

1 Introduction

Experimental economists have been studying markets for almost fifty years (Smith, 1962, 1965). There is now a large literature on the properties of experimental markets. One of the most robust findings is that, even with relatively small numbers of traders, trade asymptotically approximates the efficient, perfectly competitive equilibrium. Most of the experimental literature assumes that markets take the form of centralized auctions, in which every trader can communicate directly with every other trader. It is true that the competitive auction market serves as the standard paradigm in economics. And there are examples of exchanges, such as the NYSE, that come close to this theoretical ideal. However, there are many other markets where intermediation and decentralized trade are the norm. Economists sometimes assume that decentralized markets behave “as if” they were centralized auction markets and, in theoretical models, it can be shown that decentralized trade does indeed lead to the perfectly competitive outcome under certain conditions. But empirical evidence on this subject is scarce.

In this paper, we exploit the methods of experimental economics to explore the properties of a simple model of decentralized trade. Empirical research can tap either real-world data from large-scale markets or small-scale laboratory data. The strengths of data from the real world are its relevance and availability. Its main weakness is that in real-world settings we observe behavior, but not preferences, technologies, or private information. In the laboratory, by contrast, we can control subjects’ preferences, technology and private information. Consequently, laboratory data are especially useful for testing the efficiency of different market institutions and for comparing market structures and institutions. The clarity that is achieved by putting a market under the microscope is well worth the effort and the necessary simplification.

We begin by extending the usual market paradigm by introducing a **network structure**. The network determines the possible patterns of trade and communication in the market. A centralized auction market corresponds to the special case of a **complete** network, in which every trader is connected to every other trader, that is, every trader can communicate and trade with every other trader. Our market, by contrast, assumes an **incomplete** network, in which many links are missing. Traders can only communicate and trade with a subset of other traders, the ones to whom they are connected by the network. The incompleteness of the network represents a potentially

serious source of friction in the market. It is far from obvious that efficient trade will be achieved when these frictions are present.

The market we study has a small number of traders, the networks are simple, and the trading mechanism remains close to the well known auction paradigm. The experimental computer platform is stable and easy to understand and provides us with an insight into the behavior of market participants. Using this design, we can see how useful the theory is in interpreting the observed behavior and study the efficiency of pricing and trade in a variety of networks.

A single treatment will serve to illustrate the experimental design. In this treatment, there are nine subjects arranged in the rectangular array illustrated in Figure 1. There are nine nodes, arranged in three rows and three columns. Each node represents a human trader and the edges between the nodes indicate trading possibilities. In addition to the human traders, there is a computer-generated seller (CGS) and a computer-generated buyer (CGB). The network architecture in Figure 1 indicates that trades are restricted to adjacent rows but, subject to these constraints, all possible trading links are present. That is, each member of the top row can trade with the CGS and with every member of the middle row; each member of the middle row can trade with every member of the top and bottom rows; and each member of the bottom row can trade with every member of the middle row and with the CGB.

[Figure 1 here]

The CGS is endowed with a single unit of an indivisible asset. The nine traders are endowed with 100 tokens each. Buyers use these tokens to pay for the asset and sellers receive these tokens in exchange for the asset. The CGB is also assumed to have an endowment of 100 tokens. The asset has no value to the CGS or to the nine traders. The CGB values the asset at 100 tokens. So the surplus (gains from trade) generated by transferring the asset from the CGS to the CGB is equal to 100 tokens. Each trader simultaneously chooses a bid (the price at which he is willing to buy the asset) and an ask (the price at which he is willing to sell the asset). The bids and asks must lie between 0 and 100 tokens. The ask of the CGS is fixed at 0 and the bid of the CGB is fixed at 100.

