

Chapter 11

Network Models of Markets with Intermediaries

11.1 Price-Setting in Markets

In Chapter 10 we developed an analysis of trade and prices on a bipartite graph consisting of buyers, sellers, and the edges connecting them. Most importantly, we showed that market-clearing prices exist, and that trade at these prices results in maximal total valuation among the buyers and sellers; and we found a procedure that allowed us to construct market-clearing prices. This analysis shows in a striking way how prices have the power to direct the allocation of goods in a desirable way. What it doesn't do is provide a clear picture of where prices in real markets tend to come from. That is, who sets the prices in real markets, and why do they choose the particular prices they do?

Auctions, which we discussed in Chapter 9, provide a concrete example of price determination in a controlled setting. In our discussion of auctions, we found that if a seller with a single object runs a second-price sealed-bid auction — or equivalently an ascending-bid auction — then buyers bid their true values for the seller's object. In that discussion, the buyers were choosing prices (via their bids) in a procedure selected by the seller. We could also consider a *procurement auction* in which the roles of buyers and sellers are reversed, with a single buyer interested in purchasing an object from one of several sellers. Here, our auction results imply that if the buyer runs a second-price sealed-bid auction (buying from the lowest bidder at the second-lowest price), or equivalently a descending-offer auction, then the sellers will offer to sell at their true costs. In this case, the sellers are choosing prices (their offers) in a procedure selected by the buyer.

D. Easley and J. Kleinberg. *Networks, Crowds, and Markets: Reasoning about a Highly Connected World*. To be published by Cambridge University Press, 2010. Draft version: October 23, 2009.

But who sets prices, and who trades with whom, if there are many buyers and many sellers? To get a feel for what happens, let's look first at how trade takes place in an actual market.

Trade with Intermediaries. In a wide range of markets, individual buyers and sellers do not interact directly with each other, but instead trade through intermediaries — brokers, market-makers, or middlemen who set the prices. This is true in settings that range from the trade of agricultural goods in developing countries to the trade of assets in financial markets.

To get a sense for how markets with intermediaries typically work, let's focus on the latter example, and consider how buyers and sellers interact in the stock market. In the U.S., buyers and sellers trade over a billion shares of stock daily. But there is no one market for trade in stocks in the U.S. Instead, trade occurs on multiple exchanges such as the New York Stock Exchange (NYSE) or the NASDAQ-OMX, as well as on alternative trading systems such as those run by Direct Edge, Goldman Sachs, or Investment Technologies Group (ITG), which arrange trades in stocks for their clients. These markets operate in various ways: some (such as NYSE or NASDAQ-OMX) determine prices that look very much like our market-clearing prices from Chapter 10, while others (like Direct Edge, Goldman, or ITG) simply match up orders to buy and sell stocks at prices determined in other markets. Some have people (called specialists in the NYSE) directly involved in setting prices, while others are purely electronic markets with prices set by algorithms; some trade continuously throughout the day, while others trade less frequently as they wait for batches of buy and sell orders to arrive; some allow anyone at least indirect access to the market, while others restrict the group of buyers and sellers that they will deal with (often to large institutional traders).

Many of these markets create something called an *order book* for each stock that they trade. An order book is simply a list of the orders that buyers and sellers have submitted for that stock. A trader might for instance submit an order to sell 100 shares if the price is at \$5 or more per share; another trader might submit an order to sell 100 shares if the price is at \$5.50 or more per share. Two other traders might submit orders to buy 100 shares if the price is no more than \$4 per share, and to buy 100 shares if the price is no more than \$3.50 per share. Orders of this type are called *limit orders*, since they are commitments to buy or sell only once the price reaches some limit set by the trader. If these were the only orders that existed, then the order book for this stock would look like Figure 11.1(a).

The highest outstanding offer to buy the stock is referred to as the current *bid* for the stock, while the lowest outstanding offer to sell it is referred to as the *ask*. If the market uses a specialist then this person knows the book of orders, and may choose to submit his or her own better offer to buy or sell out of inventory of the stock, which becomes the bid or ask respectively. For example, if Figure 11.1(a) describes the order book, then the specialist

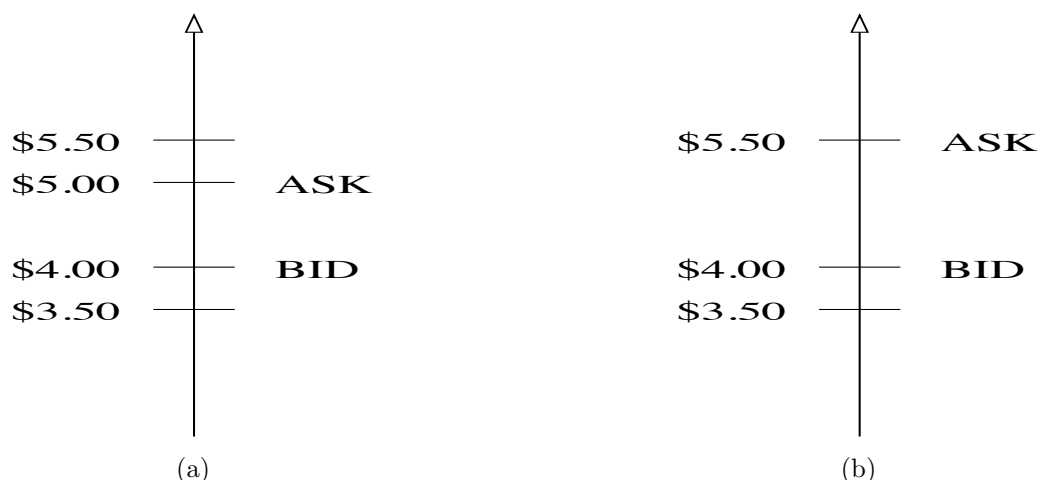


Figure 11.1: (a) A book of limit orders for a stock with a bid of \$4 and an ask of \$5. (b) A book of limit orders for a stock with a bid of \$4 and an ask of \$5.50.

may choose to display a bid of \$4.25, based on his or her own better offer, and an ask of \$5.00. These are the prices displayed to the trading public.

Retail traders (small traders who buy and sell stocks using their own assets) most often do not submit limit orders; instead they typically submit orders to buy or sell at the existing quotes — the current bid and ask. This type of order to trade immediately at market prices is called a *market order*. For example, if a trader submits a market order to buy 100 shares of a stock with the order book described by Figure 11.1(a), then the seller whose limit order was placed at \$5.00 sells 100 shares and the buyer who submitted the market order buys 100 shares at \$5.00. (Note that the seller can be either a member of the public or the specialist.) The new order book would then be as displayed Figure 11.1(b), and new ask would be \$5.50. This process continues throughout the trading day with new limit orders, specialist offers, and market orders arriving over time, and transactions being performed.

Of course, orders to buy or sell are not always for 100 shares, and in fact order sizes vary greatly. For example, if the order book is as depicted in Figure 11.1(a), and a market order to buy 200 shares arrives then both sellers on the book sell at their ask prices. The buyer will buy 100 shares at \$5.00 and 100 shares at \$5.50. We can think of this order as “walking up the book,” since executing it exposes multiple orders at different prices.

Large mutual funds such as Fidelity or Vanguard, and other institutional traders such as banks, pension funds, insurance companies and hedge funds, buy and sell a very large number of shares each day. They don’t really want to trade many small lots of shares with retail traders and, as in our 200-share example, walk up or down the book. They also don’t want to submit a single large limit order to the market, as then other market participants

will know their trading desires, and can take advantage of them.¹ Instead of submitting a single large market or limit order these traders use a variety of orders and trading venues. They typically split their order into many pieces and trade these pieces over the trading day, or over several days, in order to minimize the impact of their trading desire on the price. One way in which large traders hide their trading desires is to submit pieces of it to many different trading systems to which they have access. One particularly interesting group of trading systems are called *dark pools*. Examples of these alternative trading systems are Goldman Sachs's Sigma-X and the systems run by ITG. Access to these systems is limited and orders submitted to these systems are not displayed to the public. Instead these systems simply match orders submitted by their clients at prices established in the public market, and charge their clients a fee for the service. This is a relatively new, but growing segment of the market; in April 2009, for example, approximately 9% of the trade in U.S. equities was done on dark pools.

