

Trading Networks with Price-Setting Agents

Larry Blume *

David Easley †

Jon Kleinberg ‡

Éva Tardos §

Abstract

In a wide range of markets, individual buyers and sellers often trade through intermediaries, who determine prices through 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 via 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. In this model, traders set prices strategically, and then buyers and sellers react to the prices they are presented. We show that the resulting game always has a subgame perfect Nash equilibrium, and that all equilibria lead to an efficient (i.e. socially optimal) allocation of goods. We extend these results to a more general type of matching market, such as one finds in the matching of job applicants and employers. Finally, we consider 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.

Our work differs from recent studies of how price is affected by network structure through our modeling of price-setting as a strategic activity carried out by a subset of agents in the system, rather than studying prices set via competitive equilibrium or by a truthful mechanism.

1 Introduction

In a range of settings where markets mediate the interactions of buyers and sellers, one observes several recurring properties: 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. For example, in financial markets much of

the trade between buyers and sellers is intermediated by a variety of agents ranging from brokers to specialists on the floor of the NYSE to electronic trading systems. There is no one market for a stock; trade in a single asset may occur simultaneously on the NYSE, on NASDAQ, 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. The price at which the asset trades may differ across these trading venues. In fact, there is no “price” as different traders pay or receive different prices. In many settings there is also a gap between the price a buyer pays for an asset, the *ask price*, and the price a seller receives for the asset, the *bid price*.

In this paper, we develop a framework in which such phenomena emerge from a game-theoretic model of trade, with buyers, sellers, and traders interacting on a network. The edges of the network connect traders to buyers and sellers, and thus represent the access that different market participants have to one another. The traders serve as intermediaries in a two-stage trading game: they strategically choose bid and ask prices to offer to the sellers and buyers they are connected to; the sellers and buyers then react to the prices they face. Thus, 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. We also analyze how trader profits depend on the network structure, essentially characterizing in graph-theoretic terms how a trader’s payoff is determined by the amount of competition it experiences with other traders.

Before describing the details of the model, we briefly mention some of the work in economics, finance, and algorithmic game theory to which our work is connected. A more extended discussion of this related work is provided at the end of this section. The most widely-used approach for analyzing buyer-seller interactions is the Walrasian model, which provides equilibrium prices without focusing on where these prices come from. There are several bodies of research in economics and finance which examine how prices are set, rather than just pursuing the goal of determining equilibrium prices. These include the literatures on imperfect competition [13] and equilibrium search [2]. Models in market microstructure [11] and information economies [6] incorporate

*Dept. of Economics, Cornell University, Ithaca NY 14853. Email: lb19@cornell.edu. Supported in part by NSF grant CCF-0325453.

†Dept. of Economics, Cornell University, Ithaca NY 14853. Email: dae3@cornell.edu. Supported in part by NSF grant CCF-0325453.

‡Dept. of Computer Science, Cornell University, Ithaca NY 14853. Email: kleinber@cs.cornell.edu. Supported in part by a John D. and Catherine T. MacArthur Foundation Fellowship, a Google Research Grant, a Yahoo! Research Alliance Grant, and NSF grants CCF-0325453, IIS-0329064, CNS-0403340, and BCS-0537606.

§Dept. of Computer Science, Cornell University, Ithaca NY 14853. Email: eva@cs.cornell.edu. Supported in part by a Yahoo! Research Alliance Grant, NSF grant CCF-0325453, and ONR grant N00014-98-1-0589.

price-setting agents, but there is little research in these literatures examining the effect of network-type constraints on who can trade with whom. The use of networks restricting the pairs of buyers and sellers who can trade has been considered in the context of competitive equilibrium [5, 7], mechanism design [1, 3, 4, 8, 9], and the core [12]. A fundamental difference in our work is that we consider a network that includes buyers and sellers together with traders who serve as price-setting agents; and we model the formation of prices as a strategic process carried out by these traders, rather than as the result of a centrally controlled or exogenous mechanism.