Once the bids and asks have been determined, trades are executed as follows. Beginning at the top of the network, the CGS and the top row

exchange the asset. The asset goes to the trader with the highest bid. If two or more traders choose the highest bid, the asset is allocated randomly between them (with equal probabilities). The top-row seller (the trader who bought the asset from the CGS) sells the asset to the middle-row trader with the highest bid that is at least as high as the seller's ask. Again, ties are broken randomly. If every bid is less than the seller's ask, no trade takes place and the game ends with the seller holding the asset. Exchange between the middle and bottom rows is executed similarly. Finally, if the asset reaches the bottom row, the asset will be transferred to the CGB because the CGB's bid of 100 is at least as great as the seller's ask. When the asset is traded, the transaction price (i.e., the price paid for the asset) is a weighted average of the bid and the ask. The corresponding amount of tokens is transferred from the buyer to the seller.

This example gives a good sense of the distinctive features of the experimental design. First, it defines a normal form game. This allows us to make precise theoretical predictions and compare them to the observed play of the game. Secondly, because the game is played in normal form (and the subjects choose strategies simultaneously), it can be played repeatedly in a relatively short amount of time, generating a large data set. Thirdly, the platform is sufficiently flexible to allow us to study a variety of network architectures, transaction pricing rules, and payoff functions. The experimental design section discusses in more detail the treatments that comprise our design.

Our results can be summarized under three headings:

- *Convergence*. Since the underlying trading game is essentially a Bertrand pricing game, in any Nash equilibrium trade is efficient and the transaction prices are equal to 100 and trading profits are zero, except in the top row, where the CGS is restricted to ask 0 and profits are 50. In all treatments, the observed transaction prices start low and rise monotonically towards the equilibrium price. The speed of convergence varies across treatments. In some treatments, we do not observe complete convergence when the experiment ends after 30 trading periods, but overall the predictive power of the theory is impressive.
- *Efficiency*. Strategic uncertainty (about what other subjects will do) inevitably requires a period of learning and during this period trades may not be completed. Further, even later in the game, trade may break down if subjects make mistakes about the prices that are likely

to be bid or asked by their opponents. On the whole, trade is asymptotically efficient in the sense that it tends to be lower in the early trading periods and rises as subjects become more confident about the behavior of other subjects and as the prices bid and asked converge to the competitive equilibrium prices. Given the incompleteness of the networks, which requires intermediate trades, and the strategic form of the game, which does not allow for recontracting, the subjects' ability to coordinate on an efficient outcome is quite striking.

- Sensitivity. We study a number of variants to test the sensitivity of the results to the amount of competition (the number of columns in the network), the pricing rule (the weights used in calculating the transaction price), and the payoff function (the amount of trading losses deducted from subjects' earnings). Although convergence and efficiency are somewhat sensitive to each of these changes, the equilibrium properties continue to have predictive power. Among other things, we note that (i) less competition may lead to slower convergence and lower efficiency, and (ii) trading losses and the bid-pricing rule reduce competition by making bidders less aggressive and thus lower efficiency.

The rest of the paper is organized as follows. A discussion of the related literature is provided in Section 2. We describe the theoretical model and the experimental design in Section 3. The results are contained in Section 4. Some concluding remarks and important topics for further research are contained in Section 5.

2 Related literature

Kranton and Minehart (2001) introduce the first model of exchange in networks. Gale and Kariv (2007) develop a model of financial networks and use this model to investigate the role of costly intermediation ("frictions") and network architecture ("incompleteness") in determining the efficiency of markets and the possibility of market breakdown. Gale and Kariv (2007) show that, in the limit as the period length goes to zero and the market becomes frictionless, the market outcome is efficient. Bosch-Domènech and Sunder (2000) study an economy consisting of multiple interconnected markets. The trading mechanism is the double auction. With the help of numerical simulations, they show that, under certain conditions, equilibrium prices can be

induced by zero-intelligence traders making random bids and asks, though they might fail to achieve an efficient final allocation. In this paper, we test the efficiency of pricing and trade using a variety of network architectures. On the whole, the level of efficiency is very high, although it is sensitive to the specification of the transaction price.