As you might imagine, the actual structure of the stock market is very complex and rapidly evolving. There are many trading systems, many types of orders that buyers and sellers can use, and a wide variety of market participants. The questions of how prices evolve over time and how they relate to the underlying fundamental value of the assets being traded are also important, and we have ignored these issues so far. We will discuss some aspects of the evolution of prices and their relation to the underlying values in Chapter 22; more detailed analyses of the stock market are carried out in a number of books [205, 207, 327].

The collection of different trading venues for stocks ultimately results in a variety of markets with restricted participation. So when we take into account the full set of trading options for all market participants — both large and small — we see a network structure emerge, connecting buyers and sellers to different possible intermediaries. A fundamental question is how to reason about trade when there are multiple markets connected by a network in this way. In the next section, we develop a network model for trade which abstracts away the specific details of the stock market, focusing on the general issue of how the underlying structure constrains who can trade with whom, and how prices are set by market participants.

¹A large order to buy, for example, may provide information to other market participants suggesting that the stock is currently undervalued and its price is likely to increase. These other market participants may then jump into the market, perhaps getting ahead of the execution of some part of the large order, and drive the price up quickly. This would harm the trader who submitted the large order as he may then have to pay more than expected for his order. This is related to broader questions about the role of information in markets, a topic we discuss in Chapter 22.

11.2 A Model of Trade on Networks

Our network model will be based on three fundamental principles that we saw in discussing the stock market: 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. Rather, the prices that each buyer and seller commands are determined in part by the range of alternatives that their respective network positions provide.

Before specifying the model, let's first look at another example of trade, in a very different setting, that exhibits these properties. This is the market for agricultural goods between local producers and consumers in a developing country. In many cases there are middlemen, or traders, who buy from farmers and then resell to consumers. Given the often poor transportation networks, the perishability of the products and limited access to capital by farmers, individual farmers can sell only to a limited number of intermediaries [45, 152]. Similarly, consumers can buy from only a limited number of intermediaries. A developing country may have many such partially overlapping local markets existing alongside modern, more global markets.

We can use a graph to describe the trading opportunities available to sellers, buyers, and middlemen (traders). Figure 11.2 depicts a simple example of such a trading network, superimposed on its geographic setting. Here we've labeled seller nodes with S, buyer nodes with B, and trader nodes with T; and we've placed an edge between any two agents who can trade with each other. Notice in this example that the seller and buyer on the right-hand margin of the picture only have access to the trader on their side of the river. The buyer at the top of the figure has access to both traders — perhaps he has a boat. You might imagine that the extra trading opportunities available to this buyer, and the similar extra trading opportunities available to the seller on the west bank of the river, would result in better prices for them. We will see that this is exactly what happens in the trading outcomes determined by our model in networks of this type.

Network Structure. We now describe a simple model of trade on a network which is general enough to incorporate important features of the trading and price-setting process for commodities as varied as financial assets traded in developed countries and agricultural goods in developing countries [62].

For the simplest form of the model, we don't try to address the issue of multiple goods for sale, or multiple possible quantities; instead, we assume there is a single type of good that comes in indivisible units. Each seller i initially holds one unit of the good which he values at v_i ; he is willing to sell it at any price that is at least v_i . Each buyer j values one copy of the good at v_j , and will try to obtain a copy of the good if she can do it by paying no more than v_j . No individual wants more than one copy of the good, so additional copies are valued at 0. All buyers, sellers, and traders are assumed to know these valuations. As a




Figure 11.2: Trading networks for agricultural markets can be based on geographic constraints, giving certain buyers (nodes labeled B) and sellers (nodes labeled S) greater access to traders (nodes labeled T).

result, this model is best thought of describing interaction between individuals who have a history of trade with each other, and hence know each other's willingness to pay for goods.

Trade takes place on a network that represents who can trade with whom. As in the example depicted in Figure 11.2 the nodes consist of buyers, sellers, and traders, with each edge representing an opportunity for trade. Since we are assuming that the traders act as intermediaries for the possible seller-buyer transactions, we require that each edge connects a buyer or seller to a trader. In Figure 11.3, we depict the same graph from Figure 11.2, redrawn to emphasize these features of the network model. (In all of our figures depicting trading networks we will use the following conventions. Sellers are represented by circles on the left, buyers are represented by circles on the right, and traders are represented by squares in the middle. The value that each seller and buyer places on a copy of the good is written next to the respective node that represents them.)

Beyond the fact that we now have intermediaries, there are a few other differences between this model and our model of matching markets from Chapter 10. First, we are assuming




Figure 11.3: A standardized view of the trading network from Figure 11.2: Sellers are represented by circles on the left, buyers are represented by circles on the right, and traders are represented by squares in the middle. The value that each seller and buyer places on a copy of the good is written next to the respective node that represents them.

that buyers have the same valuation for all copies of a good, whereas in matching markets we allowed buyers to have different valuations for the goods offered by different sellers. The model in this chapter can be extended to allow for valuations that vary across different copies of the good; things become more complicated, but the basic structure of the model and its conclusions remain largely the same. A second difference is that the network here is fixed and externally imposed by constraints such as geography (in agricultural markets) or eligibility to participate (in different financial markets). In matching markets, we began the chapter with fixed graphs such as this, but then focused the core of the analysis on preferred-seller graphs that were determined not by external forces but by the preferences of buyers with respect to an evolving set of prices.

Prices and the Flow of Goods. The flow of goods from sellers to buyers is determined by a game in which traders first set prices, and then sellers and buyers react to these prices.

Specifically, each trader t offers a *bid price* to each seller i that he is connected to; we will denote this bid price by b_{ti} . (The notation indicates that this is a price for a transaction between t and i). This bid price is an offer by t to buy i 's copy of the good at a value of b_{ti} . Similarly, each trader t offers an *ask price* to each buyer j that he is connected to. This ask price, denoted a_{tj} , is an offer by t to sell a copy of the good to buyer j at a value of a_{tj} . In Figure 11.4(a), we show an example of bid and ask prices on the graph from Figure 11.3.

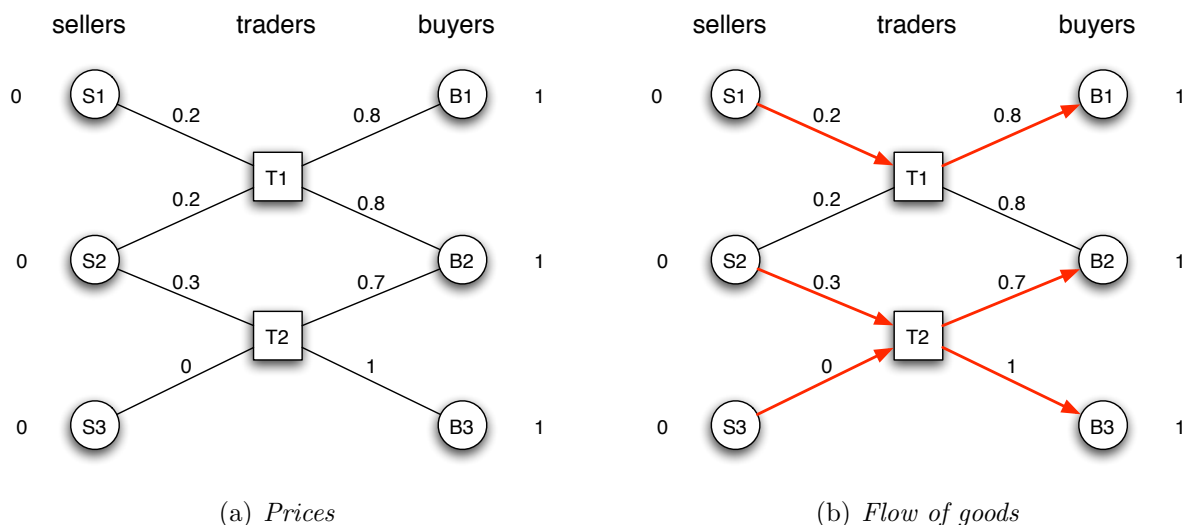


Figure 11.4: (a) Each trader posts bid prices to the sellers he is connected to, and ask prices to the buyers he is connected to. (b) This in turn determines a flow of goods, as sellers and buyers each choose the offer that is most favorable to them.