The Basic Model: Indistinguishable Goods. Our goal in formulating the model is to capture 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 most basic version of the model, we consider identical goods, one copy of which is initially held by each seller. Buyers and sellers each have a value for one copy of the good, and we assume that these values are common knowledge. We will subsequently generalize this to a setting in which goods are distinguishable, buyers can value different goods 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).

Thus, to start with the basic model, there is a single type of *good*; the good comes in indivisible units; and each seller initially holds one unit of the good. All three types of agents value money at the same rate; and each $i \in B \cup S$ additionally values one copy of the good at θ_i units of money. No agent wants more than one copy of the good, so additional copies are valued at 0. Each agent has an initial endowment of money that is larger than any 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. (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.) 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:

- 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 .)
- 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.)
- For each trader t , with accepted offers from sellers i_1, \dots, i_s and buyers j_1, \dots, 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.

This defines the basic elements of the game. The equilibrium concept we use is subgame perfect Nash equilibrium.

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, \dots from top to bottom; traders are drawn as squares in the middle column and will be named t_1, t_2, \dots from top to bottom; and buyers are drawn as circles in the rightmost column and will be named j_1, j_2, \dots 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 a standard second-price auction arises naturally from our model. Suppose the buyer valuations from top 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

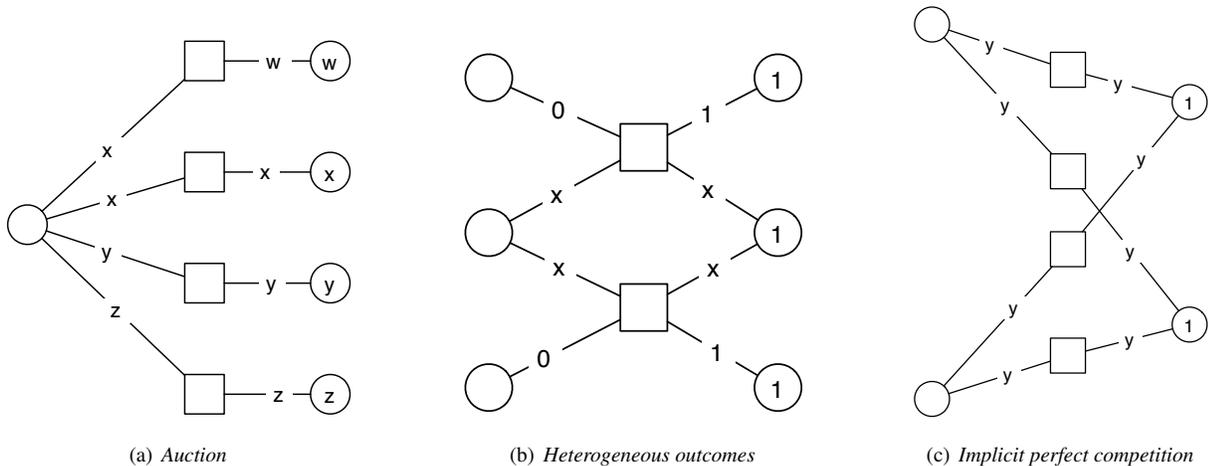


Figure 1: (a) An auction, mediated by traders, in which the buyer with the highest valuation for the good ends up with it. (b) A network in which the middle seller and buyer benefit from perfect competition between the traders, while the other sellers and buyers have no power due to their position in the network. (c) 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.

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, we can consider this particular instance as 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.

In Figure 1(b), 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 two traders drive the bid-ask spreads to 0 in competing for their business. 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.

The previous example shows perhaps the most natural 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 that 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(c) 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 . 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.

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 a bid to each seller that specifies both a price *and* a buyer, and offering an ask to each buyer that specifies both a price and a seller. (We can also handle a model in which a trader offers bids (respectively, asks) in the form of vectors, essentially specifying a “menu” with a price attached to each buyer (resp. seller).) Each buyer and seller selects an offer from an adjacent trader, and the payoffs to all agents are determined

as before.