The present paper contributes to the enormous body of work on experimental markets. Following the seminal papers of Forsythe, Palfrey and Plott (1982, 1984), and Plott and Sunder (1982, 1988), numerous experimental papers analyze many aspects of asset markets.¹ In these experiments, a double auction or bid-ask market is typically used. The main conclusion from this large body of experiments is that the double auction market produces efficient allocations and prices, even with a very small number of traders, thus providing experimental evidence that markets are an efficient mechanism for allocating resources.

In contrast to the existing literature, our paper contributes to the systematic experimental study of the efficiency of trade in networks. Although network experiments in economics are recent, there is now a substantial experimental literature on the economics of networks.² To the best of our knowledge, previous contributions have been quite different from ours. The most closely related paper is by Charness, Corominas-Bosch and Fréchette (2007), who investigate how the network structure affects the outcomes and dynamics of ultimatum bargaining. Following the model of Corominas-Bosch (2004), they decompose a network of buyers and sellers into two simple sub-graphs and test whether it matters how a single edge is added between these two groups of traders.

3 Theory, predictions and design

In this section, we describe the theory on which the experimental design is based and the design itself.

¹See Sunder (1995) for a comprehensive, if now somewhat dated, discussion of the experimental work on asset markets.

²Kosfeld (2004) surveys the experimental work in economics.

3.1 The trading game

The trading game consists of a finite number of players, indexed by $k = 1, \dots, K$, arranged in a rectangular network consisting of m rows and n columns. An example of a 3×3 network is illustrated in Figure 1 above. Holding the trading protocol, the pricing rule and the payoff function constant, adding rows increases the amount of intermediation required to capture the surplus available, whereas adding columns (i.e., adding players in each row) increases the amount of competition. A single player is located at each node and the edges connecting the nodes indicate that the corresponding players can trade with each other. In addition to the human players, there are two computer-generated players, called the computer-generated seller (CGS) and the computer-generated buyer (CGB). The CGS has one unit of an indivisible asset which he is willing to sell for zero tokens and the CGB is willing to buy the asset for $v > 0$ tokens.

The networks we consider are symmetric and satisfy the following properties. Only the players in the first row $i = 1$ can purchase the asset from the CGS. The players in row $i > 1$ can buy the asset from any of the players in row $i - 1$ and the players in row $i < m$ can sell the asset to any of the players in row $i + 1$. Only players in the last row $i = m$ can sell the asset to the CGB. A strategy for each player k consists of the announcement of a bid price (b_k) at which he would be willing to purchase one unit of the asset and an asking price (a_k) at which he would be willing to sell one unit of the asset. The bid and ask prices are restricted to the interval $[0, v]$ so the strategy set for player k is simply $S_k = [0, v] \times [0, v]$ and the set of strategy profiles is $S = S_1 \times \dots \times S_K$.

Trades are executed as follows. The asset is transferred from the CGS to the row 1 player who has the highest bid. If there is more than one player with the highest bid, the asset is allocated randomly among the winning bidders. The player who receives the asset pays the CGS an amount equal to αb_k where b_k is the winning bid and $0 \leq \alpha \leq 1$ is a constant. In each row $1 < i < m$, trade is only possible if the asset is held by one of the players in row $i - 1$, whom we call the seller. Trade takes place if at least one bid in row i is greater than or equal to the seller's asking price. The asset is transferred to the highest bidder. If more than one player has the highest bid, the asset is allocated randomly among the winning bidders. The player who receives the asset transfers an amount equal to $\alpha b_k + (1 - \alpha) a_{k'}$, where b_k is the winning bid and $a_{k'}$ is the seller's asking price. If no bid is at least

as high as the seller's asking price, no trade takes place and the asset remains with the seller. If a player in row m receives the asset, he sells it to the CGB for a price equal to $\alpha v + (1 - \alpha) a_k$, where a_k is the seller's asking price.