Once traders announce prices, each seller and buyer chooses at most one trader to deal with — each seller sells his copy of the good to the trader he selects (or keeps his copy of the good if he chooses not to sell it), and each buyer purchases a copy of the good from the trader she selects (or receives no copy of the good if she does not select a trader). This determines a flow of goods from sellers, through traders, to buyers; Figure 11.4(b) depicts such a flow of goods, with the sellers' and buyers' choices of traders indicated by the edges with arrows on them.

Because each seller has only one copy of the good, and each buyer only wants one copy, at most one copy of the good moves along any edge in the network. On the other hand, there is no limit on the number of copies of the good that can pass through a single trader node. Note that a trader can only sell as many goods to buyers as he receives from sellers; we will include in the model a large penalty imposed on a trader who defaults on an offer to sell to a buyer as a result of not having enough goods on hand. Due to this, there are strong incentives for a trader not to produce bid and ask prices that cause more buyers than sellers to accept his offers. There are also incentives for a trader not to be caught in the reverse difficulty, with more sellers than buyers accepting his offers — in this case, he ends up with excess inventory that he cannot sell. We will see that neither of these will happen in the solutions we consider; traders will choose bid and ask prices such that the number of goods they receive from sellers is equal to the number of goods they pass on to buyers.

Finally, notice something else about the flow of goods in this example: seller $S3$ accepts the bid even though it is equal to his value, and likewise buyer $B3$ accepts the ask even

though it is equal to his value. In fact, each of $S3$ and $B3$ is *indifferent* between accepting and rejecting the offer. Our assumption in this model is that when a seller or buyer is indifferent between accepting or rejecting, then we (as the modelers) can choose either alternative as the outcome that actually happens. Finding a way to handle indifference is an important aspect in most market models, since transactions will typically take place right at the boundary of an individual's willingness to trade. This is similar to the tie-breaking issue inherent in the formulation of market-clearing prices in Chapter 10 as well. An alternate way to handle indifference in the present case is to assume a miniscule positive amount of payoff (e.g. a penny) that is required for an agent to be willing to trade, in which case we would see bid and ask values like 0.01 and 0.99. While this makes the tie-breaking decision more explicit, the model becomes much messier and ultimately harder to reason about. As a result, we will stick to the approach where we allow trades at zero payoff, with ties broken as needed; in doing so, we will remember that this is essentially a formal way to represent the idea of a price or a profit margin being driven to (almost) zero. And if it makes things simpler to think about, whenever you see an indifferent seller or buyer choosing to transact or not, you can imagine the price being shifted either upward or downward by 0.01 to account for this decision by the seller or buyer.

Payoffs. Recall that specifying a game requires a description of the strategies and the payoffs. We have already discussed the strategies: a trader's strategy is a choice of bid and ask prices to propose to each neighboring seller and buyer; a seller or buyer's strategy is a choice of a neighboring trader to deal with, or the decision not to take part in a transaction.

The payoffs follow naturally from the discussion thus far.

- A trader's payoff is the profit he makes from all his transactions: it is the sum of the ask prices of his accepted offers to buyers, minus the sum of the bid prices of his accepted offers to sellers. (As discussed above, we also subtract a large penalty if the trader has more accepted asks than bids, but the effect of this is primarily to ensure that traders will never expose themselves to this situation in the solutions we consider.)
- For a seller i , the payoff from selecting trader t is b_{ti} , while the payoff from selecting no trader is v_i . In the former case, the seller receives b_{ti} units of money, while in the latter he keeps his copy of the good, which he values at v_i . (We will consider only cases in which all the seller v_i 's are 0.)
- For each buyer j , the payoff from selecting trader t is $v_j - a_{tj}$, while the payoff from selecting no trader is 0. In the former case, the buyer receives the good but gives up a_{tj} units of money.

So for example, with prices and the flow of goods as in Figure 11.4(b), the payoff to the first trader is $(0.8 - 0.2) = 0.6$ while the payoffs to the second trader is $(0.7 + 1 - 0.3 - 0) = 1.4$.




Figure 11.5: Relative to the choice of strategies in Figure 11.4(b), trader $T1$ has a way to improve his payoff by undercutting $T2$ and performing the transaction that moves $S2$'s copy of the good to $B2$.

The payoffs to the three sellers are 0.2, 0.3, and 0 respectively, while the payoffs to the three buyers are $1 - 0.8 = 0.2$, $1 - 0.7 = 0.3$, and $1 - 1 = 0$ respectively,

The game we've defined here has a further important feature, which forms a contrast with other games we have discussed earlier. In earlier games, all players moved (i.e. executed their chosen strategies) simultaneously, while in this game the moves happen in two stages. In the first stage, all the traders simultaneously choose bid and ask prices. In the second stage, all the sellers and buyers then simultaneously choose traders to deal with. For us, this two-stage structure will not make things too complicated, particularly since the second stage is extremely simple: the best response for each seller and buyer is always simply to choose the trader with the best offer, and so we can essentially view the sellers and buyers as “drones” who are hard-wired to follow this rule. Still, we will have to take the two-stage structure into account when we consider the equilibria for this game, which we do next.

Best Responses and Equilibrium. Let's think about the strategies that the two traders have chosen in Figure 11.4(b). The upper trader $T1$ is making several bad decisions. First, because of the offers he is making to seller $S2$ and buyer $B2$, he is losing out on this deal to the lower trader $T2$. If for example he were to raise his bid to seller $S2$ to 0.4, and lower his ask to buyer $B2$ to 0.6, then he'd take the trade away from trader $T2$: seller $S2$ and buyer $B2$ would both choose him, and he'd make a profit of 0.2.

Second, and even more simply, there is no reason for trader $T1$ not to lower his bid to

seller $S1$, and raise his ask to buyer $B1$. Even with worse offers, $S1$ and $B1$ will still want to deal with $T1$, since they have no other options aside from choosing not to transact. Given this, $T1$ will make more money with a lower bid to $S1$ and a higher ask to $B1$. Figure 11.5 shows the results of a deviation by the upper trader that takes both of these points into account; his payoff has now increased to $(1 + 0.6 - 0 - 0.4) = 1.2$. Note that seller $S1$ and buyer $B1$ are now indifferent between performing the transaction or not, and as discussed earlier, we are assuming that sellers and buyers will perform transactions in such situations.

This discussion motivates the equilibrium concept we will use for this game, which is a generalization of Nash equilibrium. As in the standard notion of Nash equilibrium from Chapter 6, it will be based on a set of strategies such that each player is choosing a best response to what all the other players are doing. However, the definition also needs to take the two-stage structure of the game into account.

To do this, we first think about the problem faced by the buyers and sellers in the second stage, after traders have already posted prices. Here, we have a standard game among the buyers and sellers, and each of them chooses a strategy that is a best response to what all other players are doing. Next, we think about the problem faced by the traders, in deciding what prices to post in the first stage. Here, each trader chooses a strategy that is a best response both to the strategies the sellers and buyers will use (what bids and asks they will accept) and the strategies the other traders use (what bids and asks they post). So everyone is employing a best response just as in any Nash equilibrium. The one difference here is that since the sellers and buyers move second they are required to choose optimally given *whatever* prices the traders have posted, and the traders know this. This equilibrium is called a *subgame perfect Nash equilibrium*; in this chapter, we will simply refer to it as an *equilibrium*.²

The two-stage nature of our game here is particularly easy to think about, since the behavior of sellers and buyers is very simple. Thus, for purposes of reasoning about equilibria, we can mainly think about the strategies of the traders in the first stage, just as in a simultaneous-move game, knowing that sellers and buyers will simply choose the best offers (possibly with tie-breaking) once the traders post prices.