This general framework captures matching markets [9, 12]: 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.

Our Results. Our results will identify general forms of some of the principles noted in the examples discussed above — including the question of which buyers end up with the good; the question of how payoffs are differently realized by sellers, traders, and buyers; and the question of what structural properties of the network determine whether the traders will make positive profits.

To make these precise, we introduce the following notation. Any outcome of the game determines a final allocation of goods to some of the agents; this can be specified by a collection M of triples (i_e, t_e, j_e) , where $i_e \in S$, $t_e \in T$, and $j_e \in B$; moreover, each seller and each buyer appears in at most one triple. The meaning is for each $e \in M$, the good initially held by i_e moves to j_e through t_e . (Sellers appearing in no triple keep their copy of the good.) We say that the *value* of the allocation is equal to $\sum_{e \in M} \theta_{j_e i_e} - \theta_{i_e j_e}$. Let θ^* denote the maximum value of any allocation M that is feasible given the network.

We show that every instance of our game has an equilibrium, and that in every such equilibrium, the allocation has value θ^* — in other words, it achieves the best value possible. Thus, equilibria in this model are always *efficient*, in that the market enables the “right” set of people to get the good, subject to the network constraints. We establish the existence and efficiency of equilibria by constructing a linear program to capture the flow of goods through the network; the dual of this linear program contains enough information to extract equilibrium prices.

By the definition of the game, the value of the equilibrium allocation is divided up as payoffs to the agents, and it is interesting to ask how this value is distributed — in particular how much profit a trader is able to make based on its position in the network. We find that, although all equilibria have the same value, a given trader's payoff can vary across different equilibria. However, we are able to characterize the maximum and minimum amounts that a given trader is able to make, where these maxima and minima are taken over all equilibria, and we give an efficient algorithm to compute this. In particular, our results here imply a clean combinatorial characterization of when a given trader t can achieve non-zero payoff: this occurs if and only there is some edge e incident to t that is *essential*, in the sense that deleting e reduces the value of the optimal allocation θ^* . We also obtain results for the sum of all trader profits.

Related Work. The standard baseline approach for analyzing the interaction of buyers and sellers is the Walrasian model in which anonymous buyers and sellers trade a good at a single market clearing price. This reduced form of trade, built on the idealization of a market price, is a powerful model which has led to many insights. But it is not a good model to use to examine where prices come from or exactly how buyers and sellers and trade with each other. The difficulty is that in the Walrasian model there is no agent who sets the price, and agents don't actually trade with each other. In fact there is no market, in the everyday sense of that word, in the Walrasian model. That is, there is no physical or virtual place where buyers and sellers interact to trade and set prices. Thus in this simple model, all buyers and sellers are uniform and trade at the same price, and there is also no role for intermediaries.

There are several literatures in economics and finance which examine how prices are set rather than just determining equilibrium prices. The literature on imperfect competition is perhaps the oldest of these. Here a monopolist, or a group of oligopolists, choose prices in order to maximize their profits (see [13] for the standard textbook treatment of these markets). A monopolist uses its knowledge of market demand to choose a price, or a collection of prices if it discriminates. Oligopolists play a game in which their payoffs depend on market demand and the actions of their competitors. In this literature there are agents who set prices, but the fiction of a single market is maintained. In the equilibrium search literature, firms set prices and consumers search over them (see [2]). Consumers do end up paying different prices, but all consumers have access to all firms and there are no intermediaries. In the general equilibrium literature there have been various attempts to introduce price determination. A standard proof technique for the existence of competitive equilibrium involves a price adjustment mechanism in which prices respond to excess demand. The Walrasian auctioneer is often introduced as a device to explain how this process works, but this is a fundamentally a metaphor for an iterative price-updating algorithm, not for the internals of an actual market. More sophisticated processes have been introduced to study the stability of equilibrium prices or the information necessary to compute them. But again there are no price-setting agents here.