A player's payoff is equal to his trading profit, that is, the amount he receives from selling the asset minus the amount he pays for it. A player who does not manage to buy the asset receives a payoff of zero. A player who buys the asset for a positive price and fails to sell it receives a negative payoff. Denote a typical player's strategy by $\sigma_k = (a_k, b_k)$ and a strategy profile by $\sigma = \{\sigma_k\} = (a, b)$, where $a = \{a_k\}$ and $b = \{b_k\}$. Denote player k 's payoff by $\pi_k(\sigma) = \pi_k(a, b)$. A Nash equilibrium is a strategy profile σ^* such that for any player k ,

$$\pi_k(\sigma^*) \geq \pi_k(\sigma_k, \sigma_{-k}^*)$$

for any $\sigma_k \in S_k$. It is not hard to see that the usual Bertrand competition result holds.

Theorem 1 (Nash equilibrium) Suppose that there are at least two players in each row ($n \geq 2$). Then, in any Nash equilibrium, the asset passes, by means of a sequence of trades, from the CGS to the CGB – in other words, the market is *efficient* – and in each equilibrium trade the transaction price is equal to v , except for the *first* row, where the transaction price is αv .

The normal-form game described above has an important advantage from the point of view of laboratory experiments. Since all strategies are chosen simultaneously and the outcome is calculated by the computer, the normal-form game can be played much more quickly than corresponding extensive-form games. This allows us to gather a large amount of data in a reasonable period of time. There is a close relationship between the normal form and the "natural" extensive form game. Suppose that trades occur sequentially, with perfect information at each stage. First, the CGS and the first row buyers announce their bids and trade is executed in the usual way. Then the first row seller (assuming trade has been completed) and the second row buyers announce their ask and bids, respectively, and trade is executed in the usual way. This procedure is repeated until trade fails to take place or the asset reaches the CGB. The difference between this extensive-form game and our normal-form game is that, in the former, traders know the history of trade and condition their strategies on it, whereas, in the latter, strategies are necessarily history independent. In fact, if we restrict attention to subgame

perfect equilibria in which traders use history independent strategies, the equilibrium outcome of the extensive-form game will be the same as the outcome of some Nash equilibrium of the normal-form game.

3.2 Experimental design

Our experimental design employs three network structures (3×3 , 3×2 , and 2×3), two pricing rules ($\alpha = 0.5$, and $\alpha = 1$) and two payoff functions (defined in equations (1) and (2) below). Each treatment consists of a network structure, a pricing rule and a payoff function. We study three combinations of payoff functions and pricing rules and these are applied to each of the three network structures. The groups of three treatments based on a single payoff function and pricing rule and the three different networks are referred to as the baseline, bid-price and loss treatments described below. The ask of the CGS is always fixed at zero and the bid of the CGB is fixed at $v = 100$. The human traders are endowed with 100 tokens each. Figure 2 summarizes the experimental design.

[Figure 2 here]

The three **baseline treatments** use the mean-price rule ($\alpha = 0.5$) and the payoff function

$$\text{payoff} = 10 + \max\{0, \text{trading profits}\}, \quad (1)$$

where trading profits (positive or negative) are defined as the difference between the revenue from selling the asset (zero if the asset was not sold) and the cost of buying the asset (zero if the asset was not purchased).

The payoff function (1) can be interpreted as a model of professional traders who receive bonuses when trading profits are positive, but do not suffer losses when trading profits are negative. As a matter of experimental design, the payoff function represents a compromise between two conflicting objectives. In experimental markets, subjects whose earnings at the end of a trading period are negative are usually considered “bankrupt” and barred from trading in future periods. Unfortunately, bankruptcy is impractical in a trading network where each subject occupies a different node: eliminating bankrupt traders would break up the network and disrupt trade. The risk of bankruptcy can be avoided if we treat the entire endowment of 100 tokens as the subject’s property in each period. Then, since a subject cannot lose

more than 100 tokens in a single period, he can never be bankrupt. But this creates another problem: now a subject can earn 100 tokens by not trading. What incentive does he have to risk losses by trading? The payoff function (1) is a compromise that avoids bankruptcy and still provides an incentive to trade, by subtracting the endowment from the trader's earnings and also limiting the losses.