In the next section, we'll work out the set of possible equilibria for the trading network in Figures 11.3–11.5, by first dissecting the network into simpler “building blocks.” In particular, these building blocks will correspond to two of the basic structures contained within the network in Figures 11.3–11.5: buyers and sellers who are *monopolized* by having only a single trader they can deal with, and buyers and sellers who benefit from *perfect competition* between multiple traders. In the process, we'll see that network structure and

²The word “subgame” refers to the fact that once traders post prices, the buyers and sellers are faced with a new free-standing game in the second stage. The word “perfect” refers to the requirement that in the subgame, the players who have choices remaining are required to behave optimally given the choices that have already been made. This concept is discussed in more detail in Section 6.10.




Figure 11.6: A simple example of a trading network in which the trader has a monopoly and extracts all of the surplus from trade.

access to alternatives can significantly affect the power of participants in the market.

11.3 Equilibria in Trading Networks

We now discuss the process of analyzing equilibria in trading networks. We begin with simple network structures and build up to the example from the previous section. Following our plan, we'll begin by considering simple networks corresponding to monopoly and perfect competition.

Monopoly. Buyers and sellers are subject to monopoly in our model when they have access to only a single trader. Perhaps the simplest example of this is depicted in Figure 11.6. Here we have one seller who values the good at 0, one trader, and one buyer who values the good at 1.

In this trading network the trader is in a monopoly position relative to both the seller and the buyer (there is only one trader available to each of them). The only equilibrium is for the trader to set a bid of 0 to the seller and an ask of 1 to the buyer; the seller and buyer will accept these prices and so the good will flow from the seller to the trader and then on to the buyer. Note that we are using the indifference of the seller and buyer as in the example from the previous section: since the seller and buyer are indifferent between engaging in a transaction or not, we as the modelers are choosing the outcome and having them perform the transaction.

To see why this is the only equilibrium, we simply notice that with any other bid and ask between 0 and 1, the trader could slightly lower the bid or raise the ask, thereby performing the transaction at a higher profit.

Perfect Competition. Now let's look at a basic example showing perfect competition between two traders, as depicted in Figure 11.7.




Figure 11.7: A trading network in which there is perfect competition between the two traders, $T1$ and $T2$. The equilibrium has a common bid and ask of x , where x can be any real number between 0 and 1.

In Figure 11.7 there is competition between traders $T1$ and $T2$ to buy the copy of the good from $S1$ and sell it to $B1$. To help in thinking about what forms an equilibrium, let's first think about things that are out of equilibrium, in a manner similar to what we saw in Figure 11.5. In particular, suppose trader $T1$ is performing the trade and making a positive profit: suppose his bid to the seller is some number b , and his ask to the buyer is a number $a > b$. Since $T2$ is not performing the trade, he currently has a payoff of zero. But then it must be that $T2$'s current strategy is not a best response to what $T1$ is doing: $T2$ could instead offer a bid slightly above b and an ask slightly below a , thereby taking the trade away from $T1$ and receiving a positive payoff instead of zero.

So it follows that whichever trader is performing the trade at equilibrium must have a payoff of 0: he must be offering the same value x as his bid and ask. Suppose that trader $T1$ is performing the trade. Notice that this equilibrium involves indifference on his part: he is indifferent between performing the trade at zero profit and not performing the trade. As in the earlier case of indifference by sellers and buyers, we assume that we (the modelers) can choose an outcome in this case, and we will assume that the transaction is performed. Here too, we could handle indifference by assuming a minimum increment of money (e.g. 0.01), and having the transaction take place with a bid and ask of $x - 0.01$ and x respectively, but again handling indifference via zero payoffs (and keeping in mind that they are designed to model profits that come arbitrarily close to 0) makes the analysis simpler without affecting the outcome.

Next we want to argue that the trader not performing the trade at equilibrium ($T2$ in

this case) must also have bid and ask values of x . First, notice that in equilibrium, we cannot have a trader buy the good from the seller without also selling it to the buyer; therefore, $T2$ must be offering a bid $b \leq x$, (or else the seller would sell to $T2$) and an ask $a \geq x$. (or else the buyer would buy from $T2$). But if this bid and ask were not the same — that is, if $a < b$ — then $T1$ could lower his bid or raise his ask so that they still lie strictly between a and b . In that case $T1$ could perform the trade while making a positive profit, and hence his current strategy of bidding and asking x would not be a best response to what $T2$ is doing.

So the equilibrium occurs at a common bid and ask of x . What can we say about the value of x ? It clearly has to be between 0 and 1: otherwise either the seller wants to sell but the buyer wouldn't want to buy, or conversely the buyer wants to buy but the seller wouldn't want to sell. In fact this is all that we can say about x . Any equilibrium consists of a common bid and ask by each trader, and flow of goods from the seller to the buyer through one of the traders. A key feature of the equilibrium is that the seller sells to the same trader that the buyer buys from: this is another kind of coordination in the face of indifference that is reminiscent of the tie-breaking issues in market-clearing prices from Chapter 10. It is also interesting that while traders make no profit in any equilibrium, the choice of equilibrium — captured in the value of x — determines which of the seller or buyer receives a higher payoff. It ranges from the extreme cases of $x = 0$ (where the buyer consumes all the available payoff) and $x = 1$ (where the seller consumes it) to the intermediate value of $x = \frac{1}{2}$ (where the seller and buyer receive equal payoffs). In the end, the choice of equilibrium reflects something about the relative power of the seller and buyer that can only be inferred by looking outside the formulation of the trading game — the game itself can determine only the range of possible equilibria.

The Network from Section 11.2. Using the networks in Figures 11.6 and 11.7 as building blocks, it is not hard to work out the equilibria in the example from Section 11.2. This is illustrated in Figure 11.8. Sellers $S1$ and $S3$, and buyers $B1$ and $B3$, are monopolized by their respective traders, and so in any equilibrium these traders will drive the bids and asks all the way to 0 and 1 respectively.

Seller $S2$ and buyer $B2$, on the other hand, benefit from perfect competition between the two traders. Here the argument follows what we used in analyzing the simpler network in Figure 11.7: the trader performing the transaction must have bid and ask values equal to the same number x (for some real number x between 0 and 1), or else the other trader could take the trade away from him; and given this, the other trader must also have bid and ask values equal to x .

These types of reasoning are useful in analyzing other more complex networks as well. When you see a seller or buyer connected to only a single trader, they will receive zero payoff in any equilibrium, since the trader will drive the bid or ask to as extreme a value as possible.




Figure 11.8: The equilibria for the trading network from Section 11.2. This network can be analyzed using the ideas from the simpler networks representing monopoly and perfect competition.

On the other hand, when two traders both connect the same seller and buyer, then neither can make a positive profit in conveying a good from this seller to this buyer: if one trader performed the trade at a positive profit, the other could undercut them.

We now consider an example illustrating how the network structure can also produce more complex effects that are not explained by these two principles.

Implicit Perfect Competition. In our examples so far, when a trader makes no profit from a transaction, it is always because there is another trader who can precisely replicate the transaction — i.e., a trader who is connected to the same seller and buyer. However, it turns out that traders can make zero profit for reasons based more on the global structure of the network, rather than on direct competition with any one trader.

The network in Figure 11.9 illustrates how this can arise. In this trading network there is no direct competition for any one “trade route” from a seller to a buyer. However, in any equilibrium, all bid and ask prices take on some common value x between 0 and 1, and the goods flow from the sellers to the buyers. So all traders again make zero profit.

It is easy to see that this is an equilibrium: we can simply check that each trader is using a best response to all the other traders’ strategies. It takes a bit more work to verify that in every equilibrium, all bid and ask prices are the same value x : this is most easily done by checking alternatives in which some trader posts a bid that is less than the corresponding ask, and identifying a deviation that arises.




Figure 11.9: 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.