In the finance literature the work on market microstructure does have price-setting agents (specialists), parts of it do determine separate bid and ask prices, and different agents receive different prices for the same asset (see [11] for a treatment of microstructure theory). Work in information economics has identified similar phenomena (see e.g. [6]). But there is little research in these literatures examining the effect of restrictions on who can trade with whom.

There have been several approaches to studying how network structure determines prices. These have posited price determination through definitions based on competitive equilibrium or the core, or through the use of truthful mecha-

nisms. In briefly reviewing this work, we will note the contrast with our approach, in that we model prices as arising from the strategic behavior of agents in the system.

In recent work, Kakade et al. [7] have studied the distribution of prices at competitive equilibrium in a bipartite graph on buyers and sellers, generated using a probabilistic model capable of producing heavy-tailed degree distributions [10]. Even-Dar et al. [5] build on this to consider the strategic aspects of network formation when prices arise from competitive equilibrium.

Leonard [9], Babaioff et al. [1], and Chu and Shen [3] consider an approach based on mechanism design: buyers and sellers reside at different nodes in a graph, and they incur a given transportation cost to trade with one another. Leonard studies VCG prices in this setting; Babaioff et al. and Chu and Shen additionally provide a budget-balanced mechanism. Since the concern here is with truthful mechanisms that operate on private valuations, there is an inherent trade-off between the efficiency of the allocation and the budget-balance condition.

In contrast, our model has known valuations and prices arising from the strategic behavior of traders. Thus, the assumptions behind our model are in a sense not directly comparable to those underlying the mechanism design approach: while we assume known valuations, we do not require a centralized authority to impose a mechanism. Rather, price-setting is part of the strategic outcome, as in the real markets that motivate our work, and our equilibria are simultaneously budget-balanced and efficient — something not possible in the mechanism design frameworks that have been used.

Demange, Gale, and Sotomayor [4], and Kranton and Minehart [8], analyze the prices at which trade occurs in a network, working within the framework of mechanism design. Kranton and Minehart 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, we can show how the basic model of Kranton and Minehart can be encoded as an instance of our game, with traders producing prices at equilibrium matching the prices produced by their auction mechanism.¹

Finally, the classic results of Shapley and Shubik [12] 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 focus is different as they do not consider agents that seek to set prices.

¹Kranton 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.

2 Markets with Pair-Traders

For understanding the ideas behind the analysis of the general model, it is very useful to first consider a special case with a restricted form of traders that we refer to as *pair-traders*. In this case, each trader is connected to just one buyer and one seller. (Thus, it essentially serves as a “trade route” between the two.) The techniques we develop to handle this case will 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 certain linear program, use this relation to show that all equilibria result in efficient allocation of the goods, and show that a pure equilibrium 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. We then characterize the minimum and maximum profits traders can make. In the next section, we extend the results to traders that may be connected to any subset of sellers and buyers.

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. For each buyer and seller i , we will use $adj(i)$ to denote the set of traders who can trade with i .

2.1 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 following linear program.

$$\begin{aligned} \max SV(x) &= \sum_{t \in T: t=(i,j)} x_t(\theta_j - \theta_i) \\ x_t &\geq 0 \quad \forall t \in T \\ \sum_{t \in adj(i)} x_t &\leq 1 \quad \forall i \in S \\ \sum_{t \in adj(j)} x_t &\leq 1 \quad \forall j \in B \end{aligned}$$

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 here 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.

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

Lemma 2.1 *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 $t = (i, j)$, the value $p_i + \theta_i$ is the price that seller i sold for (or θ_i if seller i decided to keep the good). Offering a bid $\beta_t > p_i + \theta_i$ would get the seller to sell to trader t . Similarly, $\theta_j - p_j$ is the price that buyer j bought for (or θ_j if he didn't buy), and 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. ■

Now we are ready to prove our first theorem:

Theorem 2.2 *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 $\sum_{t \in \text{adj}(i)} x_t \leq 1$. 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 $\sum_{t \in \text{adj}(i)} x_t = 1$, 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 always exist.