Within the baseline treatments, we are testing the sensitivity of the subjects' behavior with respect to variations in the network structure. The remaining treatments are designed to test their sensitivity to changes in the pricing rule and the payoff function. The three **bid-price treatments** test the robustness of the results of the baseline treatments to a change in the definition of the transactions price by setting it equal to the winning bid ($\alpha = 1$). The **loss treatments** test the robustness of the results of the baseline treatments to a change in the definition of the payoff function by substituting the payoff function

$$\text{payoff} = 50 + \max\{-40, \text{trading profit}\}. \quad (2)$$

for the definition in equation (1).

Compared to the function defined by (1), the payoff function in (2) has a higher constant term and a higher limit on the trading losses that can be deducted, as depicted in Figure 3 below. If a subject makes a trading loss of more than 40 tokens, his payoff will be equal to 10 tokens. If the trading profit is non-negative, his payoff will be at least 50 tokens, that is, 40 tokens more than under payoff function (1). In effect, we are increasing the amount of the endowment the trader can keep to 50 tokens while increasing the maximum loss the trader can bear to 40 tokens. The fact that the subject can now earn 50 tokens for sure by not trading creates a significant risk of loss from trading. Under payoff function (2), payoffs are higher, other things being equal, but the incentive to trade may be smaller.

[Figure 3 here]

3.3 Experimental procedures

All the experimental sessions were conducted at the Center for Experimental Social Science (C.E.S.S.) at New York University (NYU). The subjects were recruited from the undergraduate student body of the College of Arts and

Sciences at NYU. Subjects read the instructions silently, after which the instructions were read aloud by one of the experiment administrators. Subjects were invited to ask questions during the verbal instruction period. No subject reported difficulty understanding the procedures or using the computer interface. Each experimental session lasted a little more than one hour. A \$5 participation fee and subsequent earnings, which averaged about \$20, were paid in private at the end of the session.

Each experimental session consisted of 30 independent trading periods. The experimental treatment was held constant throughout a given experimental session. At the beginning of each session, each subject was randomly and independently assigned to one row of the network. This determined his type, which remained constant throughout the experiment. At the beginning of each trading period, the computer would randomly form networks by assigning subjects to the various nodes in the network. Top row subjects were assigned to top row nodes, middle row types to middle row nodes, and so on, but the assignments were otherwise random and subjects had an equal probability of being selected for each node and network. Subjects were informed of the network structure, the trading protocol, the pricing rule, and the payoff function.

At the beginning of a trading period, subjects would be informed of the position in the network to which they were assigned and then would be asked to choose a bid price (the price at which they were willing to buy one unit of an asset) and an ask price (the price at which they were willing to sell one unit of the asset). Each subject had an initial endowment of 100 tokens and was allowed to choose any number (including decimals) between 0 and 100 as a bid or ask price. Subjects knew the asking price of the CGS (0 tokens) and the bid price of the CGB (100 tokens).

The computer program dialog window is shown in the sample experimental instructions which are reproduced in Online Appendix I. The main features of the computer interface are: the large window at the left of the screen, which displays the network and the price and trading information; the View Results button in the top right corner of the screen, which allows subjects to recall the price and trading information from any previous trading period; the Bid and Ask fields at the right of the screen, where subjects enter the prices at which they are willing to buy and sell; and the message window in the lower left corner of the screen.

In each period, subjects are required to enter their bids and asks in the respective fields and click the Submit button. After all subjects have entered

a bid and ask price, the computer executes the feasible trades according to the trading protocol. Subjects are then shown the resulting transactions and bid and ask prices in their network. When they are ready to begin the next trading period, subjects click the OK button. When all subjects have clicked the OK button, the next trading period begins. After this process has been repeated 30 times, the experiment ends, the computer selects one trading period at random, where each period has an equal probability of being chosen, and the subject is paid an amount based on the number of tokens earned in that period. Payoffs are calculated in terms of tokens and then converted into dollars. Each token is worth \$1.