11.4 Further Equilibrium Phenomena: Auctions and Ripple Effects

The network model we’ve been considering is expressive enough that it can represent a diverse set of other phenomena. Here we consider two distinct examples: the first showing how the second-price auction for a single item arises from a trading network equilibrium, and the second exploring how small changes to a network can produce effects that ripple to other nodes.

Second-price auctions. Figure 11.10 shows how we can represent the structure of a single-item auction using a trading network. Suppose there is a single individual $S1$ with an item to sell, and four potential buyers who value the item at values w, x, y, z , listed in descending order $w > x > y > z$. We use four buyers in this example but our analysis would work for an arbitrary number of buyers.

In keeping with our model in which trading happens through intermediaries, we assume that each buyer is represented by a distinct trader — essentially, someone who serves as the buyer’s “proxy” for the transaction. This gives us a trading network as depicted in Figure 11.10(a).

Now let’s consider a possible equilibrium for this network. Trader $T1$ has the ability to

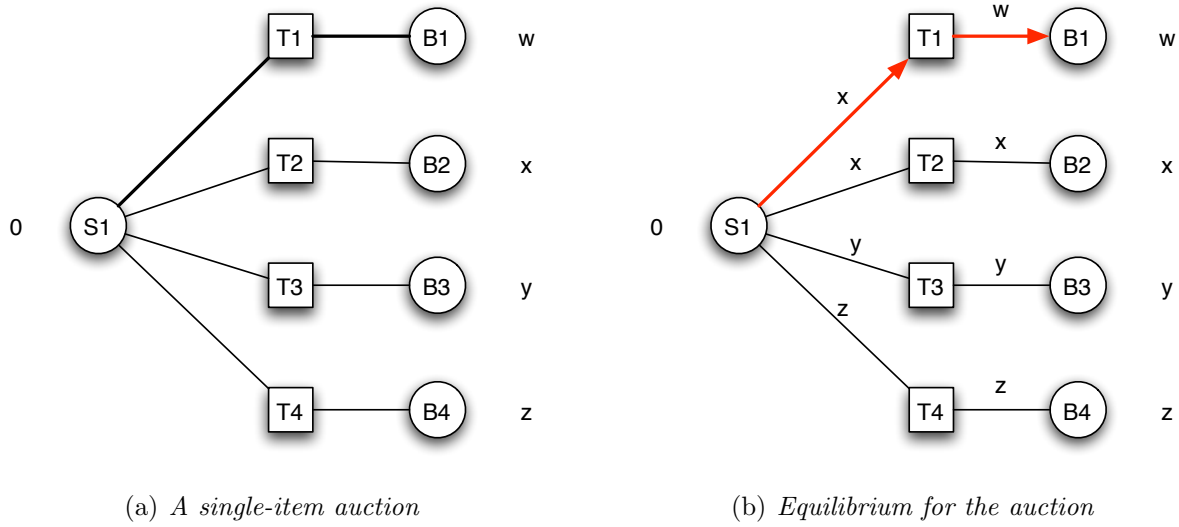


Figure 11.10: (a) A single-item auction can be represented using a trading network. (b) Equilibrium prices and flow of goods. The resulting equilibrium implements the second-price rule from Chapter 9.

outbid all the other traders, since he has ability to sell to his buyer for an ask up to w . In equilibrium, he will outbid them by the minimum he needs to in order to make the trade, which he can do by offering x to outbid $T2$. Here we use indifference to assume that the sale at x will go to $T1$ rather than $T2$, that buyer $B1$ will buy from $T1$ at a price of w , and that buyers $B2$ through $B4$ will choose not to buy the good from their respective traders.

We therefore get the equilibrium depicted in Figure 11.10(b). Notice how this equilibrium has exactly the form of a second-price auction, in which the item goes to the highest bidder, with the seller receiving the second-highest valuation in payment.³ What's interesting is that the second-price rule wasn't in any sense "built in" to the formulation of the auction; it emerged naturally as an equilibrium in our network representation.

³With a little more work, we can describe the full set of equilibria for this network, and show that the second-price rule is unique over equilibria that avoid a certain "pathological" structure, as follows. In any equilibrium, the good flows from the seller to buyer $B1$, and each trader offers to sell the good to his monopolized buyer for the buyer's value. In the equilibrium we consider, $T1$ and $T2$ both bid x . However, there are other bids that can be part of an equilibrium: Essentially, as long as one of traders $T2, T3$, or $T4$ bids between x and w , and $T1$ matches this bid, we have an equilibrium. If the highest bid among $T2, T3$, and $T4$ is strictly greater than x , then we have a situation in which this high-bidding trader has a *crossing* pair of bid and ask values: his bid is higher than his corresponding ask. This is an equilibrium, since $T1$ still makes the trade, the trader with the crossing bid/ask pair doesn't lose money, and no one has an incentive to deviate. However, it is a pathological kind of equilibrium since there is a trader who is offering to buy for more than he is offering to sell [62]. Thus, if we consider only equilibria without such crossing pairs of bid/ask values, then $T2$ bids the "second-price value" x , and this is what the good sells for. So the second-price rule is unique over equilibria without crossing pairs.

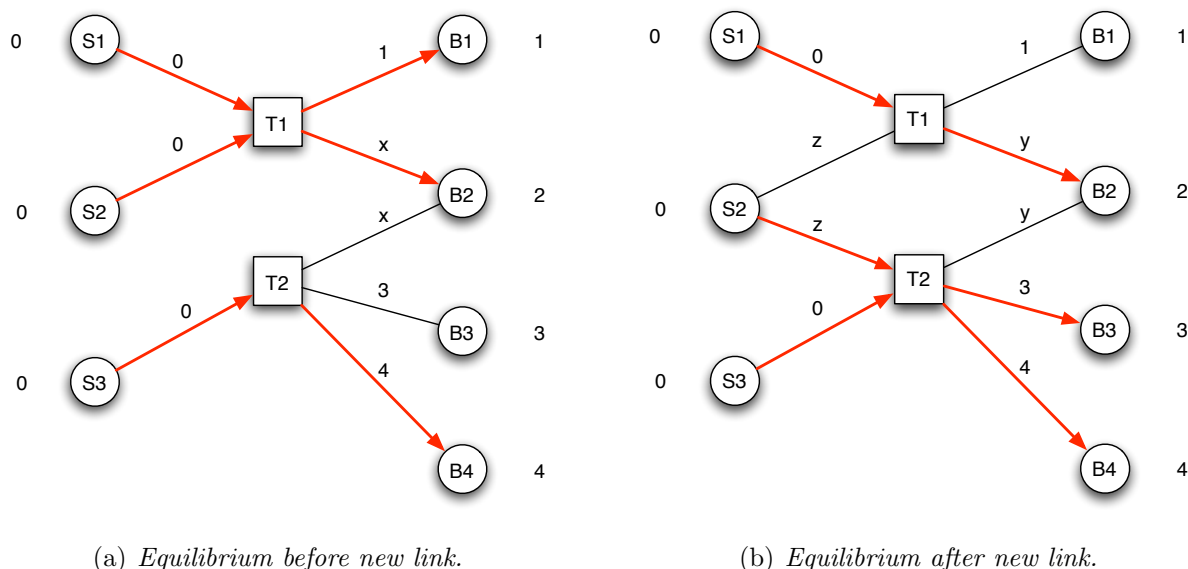


Figure 11.11: (a) Equilibrium before the new $S2$ - $T2$ link is added. (b) When the $S2$ - $T2$ edge is added, a number of changes take place in the equilibrium. Among these changes is the fact that buyer $B1$ no longer gets a copy of the good, and $B3$ gets one instead.

Ripple Effects from Changes to a Network. Our network model also allows us to explore how small changes to the network structure can affect the payoffs of nodes that are not directly involved in the change. This suggests a way of reasoning about how “shocks” to highly interconnected trading networks can ripple to more distant parts of the network.

This is a very general issue; we consider it here in a specific example where we can see concretely how such effects can arise. Consider the pair of networks in Figure 11.11: the second network is obtained from the first simply by adding a link from $S2$ to $T2$. We first work out the equilibria for each of these networks, and then consider how they differ.