Theorem 2.3 *For any efficient trade between buyers and sellers there is a pure equilibrium of bid-ask values that supports this trade.*

Proof. 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 prices, 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. ■

Note that a seller or buyer i connected to a single trader t cannot have profit at equilibrium, so possible equilibrium profits are in one-to-one correspondence with dual solutions for which $p_i = 0$ whenever i is monopolized by one trader.

A disappointing feature of the equilibrium created by this proof is that some agents t may have to create ask-bid pairs where $\beta_t > \alpha_t$, offering to buy for more than the price at which they are willing to sell. Agents that make such crossing bid-ask pairs never actually perform a trade, so it does not result in negative profit for the agent, but such pairs are unnatural. Crossing bid-ask pairs are weakly dominated by the strategy of offering a low bid $\beta = 0$ and an extremely high ask to guarantee that neither is accepted.

To formulate a way of avoiding such crossing pairs, we say an equilibrium is *cross-free* if $\alpha_t \geq \beta_t$ for all traders t . We now show there is always a cross-free equilibrium.

Theorem 2.4 *For any efficient trade between buyers and sellers there is a pure cross-free equilibrium.*

Proof. 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 \bar{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 \bar{E} contain 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 \bar{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 $t' \in \text{adj}(i)$ becomes tight. This change only increases the set of tight edges \bar{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 cross-free equilibrium. For all traders $t = (i, j)$ associated with an edge $t \in \bar{E}$ 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 $t = (i, j) \notin \bar{E}$ 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 $p_i + \theta_i \geq \theta_j - p_j$ ■

2.2 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, with the only change being in the objective function.

$$\max SV(x) = \sum_{t \in T: t=(i,j)} x_t(\theta_{ji} - \theta_{ij})$$

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:

$$\begin{aligned} \min \text{sum}(p, y) &= \sum_{i \in B \cup S} p_i + \sum_{t \in T} y_t \\ y_t &\geq 0 \quad \forall t \in T : \\ p_i &\geq 0 \quad \forall i \in S \cup B : \\ y_t &\geq (\theta_{ji} - p_j) - (\theta_{ij} + p_i) \quad \forall t = (i, j) \in T \end{aligned}$$

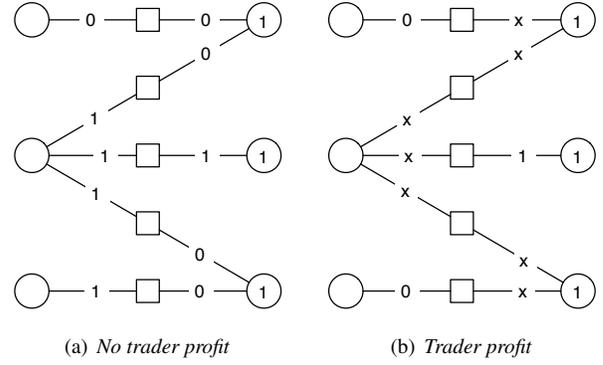


Figure 2: Left: an equilibrium with crossing bids where traders make no money. Right: an equilibrium without crossing bids for any value $x \in [0, 1]$. Total trader profit ranges between 1 and 2.

It is not hard to extend the proofs of Theorems 2.2 – 2.4 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 an equilibrium, and with a bit more care, we can also create an equilibrium with no crossing bid-ask pairs.

Theorem 2.5 *All equilibria for the case of pair-traders with distinguishable goods result in socially optimal trade. Pure non-crossing equilibria exist.*

2.3 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 2 depicts a simple example of this phenomenon.

Our goal is to understand how a trader's profit is affected by its position in the network; we will 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 $\sum_{t \in T'} y_t$). Such dual solutions will then correspond to equilibria with non-crossing prices.

Theorem 2.6 *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. This maximum profit can be obtained by a non-crossing equilibrium.*

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.