The experiments provide us with a rich set of data. The diagram below summarizes the experimental dataset. The entries of the form $a/b/c$ represent the number of subjects (a), the number of observations on different networks (b), and total of individual decisions (c).

Treatment		# of obs.
Baseline	3×3	54/180/1620
	3×2	36/180/1080
	2×3	54/270/1620
Bid-price	3×3	54/180/1620
	3×2	36/180/1080
	2×3	36/180/1080
Loss	3×3	54/180/1620
	3×2	30/150/900
	2×3	30/150/900

4 Experimental results

In this section, we present our experimental results concerning the convergence of prices to equilibrium values, the efficiency of trade, and the sensitivity of the observed behavior to variations in networks, payoff functions, and pricing rules. The aim of the analysis is to provide insights into how experimental networks behave, as well as to test the usefulness of the theory for interpreting behavior in the laboratory.

The novel feature of our design is the presence of intermediation. We first explore the effect of intermediation, measured by the number of rows in the network, and competition, measured by the number of columns in the

network, on bid and ask prices. Then we compare the levels of efficiency, measured by the fraction of completed trades, across treatments.

4.1 The data

We begin by providing an overview of some important features of the experimental data. Each experiment consists of 30 trading periods. To facilitate comparisons across treatments, instead of showing the data from each trading period, we have grouped the trading periods into terciles, corresponding to early periods (1-10), intermediate periods (11-20) and late periods (21-30). To economize on space, Online Appendix II presents the data for each tercile as a separate sub-panel. Each cell contains the mean and standard deviation over the tercile, the treatment, and the row of the $m \times n$ network used in the treatment. Appendix II also presents the data, trading period by trading period, in graphical form.³

Panel A shows the average winning bids and Panel B shows the average maximum bids. If the winning bid is less than the ask, no trade occurs and the winning bid is not defined. In that case, the average winning bid will be different from the average maximum bid because of the missing observation. Panel C shows the average seller's asks. The seller's ask is the asking price of the winning bidder. Thus, we are calculating means of bid and asks for the same population of subjects in Tables 1A and 1C.⁴

Panel D shows the average transaction price. The transaction price corresponding to row i is the actual amount paid for the asset by the subject in row i . If no trade occurs, the transaction price is not defined and is not included in the average. Recall that the pricing rule in the baseline and loss treatments is $\alpha = 0.5$ and that the CGS always asks 0. These two facts imply that the transaction price in the first row is never more than 50. Finally, Panel E displays, row by row, the fraction of completed trades. The first row is excluded because there is no possibility of incomplete trade in that row.

³Since we set the transaction price equal to the bid price ($\alpha = 1$) in the bid-price treatments, we only report the evolution of the average winning bids and the fraction of completed trades.

⁴We include the average maximum bids, as well as the average winning bids, to give a better picture of the evolution of bidding behavior. Some of the volatility in the average winning bid series is caused by dropping observations when there is no trade. When it comes to asking prices, only the seller's ask is relevant. We record it here whether or not the seller is successful in making a sale.

From the data in Appendix II and our observations in the laboratory, we draw four broad conclusions:

- First, the experimental platform is stable and easy to understand. It generates a large amount of data in a short amount of time, allowing us to test the equilibrium predictions of the theory.
- Secondly, although strategic uncertainty (about what other subjects will do) necessitates a period of learning, in most treatments, subjects rapidly converge to equilibrium prices and trade is generally efficient.
- Thirdly, the time path of prices is qualitatively similar across treatments: bid and ask prices start low and rise monotonically toward their equilibrium values as subjects become more confident about the behavior of other subjects.
- Finally, although the results are qualitatively similar, the speed of convergence does vary among the different treatments. The observed differences are quite intuitive: convergence is slower, other things being equal, if there is more intermediation or less competition, if the transaction price equals the winning bid (rather than the average of the winning bid and ask), or if there is a higher limit on losses from trading that can be deducted.