In Figure 11.11(a), all sellers and buyers are monopolized except for $B2$, so their payoffs will all be zero; we use indifference to assume that $B3$ will not buy the good but $B1$ and $B4$ will. (As in all our examples, we can view this as modeling the fact that $T2$ can charge a price very slightly above 3 to $B3$, dissuading $B3$ from purchasing.) The one part that needs additional analysis is the pair of asks charged to $B2$. In equilibrium, these must be the same (otherwise, the trader making the sale could slightly raise his ask), and this common value x can be anything between 0 and 2. We exploit indifference here to assume that $B2$ buys from trader $T1$. Note that there cannot be an equilibrium in which buyer $B2$ buys from trader $T2$, since $B2$ can pay only 2 while trader $T2$ can sell the unit of the good he is able to buy at a price of and 4.

In Figure 11.11(b), once the edge from $S2$ to $T2$ has been added, we need to work out the equilibrium bids and asks charged to $S2$ and $B2$, and the flow of goods. Reasoning about

these bids and asks requires a bit more work than we've seen in previous examples, so we build up to it in a sequence of steps.

- The two bids to $S2$ must be the same as each other, since otherwise the trader getting the good could slightly lower his bid; for a similar reason, the two asks to $B2$ must also be the same as each other. Let's call the common bid z and the common ask y .
- We can next determine how the seller-trader transactions work out in equilibrium. $S2$ will sell to $T2$ rather than $T1$ in equilibrium: if $S2$ were selling to $T1$, and $T1$ were receiving a positive payoff from this transaction, then $S2$ would be selling for at most 2. In this case, $T2$ could slightly outbid $T1$ and sell $S2$'s copy of the good to $B3$. So in equilibrium, $T2$ buys two copies of the good, while $T1$ buys only one.
- Now let's figure out the possible values for the ask y . The ask y must be at least 1: otherwise, one of the traders is selling to $B2$ for a low price, and the trader performing this sale has an alternate trader whom he monopolizes, and from whom he would get a higher payoff. This can't happen in equilibrium, so y is at least 1.

Also, the ask y cannot be above 2 in equilibrium: in this case, $B2$ would not buy, and so $T1$ could perform a payoff-improving change in strategy by lowering his ask to $B2$, thereby getting $B2$ to buy from $T1$ for a price between 1 and 2.

- Next, we determine how the trader-buyer transactions work out in equilibrium. We've already concluded that $T2$ is buying two copies of the good, and so he maximizes his payoff by selling them to $B3$ and $B4$. Therefore $T2$ is not selling to $B2$ in equilibrium. Since, the ask y is at least 1, trader $T1$ will buy from $S1$ and sell to $B2$.
- Finally, what do we know about the value of z ? It has to be at least 1, or else $T1$ could outbid $T2$ for $S2$'s copy of the good, and receive a positive payoff by selling it to $B1$. It also has to be at most 3, or else $T2$ would prefer not to buy it from $S2$.

This sums up the analysis: in equilibrium, the bid z can be anything between 1 and 3, the ask y can be anything between 1 and 2, and the goods flow from $S1$ through $T1$ to $B2$, and from $S2$ and $S3$ through $T2$ to $B3$ and $B4$. Notice that this flow of goods maximizes the total valuation of the buyers who obtain the good, given the constraints on trade imposed by the network. We will see later in this chapter that any equilibrium has this efficiency property.

Let's consider at a high level what's going on in this pair of examples. In Figure 11.11(a), trader $T2$ has access to a set of buyers who want copies of the good very badly (they value it highly), but his access to sellers is very limited. $T1$ on the other hand is able to use all his available trading opportunities. The market in this respect has a "bottleneck" that is restricting the flow of goods.

Once $S2$ and $T2$ form a link, creating the network in Figure 11.11(b), a number of things change. First, and most noticeably, buyer $B3$ now gets a copy of the good while $B1$ doesn't. Essentially, the bottleneck in the market has been broken open, so the high-value buyers can now obtain the good at the expense of the low-value buyers. From $B1$'s perspective, this is a “non-local” effect: a link formed between two nodes, neither of which are neighbors of hers, has caused her to no longer be able to obtain the good.

There are other changes as well. Seller $S2$ is now in a much more powerful position, and will command a significantly higher price (since y is at least 1 in any equilibrium). Moreover, the range of possible equilibrium asks to $B2$ has been reduced from the interval $[0, 2]$ to the interval $[1, 2]$. So in particular, if we were previously in an equilibrium where the ask to $B2$ was a value $x < 1$, then this equilibrium gets disrupted and replaced by one in which the ask is a higher number $y \geq 1$. This indicates a subtle way in which $B2$ was implicitly benefitting from the weak position of the sellers, which has been now strengthened by the creation of the edge between $S2$ and $T2$.

This is a simple example, but it already illustrates some of the complexities that can arise when the structure of a trading network changes to alleviate (or create) bottlenecks for the flow of goods. With more work, one can create examples where the effects of changes in the network ripple much farther through the structure.

This style of reasoning also points to a line of questions in which we view the network as something malleable, and partially under the control of the market participants. For example, how much should $S2$ and $T2$ be willing to spend to create a link between each other, shifting the network structure from the one in Figure 11.11(a) to the one in Figure 11.11(b)? More generally, how should different nodes evaluate the trade-offs between investing resources to create and maintain links, and the benefits they get in terms of increased payoffs? This is a question that has been considered in other models of trading networks [149, 258], and it is part of a much broader research activity that investigates the formation of networks as a game-theoretic activity under a variety of different kinds of payoffs [20, 38, 120, 151, 224, 380].

11.5 Social Welfare in Trading Networks

When we've looked at games in earlier settings, we've considered not just equilibrium solutions, but also the question of whether these solutions are *socially optimal*. That is, do they maximize *social welfare*, the sum of the payoffs of all players?

In the context of our game, each good that moves from a seller i to a buyer j contributes $v_j - v_i$ to the social welfare. This is how much more j values the good than i , and the money that is spent in moving the good from i to j is simply transferred from one player to another, creating a net effect of zero to the total payoff. In more detail, if the good moves through trader t , who offers a bid of b_{ti} to i and an ask of a_{tj} to j , then the sum of the payoffs of i

and j , plus the portion of t 's payoffs arising from this transaction, is equal to

$$(b_{ti} - v_i) + (a_{tj} - b_{ti}) + (v_j - a_{tj}) = v_j - v_i.$$

Thus the social welfare is simply the sum of $v_j - v_i$ over all goods that move from a seller i to a buyer j . This makes sense, since it reflects how much happier, in total, the new owners of the goods are compared to the original owners of the goods. The maximum value of this quantity over all possible flows of goods — i.e., the socially optimal value — depends not just on the valuations of the sellers and buyers, but also on the network structure. Networks that are more richly connected can potentially allow a flow of goods achieving a higher social welfare than networks that are more sparsely connected, with bottlenecks that prevent desirable flows of goods.

For example, let's go back to the pair of networks in Figure 11.11. In each case, the equilibrium achieves a flow of goods that achieves the social optimum. In Figure 11.11(a), the best possible value of the social welfare is $1 + 2 + 4 = 7$, since there is no way to use the network to get copies of the goods to both $B3$ and $B4$. However, when the single edge from $S2$ to $T2$ is added, it suddenly becomes possible for both of these buyers to receive copies of the good, and so the value of the social welfare increases to $2 + 3 + 4 = 9$. This provides a simple illustration of how a more richly connected network structure can enable greater social welfare from trade.

In our discussion of social optimality we count the gains to traders as part of the social welfare (since they are part of the society of players, along with sellers and buyers). In the next section, we will consider how the total payoffs are divided between sellers, buyers, and traders, and how it depends on the network structure.

Equilibria and Social Welfare. In both the networks in Figure 11.11, the flow of goods achieving maximum social welfare can be achieved by an equilibrium.