Theorem 2.7 *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 value to society, that is, the increase his presence causes in the social welfare.*

If we allow crossing equilibria, then we can also find the minimum possible profit. Recall that in the proof of Theorem 2.3, traders only made money off of sellers or buyers that they have a monopoly over. Allowing such equilibria with crossing bids 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.

Theorem 2.8 *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.*

3 General Traders

Next we extend the results to a model where traders may be connected to an arbitrary number of sellers and buyers. For a trader $t \in T$ we will use $S(t)$ and $B(t)$ to denote the set of buyers and sellers connected to trader t . In this section we focus on the general case when goods are distinguishable (i.e. both buyers and sellers have valuations that are sensitive to the identity of the agent they are paired with in the allocation). In the full version of the paper we also discuss the special case of indistinguishable goods in more detail.

To get the optimal trade, 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: $E = \{(i, j) : adj(i) \cap adj(j) \neq \emptyset\}$. On this graph, we then solve the instance of the assignment problem that was also used in Section 2.2, with the value of edge (i, j) equal to $\theta_{ji} - \theta_{ij}$ (since the value of trading between i and j is independent of which trader conducted the trade). We will also use the dual of this linear program:

$$\begin{aligned} \min val(z) &= \sum_{i \in B \cup S} z_i \\ z_i &\geq 0 \quad \forall i \in S \cup B. \\ z_i + z_j &\geq \theta_{ji} - \theta_{ij} \quad \forall i \in S, j \in B : \\ &\quad adj(i) \cap adj(j) \neq \emptyset. \end{aligned}$$

3.1 Bids and Asks and Trader Optimization

First we need to understand what bidding model we will use. Even when goods are indistinguishable, a trader may want to price-discriminate, and offer different bid and ask values to different sellers and buyers. In the case of distinguishable goods, we have to deal with a further complication: the trader has to name the good she is proposing to sell or buy, and can possibly offer multiple different products.

There are two variants of our model depending whether a trader makes a *single* bid or ask to a seller or buyer, or she offers a *menu* of options.

- (i) 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 .
- (ii) Alternatively, we can require that each trader t can make at most one ask to each seller and one bid for each buyer, and an ask has to include the product sold, and a bid has to offer a particular buyer to sell to.

Our results hold in either model. For notational simplicity we will use the menu option here.

Next we need to understand the optimization problem of a trader t . Suppose we have bid and ask values for all other traders $t' \in T, t' \neq t$. What are the best bid and ask offers trader t can make as a 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 assume p_j is 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 $B(t)$ that maximizes the total value v_{tij} for trader t .

We will need the dual of the linear program of finding the trade of maximum profit for the trader t . We will use q_{ti} as the dual variable associated with the constraint of seller or buyer i . The dual is then the following problem.

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

We view q_{ti} as the *profit made by t from trading with seller or buyer i* . Theorem 3.1 summarizes the above discussion.

Theorem 3.1 *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 above.*

3.2 Efficient Trade and Equilibrium

Now we can prove trade at equilibrium is always efficient.

Theorem 3.2 *Every equilibrium results in an efficient allocation of the goods.*

Proof. 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 used in Section 2.2.

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 a matching problem. Let q_{ti} for $i \in B(t) \cup S(t)$ denote the optimal dual solution to this matching problem as described by Theorem 3.1.

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. ■

Next we want to show that a non-crossing equilibrium always exists. We call an equilibrium non-crossing if the bid-ask offers a trader t makes for a seller-buyer pair (i, j) never cross, that is $\beta_{tij} \leq \alpha_{tji}$ for all t, i, j .

Theorem 3.3 *There exists a non-crossing equilibrium supporting any socially optimal trade.*

Proof. Consider an optimal trade x and a dual solution z as before. To find a non-crossing equilibrium we need to 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

$$p_i + q_{ti} + p_j + q_{tj} \geq \theta_{ji} - \theta_{ij}, \quad (1)$$

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 (1) 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 equality in the above constraint (1).

