trade, the first few copies of all agents will be involved in a trade. Note that in the trade resulting from this reduction, a trader t will make multiple bids and asks to an agent: separate bid and ask values for the each successive copy of the goods. We can think of the asks and bids made by the traders as menus, analogous to the menus we used for the case of distinguishable goods: a trader t offers buyer i a bid β_{ti1} for buying one item, $\beta_{ti1} + \beta_{ti2}$ for buying two items, and so forth. Using this reduction we get the following theorem.

Theorem 4. *Trade at an equilibrium of this game with multiple indistinguishable goods is always socially optimal, and there exists an equilibrium supporting any socially optimal trade. A trader t can make positive profit if for an adjacent seller $i \in adj(t) \cap S$ decreasing the number of goods t can buy from seller i decreases the social welfare, or alternatively, if there is an adjacent buyer $j \in adj(t) \cap B$ such that decreasing the number of items t can sell to seller i by one decreases social welfare.*

This Theorem generalizes to the networked market the idea that first degree price discrimination is socially optimal. In fact it is necessary, since a monopolist facing a seller and buyer with multiple items will not conduct the optimal trade if he has to offer a single bid and ask for all copies of the good. Concerning the ability of a trader to make profit, note that it is important to have the last item of the seller or buyer improve social welfare. In the example of Figure 2 we can combine the two top sellers into one seller with two units of the good. This will make the link of between the top seller and trader essential for social welfare, and yet the trader cannot make money.

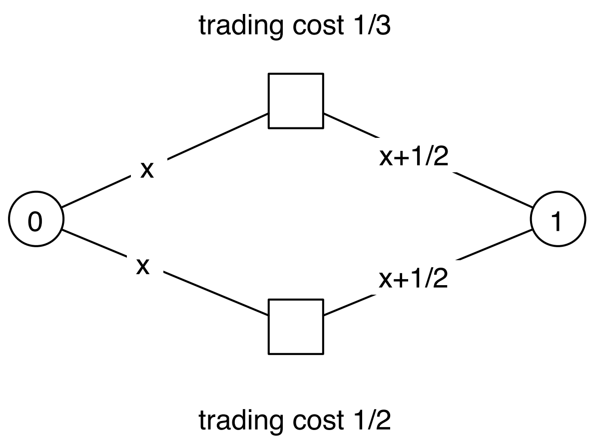

Figure 4: The upper trader t is able to perform the transaction more cheaply, and trade therefore occurs through t at a spread of $\frac{1}{2}$ — the largest spread that still ensures that in equilibrium t performs the transaction. In any equilibrium achieving this outcome, t makes a profit of $\frac{1}{2} - \frac{1}{3} = \frac{1}{6}$.

5.2. Trading Costs and Trader-Dependent Valuations

Trading Costs.. A natural extension to consider is the case in which traders incur costs to perform transactions. Suppose that trader t incurs a cost of c_{tij} for conducting trade between seller i and buyer j. Note that this trading cost will affect the social welfare: the welfare induced by trader t conducting a trade between seller i and buyer j is now $\theta_{tji} - \theta_{tij} - c_{tij}$: an increase of θ_{tji} in the welfare of buyer j, a loss of θ_{tij} in the utility of seller i and a social loss of c_{tij} for the trading cost.

We show an example of this in Figure 4. The upper trader t is able to perform the transaction more cheaply, and hence the trade goes through t at a bid price of x and an ask price of $x + \frac{1}{2}$ (for an $x \in [0, \frac{1}{2}]$). Note that this is the largest spread that ensures that in equilibrium t performs the transaction, and it provides t with a profit of $\frac{1}{2} - \frac{1}{3} = \frac{1}{6}$, the difference between the spread and the trading cost.

Thus, with this underlying graph structure, neither trader can make a profit when all trading costs are 0, but when the costs incurred by the traders differ, then the one who can perform the transaction more cheaply is able to make a profit. (Moreover, this trader is necessary for achieving the optimum social welfare, as a lower trading cost means that less of the surplus is lost due to trading costs.)