4.2 Convergence

Since in every network there are two or three bidders in each row, the theoretical model suggests that any equilibrium of the trading game is efficient and the transaction prices equal 100 tokens (except for the first row in each network in the baseline and loss treatments, where the transaction price equals 50 tokens because the seller's ask is fixed at 0 tokens).

In an experimental setting, there are many reasons why we do not at first observe the equilibrium transaction prices. Perhaps the most important reason is strategic uncertainty: a subject bidding for an asset has little information about the price at which he can re-sell the asset, unless he happens to be in the bottom row and can sell the asset to the CGB for the price of 100. Uncertainty about the resale price may cause subjects to shave their bids in order to protect themselves against the possibility of selling at a loss, or failing to sell at all.

Although subjects are randomly matched with different subjects each period, they learn from experience and their uncertainty gradually diminishes as the trading game is repeated. As a result, the effectiveness of competition may be expected to increase and cause the transaction price to converge to the equilibrium price. Result 1 summarizes the time paths of the average winning bid, ask and transaction prices across treatments.

Result 1 (convergence) In the baseline treatments, the average winning bid and ask prices are initially far below their equilibrium values, but they converge rapidly after several periods. Transaction prices, being the average of winning bids and asks, also converge rapidly to the equilibrium values.

There are small differences between the rates of convergence in the bid-price and loss treatments based on the 3×3 and 2×3 networks and the corresponding baseline treatments, especially in the second and third terciles. The rates of convergence in the bid-price and loss treatments based on the 3×2 network are much slower than in the corresponding baseline treatment and, in most cases, prices fail to converge completely by the end of 30 trading periods.

The relevant support for Result 1 comes from Appendix II. In the baseline treatments, the average winning bids converge very fast, especially in the 3×3 and 2×3 networks, where all bids are within one percent of the equilibrium bid of 100 in the second and third terciles. The average seller's asks converge more slowly, but almost all reach the neighborhood of equilibrium by the third tercile. The maximum bids should show similar time paths (in practice, the time paths are almost identical once we aggregate by tercile). Overall, given subjects' uncertainty about the possibility of reselling the asset, it takes remarkably little time for prices to converge.

The bid-price treatments differ from the baseline treatments only in setting the transaction price equal to the successful bid ($\alpha = 1$). Intuitively, this could slow convergence by making subjects less willing to bid aggressively. Nevertheless, the rates of convergence are quite similar under bid-price and average-price rules, with the exception of the bottom row of the 2×3 network, where the average winning bids are lower in the bid-price treatment than in the baseline treatment.

The loss treatments use a payoff function that deducts more trading losses from subjects' earnings than in the other treatments. Consequently, one

might expect that price convergence will be slower than in the baseline treatments – or that convergence will be incomplete – because the possibility of substantial trading losses makes subjects less willing to bid aggressively for the asset. Convergence is slower in the loss treatments based on the 3×3 and 2×3 networks than in the corresponding baseline treatments, but prices very rapidly reach the neighborhood of the equilibrium price. In the 3×2 loss treatment, by contrast, the prices are lower and the gap between the 3×2 loss treatment and the 3×2 baseline treatment often widens over time. Furthermore, in the 3×2 loss treatment, convergence to the equilibrium price is incomplete at the end of the experiment and it is not clear whether further repetitions would lead to complete convergence. We can thus conclude that full price convergence occurs in most but not all treatments.

Next, we look more closely at the sensitivity of pricing behavior to the network architecture by comparing behavior in corresponding rows across networks and in different rows within a given network, holding other aspects of the treatment constant. We first examine the pricing behavior of subjects belonging to corresponding rows across networks.

4.3 Intermediation

The 3×3 and 2×3 networks differ only in the number of rows. The more rows, the greater the amount of intermediation required to transfer the asset from the CGS to the CGB. One might expect that more intermediation would reduce the speed of convergence to the equilibrium price, but that does not appear to be the case. In fact, in all treatments, the pricing behavior is quite robust to variation in the number of rows in the network. Result 2 summarizes the behavioral regularities in this regard by comparing average winning bids and sellers' asks in rows that have similar positions relative to the CGB.