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 $t \in adj(t)$ having equality in (1). 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 (1) for the buyer j that he matches with i . If this is the only trader t that has a tight constraint in (1) 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 (1) 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 (1) 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 (1) 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 (1), 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. ■

3.3 Trader Profits

Finally we turn to the goal of understanding, in the case of general traders, how a trader's profit is affected by its position in the network.

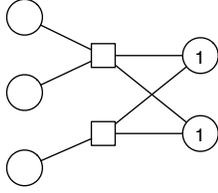


Figure 3: The top trader is essential for social welfare. Yet the only equilibrium is to have bid and ask values equal to 0, and the trader makes no profit.

First, we show how to maximize the total profit of a set of traders. 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 $z_i = p_i + \sum_{t \in \text{adj}(i)} q_{ti}$ is an optimal dual solution, and also satisfies the constraints (1) 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 3.3 shows how to turn this into an equilibrium.

Theorem 3.4 *The maximum value for $\sum_{t \in T'} \sum_i q_{ti}$ above is the maximum profit the set T' of traders can make.*

Proof. By the proof of Theorem 3.2 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 3.3. We modify that proof to start with profit vectors p and q_t for $t \in T'$, and set $q_t = 0$ for all traders $t \notin T'$. We verify that this starting solution satisfies the first three of the four required properties, and then we can follow the proof to make the fourth property true. We omit the details of this in the present version. ■

In Section 2.3 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 shown by the example in Figure 3.

However, we still get a characterization for when a trader t can make a positive profit.

Theorem 3.5 *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 — that is, if deleting agent t from $\text{adj}(i)$ decreases the value of the optimal allocation.*

Proof. 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 3.2 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 t from $\text{adj}(t)$. 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 t from $\text{adj}(t)$ 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. Now we use the same process as in the proof of Theorem 3.3 to create a non-crossing equilibrium starting with $p_i = z_i$ for all $i \in S \cup B$, and $q_{ti} = \gamma$, and all other q values 0. This creates an equilibrium with non-crossing bids where t makes at least γ profit (due to trade with seller i). ■

Finally, if we allow crossing equilibria, then we can find the minimum possible profit by simply finding a dual solution minimizing the dual variables associated with agents monopolized by some trader.

Theorem 3.6 *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.*

References

- [1] M. Babaioff, N. Nisan, E. Pavlov. Mechanisms for a Spatially Distributed Market. ACM EC Conference, 2005.
- [2] Kenneth Burdett and Kenneth Judd. Equilibrium Price Dispersions. *Econometrica*, 51/4, July 1983, 955-969.
- [3] L. Chu, Z.-J. Shen. Agent Competition Double Auction Mechanism. *Management Science*, 52/8, 2006.
- [4] G. Demange, D. Gale, M. Sotomayor. Multi-item auctions. *J. Political Econ.* 94(1986).
- [5] E. Even-Dar, M. Kearns, S. Suri. A Network Formation Game for Bipartite Exchange Economies. ACM-SIAM Symp. on Discrete Algorithms (SODA), 2007.
- [6] J. Kephart, J. Hanson, A. Greenwald. Dynamic Pricing by Software Agents. *Computer Networks*, 2000.
- [7] S. Kakade, M. Kearns, L. Ortiz, R. Pemantle, S. Suri. Economic Properties of Social Networks. NIPS 2004.
- [8] R. Kranton, D. Minehart. A Theory of Buyer-Seller Networks. *American Economic Review* 91(3), June 2001.
- [9] H. Leonard. Elicitation of Honest Preferences for the Assignment of Individuals to Positions. *J. Pol. Econ.* 1983.
- [10] M. E. J. Newman. The structure and function of complex networks. *SIAM Review*, 45:167-256, 2003.
- [11] M. O'Hara. *Market Microstructure Theory*. Blackwell Publishers, Cambridge, MA, 1995.
- [12] L. Shapley M. Shubik, The Assignment Game I: The Core. *Intl. J. Game Theory* 1/2 111-130, 1972.
- [13] Jean Tirole. *The Theory of Industrial Organization*. The MIT Press, Cambridge, MA, 1988.