In fact, the main results for our model carry over to the generalization in which traders can have non-zero trading costs. Rather than proving this directly, we derive it from a different generalization of the model, in which the valuations of both sellers and buyers depend on the identity of the trader t conducting the trade. We now describe this further generalization, and then explain how the model with trading costs arises as a special case.

Trader-Dependent Valuations.. In the extension to trader-dependent valuations, the value of a trade from seller i to buyer j conducted by trader t is θ_{jit} for buyer j, and the seller i experiences θ_{ijt} loss of utility. We use the exact same game as before, with traders making asks and bids to buyers and sellers via a menu of options: asks α_{tji}, where α_{tji} is the ask for the product of seller i, and symmetrically, a trader t can offer to each seller i a menu of bids β_{tij} for selling to different buyers j. To see that our results extend to this more general case, note that we can use the new trader-dependent valuations in the proofs and constructions above with essentially no change.

Now, we can incorporate trading costs in the model by a simple reduction: we charge the buyer (say) for the cost.[9] We will assume that an ask of β_{tji} for this trade means that if buyer j accepts this ask, he will have to pay β_{tji} to the trader t and will also be responsible for covering the trading cost c_{tij}. Under this bidding model, the value of the trade for buyer j is decreased to $\theta'_{tji} = \theta_{tji} - c_{tij}$, and we have reduced the problem to the case without trading costs.

Theorem 5. *Trade at an equilibrium of this game with trader dependent valuations and trader costs is always socially optimal, and there exists an equilibrium supporting any socially optimal trade. A trader t can make profit in an equilibrium if and only if there is a seller or buyer i adjacent to t such that the connection of trader t to agent i is essential for social welfare.*

The generalization to trader-dependent valuations not only includes transactions costs as a special case, it also includes a primitive form of production. A

[9]Charging the seller is also possible with no substantive change to our results.

trader might buy a good from seller x, transform it to something else and sell it to buyer y. Transactions costs are essentially this — a good available at one location is transformed into a good available at another location. Similarly, rough-cut stones can be transformed to polished gems, and fresh tomatoes can be transformed to bruised tomatoes. More general models of production, in which one seller's iron ore and another seller's plastic is turned into a car, take us far beyond the scope of this paper.

6. Conclusion

In this paper we replace the abstract concept of a market with a graph describing who can trade with whom, and we describe who actually does trade with whom in equilibrium. In our model all trade goes through intermediaries who attempt to make a profit by buying at their bid price and selling at their ask price. We show that in any equilibrium of our price-setting and trading game the resulting allocation of goods is efficient (as constrained by the network structure). We also characterize in graph-theoretic terms how a trader's profit is affected by the degree of competition it faces from other traders.

There are many features of real markets that our stylized model does not incorporate. We analyzed a simple model in order to gain insight into the role of network structure. But it would be interesting to enrich our analysis to endogenize the network and to consider asymmetric information about buyer and seller values. An equilibrium in our trading game generates payoffs that the agents receive given the network. If the agents know these payoffs for potential networks, and can influence the network at some cost, then we can define a multi-stage game in which the network itself is endogenous. An extension to a setting with asymmetric information would also be interesting. In our trading game, only the traders use information about buyer and seller values; the buyers and sellers do not need to know each others values. So we can handle asymmetry in information between buyers and sellers. An extension to a setting in which traders do not know buyer and seller values would be more challenging. We view our trading game as natural in settings in which the traders have long experience in trading with the buyers and sellers in a stable setting. It may perform well both if traders know individuals values and if each trader faces a large population of buyers and sellers and only knows distributions of values. More generally, with asymmetric information, our trading mechanism will not yield (full information) socially optimal trades in one shot games. But, in this setting, no mechanism will yield (full information) socially optimal trade as incentive constraints will bind; see Myerson and Satterthwaite (1983). How well our mechanism would preform in repeated interactions is an interesting question for future research.