In fact, this holds for all the examples we've seen thus far, and it's a fact that's true in general: it can be shown that in every trading network, there is always at least one equilibrium, and every equilibrium produces a flow of goods that achieves the social optimum [62]. While we won't go into the details of the proof here, it is similar in structure to the Existence and Optimality of Market-Clearing Prices that we discussed in the previous chapter. There too, without intermediaries present, we were able to show that prices achieving a certain type of equilibrium always exist (in that case the market-clearing property), and that all such prices produce an allocation that maximize social welfare.

11.6 Trader Profits

We now consider the question of how the social welfare in an equilibrium is divided up as payoffs among the sellers, buyers, and traders. In particular, the examples we've studied




Figure 11.12: Whether a trader can make a profit may depend on the choice of equilibrium. In this trading network, when $x = 1$, traders $T1$ and $T5$ make a profit, while when $x = 0$, only trader $T3$ makes a profit.

so far suggest the informal principle that as the network becomes more richly connected, individual traders have less and less power, and their payoffs go down. Understanding this more precisely points to a basic question that can be expressed in terms of our network model: what is the structural basis of perfect competition?

Our examples suggest that in order to make a profit (i.e. a positive payoff), a trader must in some way be “essential” to the functioning of the trading network. Certainly, if there is another trader who can replicate his function completely, then he cannot make a profit; nor can he make a profit in more complex settings like the implicit perfect competition in Figure 11.9. In fact, it will turn out that a version of this “essentiality” principle is true; but it is a bit more subtle than it might initially appear. To motivate it, we start with two illuminating examples.

First, whether a trader makes a profit can depend on the equilibrium: in some networks, it can be possible for a trader to make a profit in some equilibria but not others. Figure 11.12 shows how this can occur. Any choice of x between 0 and 1 will result in an equilibrium, with the traders $T2$ and $T4$ who are left out of trades serving to “lock” the value of x in place. However, when $x = 1$, traders $T1$ and $T5$ make a profit, while when $x = 0$, only trader $T3$ makes a profit. Furthermore, while every equilibrium produces a flow of goods to all three sellers, for a social welfare of 3, the amount of this social welfare that goes to the buyers and sellers — rather than the traders — varies between 1 and 2 as x ranges from 1 to 0.

The second example, in Figure 11.13, is even more counter-intuitive. Here, traders $T1$ and $T2$ both have monopoly power over their respective sellers, and yet their profits are

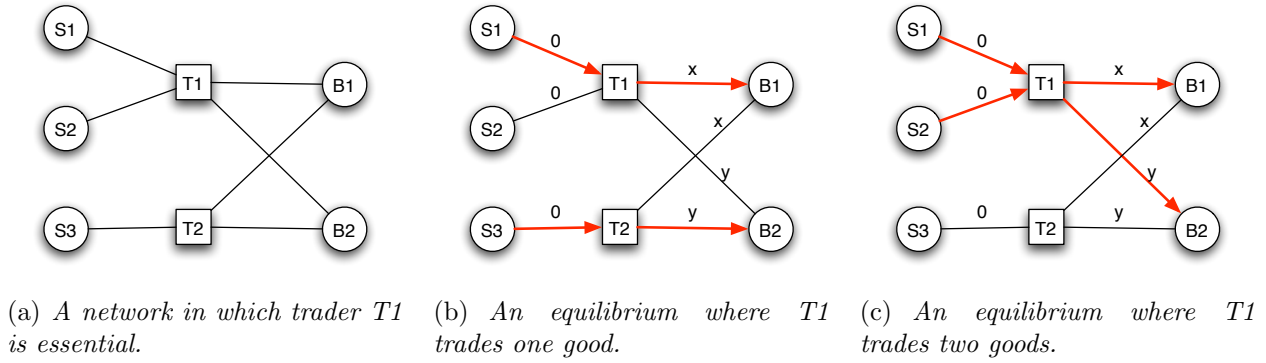


Figure 11.13: Despite a form of monopoly power in this network, neither trader can make a profit in any equilibrium.

zero in *every* equilibrium. We can verify this fact as follows. First, we notice that any equilibrium must look like one of the solutions in Figure 11.13(b) or Figure 11.13(c). The sellers are monopolized and will get bids of 0. For each buyer, the two asks must be the same, since otherwise the trader making the sale could slightly raise his ask. Now, finally, notice that if the common ask to either buyer were positive, then the trader left out of the trade on the higher one has a profitable deviation by slightly undercutting this ask.

Therefore, in this example, all bids and asks equal 0 in any equilibrium, and so neither trader profits. This happens despite the monopoly power of the traders — and moreover $T1$ fails to make a profit despite the fact that $T2$ can only perform one trade on his own. We can interpret this as a situation in which a small trader competes with a larger trader across his full set of potential buyers, despite having access to an insufficient set of sellers to actually perform all the available trades — a situation in which “the threat is stronger than its execution.” While this fits naturally within the scope of the model, examples such as this one also suggest natural extensions to the model, in which each trader has an intrinsic limit on the number of trades he can perform, and this affects the behavior of competing traders. While the possibility of such intrinsic limits haven’t played a role in our earlier examples, Figure 11.13 suggests that allowing for such limits could change the outcome in certain settings.

With these examples in mind, let’s return to the question of when, for a given trader T in a network, there exists an equilibrium in which T receives a positive payoff. It turns out that there exists such an equilibrium precisely when T has an edge e to a seller or buyer such that deleting e would change the value of the social optimum. In such a situation, we say that e is an *essential edge* from T to the other node. The proof of this statement is somewhat involved, and we refer the reader to [62] for the details. Figures 11.6 and 11.8 are examples of trading networks in which each trader has an essential edge and thus makes a

profit in equilibrium, while Figures 11.7 and 11.9 are examples in which no trader has an essential edge and no trader makes a profit.

This essential-edge condition is a stronger form of monopoly power than we saw in Figure 11.13(a). There, although deleting the node $T1$ would change the value of the social optimum, there is no single edge whose deletion would reduce the value of the social optimum below 2; rather, after the removal of any one edge, there would still be a flow of goods to both buyers. This is the crux of why $T1$ is not able to make a profit in Figure 11.13(a), despite his powerful position.

The example in Figure 11.12 also shows that this condition only implies a profit in *some* equilibrium, as opposed to every equilibrium. In Figure 11.12, the available profit essentially “slides” smoothly from one trader to another as we vary the value of x in the equilibrium.

11.7 Reflections on Trade with Intermediaries

In closing, it is useful to reflect on how our analysis of trade on networks relates to the motivating examples from the beginning of this chapter: trade in the stock market and the trade of agricultural goods in developing countries. The network model we analyzed in this chapter is an abstraction that captures some essential features of these real markets, and misses other features. Our trade model reflects the constraint that trade takes place through intermediaries and that there is differential access to these intermediaries. Equilibria in our trading networks reflect the fact that buyers and sellers in intermediated markets, such as the stock market, face a bid-ask spread. In our model, as in actual intermediated markets, the size of this spread, and how much profit intermediaries make, depends on the amount of competition between intermediaries for the trade flow.

However, there are other interesting aspects of trade in intermediated markets that are not captured by our simple network model. In particular, we do not ask where buyers’ and sellers’ values come from; nor do we ask about how they might use information revealed by bids, asks, or trades to update these values. We discuss the role of beliefs and information in the stock market in Chapter 22.

11.8 Exercises

1. Consider a trading network with intermediaries in which there is one seller S , two buyers $B1, B2$ and two traders (intermediaries) $T1, T2$. The seller is allowed to trade with either trader. The buyers can each trade with only one of the traders: buyer $B1$ can only trade with trader $T1$; and buyer $B2$ can only trade with trader $T2$. The seller has one unit of the object and values it at 0; the buyers are not endowed with the object. Buyer $B1$ values a unit at 1 and buyer $B2$ values a unit at 2.

(a) Draw the trading network, with the traders as squares, the buyers and the seller as circles, and with edges connecting nodes who are able to trade with each other. Label each node as S , $B1$, $B2$, $T1$ or $T2$.

(b) Suppose the traders offer prices as follows.