References

Babaioff, M., Nisan, N., Pavlov, E., 2005. Mechanisms for a Spatially Distributed Market. ACM EC Conference.

Barrett, C., Mutambatsere, E., 2008. Agricultural markets in developing countries. In: Durlauf, S. N. , Blume, L. E. (Eds.), The New Palgrave Dictionary of Economics, 2nd edition, Palgrave Macmillan.

Bikhchandani, S., Ostroy, J., 2002. The Package Assignment Model. J. Econ. Theory 107, 377–406.

Chu, L., Shen, Z.-J., 2006. Agent Competition Double Auction Mechanism. Management Science 52, 1215–1222.

Crawford, V., Knoer, E., 1981. Job Matching with Heterogeneous Firms and Workers. Econometrica 49, 437–4504.

Demange, G., Gale, D., Sotomayor, M., 1986. Multi-item auctions. J. Polit. Economy 94, 863–872.

Edid, M. S., 1994. Farm Labor organizing: Trends and Prospects. Cornell Univsrsity Press.

Even-Dar, E., Kearns, M., Suri, S., 2007. A Network Formation Game for Bipartite Exchange Economies. ACM-SIAM Symp. on Discrete Algorithms (SODA).

Fafchamps, M., Gabre-Madhin, E., 2007. Agricultural markets in Benin and Malawi. African Journal of Agricultural and Resource Economics.

Gale, D., Kariv, S., 2007. Financial Networks. Amer. Econ. Rev.: Papers and Proceedings, pp. 99–103.

Hall, G., Rust, J., 2000. An empirical model of inventory investment by durable commodity intermediaries. Carnegie-Rochester Conference Series on Public Policy 52(1), 171–214.

Hellwig, C., 2003. Money, intermediation, and cash-in-advance constraints. unpublished, UCLA.

Kakade, S., Kearns, M., Ortiz, L., Pemantle, R., Suri, S., 2004. Economic Properties of Social Networks. Neural Information Processing Systems (NIPS).

Kelso, A., Crawford, V., 1982. Job Matching, Coalition Formation, and Gross Substitutes. Econometrica 50, 1483–1504.

Kranton, R., Minehart, D., 2001. A Theory of Buyer-Seller Networks. Amer. Econ. Rev. 91, 85–508.

Krishna, K., Imai, S., Mukhopadhyay, A., and Tan, L. H., 2004. Quota Brokers. IMF Working paper.

Lamoreaux, N. R., Sokoloff, K. L., 2002. Intermediaries in the U.S. Market for Technology, 1870-1920. NBER Working Paper No. 9017.

Leonard, H., 1983. Elicitation of Honest Preferences for the Assignment of Individuals to Positions. J. Polit. Economy 91, 461–479.

Li, Y., 1998. Middlemen and private information. Journal of Monet. Econ. 42(1), 131–159.

Myerson, R., Satterthwaite, M., 1983. Efficient Mechanisms for Bilateral Trading. J. Econ. Theory 29, 265–281.

O'Hara, M., 1995. Market Microstructure Theory. Blackwell Publishers, Cambridge, MA.

Rubinstein, A., Wolinsky, A., 1987. Middlemen. Quart. J. of Econ. 102(3), 581–93.

Sarma, A. D., Chakrabarty, D., and Gollapudi, S., 2007. Public Advertisement Broker Markets. Internet and Network Economics, Third International Workshop, WINE 2007.

Shapley, L., Shubik, M., 1972. The Assignment Game I: The Core. Intl. J. Game Theory pp. 111–130.

Spulber, D., 1999. Market microstructure: Intermediaries and the theory of the firm. Cambridge University Press.

Yavas, A., 1994. Middlemen in bilateral search markets. J. of Lab. Econ. 12(3), 406–29.