- Trader $T1$ offers a bid price of $\frac{1}{3}$ to S and an ask price of 1 to $B1$.
- Trader $T2$ offers a bid price of $\frac{2}{3}$ to S and an ask price of 2 to $B2$.

Does this set of prices form a Nash equilibrium? If you think the answer is yes, give a brief (1-3 sentence) explanation why. If you think the answer is no, describe a way in which one of the traders could change its prices so as to increase its profit.

2. Consider a trading network in which there are two buyers ($B1$ and $B2$), two sellers ($S1$ and $S2$) and one trader ($T1$). All of the buyers and the sellers are allowed to trade with the trader. The sellers each have one unit of the object and value it at 0; the buyers are not endowed with the object, but they each want one unit; buyer $B1$ attaches a value of 1 to one unit, while buyer $B2$ attaches a value of 2 to one unit.

(a) Draw the trading network, with the trader as a square, the buyers and the sellers as circles, and edges representing pairs of people who are able to trade with each other. Label the nodes as $T1$, $B1$, $B2$, $S1$, and $S2$. Find Nash equilibrium bid and ask prices. (You do not need to provide an explanation for your answer.)

(b) Suppose now that we add a second trader ($T2$) who can trade with each seller and each buyer. In the new network is there a Nash equilibrium in which each trader's bid price to each seller is 1; each trader's ask price to buyer $B1$ is 1; each trader's ask price to buyer $B2$ is 2; one unit of the good flows from $S1$ to $B1$ through trader $T1$; and, one unit of the good flows from $S2$ to $B2$ through trader $T2$? Draw the new trading network and give a brief (1-3 sentence) explanation for your answer.

3. Consider a trading network with intermediaries in which there are two sellers $S1, S2$, three buyers $B1, B2, B3$, and two traders (intermediaries) $T1, T2$. Each seller can trade with either trader. Buyer $B1$ can only trade with trader $T1$. Buyer $B2$ can trade with either trader. Buyer $B3$ can only trade with trader $T2$. The sellers each have one unit of the object and value it at 0; the buyers are not endowed with the object. Buyer $B1$ values a unit at 1, buyer $B2$ values a unit at 2 and buyer $B3$ values a unit at 3.

(a) Draw the trading network, with the traders as squares, the buyers and the seller as circles, and with edges connecting nodes who are able to trade with each other. Label each node as $S1$, $S2$, $B1$, $B2$, $B3$, $T1$ or $T2$.

(b) Suppose prices and good flows are as follows.

- Trader $T1$ offers a bid price of 1 to each seller, an ask price of 1 to $B1$ and ask price of 2 to $B2$.
- Trader $T2$ offers a bid price of 1 to each seller, an ask price of 2 to $B2$ and an ask price of 3 to $B3$.
- One unit of the good flows from seller $S1$ to buyer $B2$ through trader $T1$ and one unit of the good flows from seller $S2$ to buyer $B3$ through trader $T2$.

(If it is useful, it is okay to write these prices and good flows on the picture you drew for part (a), provided the picture itself is still clear.) Do these prices and good flows form a Nash equilibrium? If you think the answer is yes, give a brief (1-3 sentence) explanation why. If you think the answer is no, describe a way in which one of the traders could change its prices so as to increase its profit.

4. Consider a trading network with intermediaries in which there is one buyer, one seller and two traders (intermediaries). The buyer and the seller each are allowed to trade with either trader. The seller has one unit of the object and values it at 0; the buyer is not endowed with the object but attaches a value of 1 to one unit of it.

Draw the trading network, with traders as squares, the buyer and the seller as circles, and edges representing pairs of people who are able to transact directly. Then describe what the possible Nash equilibrium outcomes are, together with an explanation for your answer.

5. Consider a trading network with intermediaries in which there is one seller S , two buyers $B1, B2$ and two traders (intermediaries) $T1, T2$. The seller is allowed to trade with either trader. The buyers can each trade with only one of the traders: buyer $B1$ can only trade with trader $T1$; and buyer $B2$ can only trade with trader $T2$. The seller has one unit of the object and values it at 0; the buyers are not endowed with the object. Buyer $B1$ values a unit at 3 and buyer $B2$ values a unit at 1.

(a) Draw the trading network, with the traders as squares, the buyers and the seller as circles, and with edges connecting nodes who are able to trade with each other. Label each node as $S, B1, B2, T1$ or $T2$.

(b) Find Nash equilibrium bid and ask prices for this trading network. How much profit do the traders make?

(c) Suppose now that we add edges representing the idea that each buyer can trade with each trader. Find a Nash equilibrium in this new trading game. What happens to trader profits? Why?

6. Consider a trading network with intermediaries in which there are three sellers $S1, S2, S3$, two buyers $B1, B2$, and two traders (intermediaries) $T1, T2$. Sellers $S1$ and $S2$ can

trade only with trader $T1$; and, seller $S3$ can trade only with trader $T2$. The buyers can each trade with only one of the traders: buyer $B1$ can only trade with trader $T1$; and buyer $B2$ can only trade with trader $T2$. The sellers each has one unit of the object and value it at 0. The buyers are not endowed with the object and they each value a unit at 1.

(a) Draw the trading network, with the traders as squares, the buyers and sellers as circles, and with edges connecting nodes who are able to trade with each other. Label each node as $S1$, $S2$, $S3$, $B1$, $B2$, $T1$ or $T2$.

(b) Describe what the possible Nash equilibria are, including both prices and the flow of goods. Give an explanation for your answer.

(c) Suppose now that we add an edge between buyer $B2$ and trader $T1$. We want to examine whether this new edge changes the outcome in the game. To do this, take the equilibrium from your answer to (b), keep the prices and good flows on the edges from (b) the same as before, and then suppose that the ask price on the new $B2$ - $T1$ edge is 1, and there is no good flow on this new edge. Do these prices and good flows still form an equilibrium? If you think that the answer is yes, give a brief (1-3 sentence) explanation why. If you think the answer is no, describe a way in which one of the participants in the game would deviate.

7. Consider a trading network in which there are two buyers ($B1$ and $B2$), two sellers ($S1$ and $S2$), and two traders ($T1$ and $T2$). The sellers each have one unit of the object and value it at 0; the buyers are not endowed with the object, but they each want one unit and attach a value of 1 to one unit. Seller $S1$ and Buyer $B1$ can trade only with trader $T1$; seller $S2$ and Buyer $B2$ can each trade with either trader.

(a) Draw the trading network, with the traders as squares, the buyers and the sellers as circles, and edges representing pairs of people who are able to trade with each other. Label the nodes as $T1$, $T2$, $B1$, $B2$, $S1$, and $S2$.

(b) Consider the following prices and flow of goods:

- $T1$'s bid price to Seller $S1$ is 0, his bid price to Seller $S2$ is $1/2$, his ask price to Buyer $B1$ is 1, and his ask price to Buyer $B2$ is $1/2$.
- $T2$'s bid price to Seller $S2$ is $1/2$ and his ask price to Buyer $B2$ is $1/2$.
- One unit of the good flows from Seller $S1$ to Buyer $B1$ through Trader $T1$; and, one unit of the good flows from Seller $S2$ to Buyer $B2$ through trader $T2$.

Do these prices and goods flow describe an equilibrium of the trading game? If you think that the answer is No, then briefly describe how someone should deviate. If you

think that the answer is Yes, then briefly explain (1-3 sentences) why the answer is Yes.

(c) Suppose now that we add a third trader (T3) who can trade with Seller S1 and Buyer B1. This trader cannot trade with the other seller or buyer, and the rest of the trading network remains unchanged. Consider the following prices and goods flow:

- The prices on the old edges are unchanged from those in part (b).
- The prices on the new edges are: a bid of $1/2$ to Seller S1 by Trader T3 and an ask of $1/2$ to Buyer B1 by Trader T3.
- The goods flow is unchanged from part (b).

Do these prices and goods flow describe an equilibrium of the trading game? If you think that the answer is No, then briefly describe how someone should deviate. If you think that the answer is Yes, then briefly explain (1-3 sentences) why the answer is Yes.