Result 2 (intermediation) In all treatments, there are small differences in rates of convergence of the average winning bids and sellers' asks to their equilibrium values in the corresponding rows of the 3×3 and 2×3 networks, that is, the bottom rows of the 3×3 and 2×3 networks, and the top and middle rows of the 3×3 and 2×3 networks, respectively.

The support for Result 2 comes again from Appendix II. We reorganize the relevant information in Table 1 below. For the bid prices, the only case

where differences in the rates of convergences are substantial is during the first tercile.⁵ For the ask prices, there is also some variation in the first tercile, but in later terciles the corresponding rows have similar prices, except for the curious drop in asking prices in the middle row of the 3×3 baseline treatment in the last tercile. Apart from this small difference, the asks in corresponding rows do not differ across networks after the first tercile. Note that we do not include the asks in the bid-price treatments (middle panel), since transaction prices do not depend on the seller's ask. We also do not include the asks in the bottom rows in each network in the other treatments, since subjects very quickly realized that they could ask for 100 tokens from the CGB.

[Table 1 here]

4.4 Competition

In the 3×3 and 2×3 networks, there are three bidders ($n = 3$) in each row. In theory, Bertrand competition will guarantee an equilibrium price of 100 as long as there are at least two bidders in each row. In the laboratory, we do not necessarily expect perfectly competitive behavior when the number of bidders is small. It is therefore of particular interest to see whether prices reach the neighborhood of the competitive price when we reduce the number of subjects in each row. We compare, row by row, the pricing behavior in the 3×3 networks with the pricing behavior in the corresponding 3×2 networks. Considerable differences in pricing behavior are observed, as the next result reports.

Result 3 (competition) With few exceptions, competition increases the bid and ask prices in any given row. As a result, other things being equal, the rates of convergence to the equilibrium price are slower in the 3×2 network than in the corresponding 3×3 network. In the 3×2 bid-price and loss treatments, prices generally fail to converge to their equilibrium values by the end of 30 trading periods.

⁵As a benchmark, we also studied an auction treatment (a single session with 15 subjects), which is identical to the baseline treatment except that it only uses the 1×3 network so the successful buyer knows that he can always sell the asset to the CGB for the price of 100. The absence of uncertainty in the auction treatment guarantees aggressive bidding in line with the predictions of equilibrium, that is, convergence to the equilibrium price occurs in the first few trading periods and the bids remain at that level throughout the game, apart from occasional experimental deviations.

Support for Result 3 is also based on the data from Appendix II. In Table 2 below, we compare, row by row, the relevant data from Appendix II on the average winning bids and sellers' asks. We again do not include the asks in the bid-price treatments and in the bottom rows in each network in the baseline and loss treatments. Bid and ask prices show the same time profile when competition is high ($n = 3$) as when it is low ($n = 2$), though the levels are different and the general pattern is that competition increases the price in any given row, with the exception of the curious drop in the seller's asks in the middle row of the 3×3 baseline treatment. We thus conclude that, overall, less competition does make a difference, especially in the bid-price and loss treatments where behavior do not reach the neighborhood of equilibrium at the end of the experiment when competition is low.

[Table 2 here]

4.5 Spreads

Another interesting feature of the pricing behavior is the spread between adjacent rows of a given network, other things being equal. Again, the critical factor is the uncertainty about resale as measured by the distance from the CGB. Except for the bottom-row buyer, who can resell the asset to the CGB for the price of 100, subjects cannot be sure of the price at which they can resell the asset. This strategic uncertainty tends to depress the bids and asks, and suggests that, in the 3×3 and 3×2 networks, the transaction prices between the first row seller and the second row buyer will be lower than the transaction prices between the second row seller and the third row buyer. Our next result confirms this conjecture. The evidence is again provided in Appendix II. We present the relevant data from Appendix II in Table 3, which compares the average transaction prices across rows within a given network and tercile. Since the ask of the CGS is fixed at 0 and the bid of the CGB is fixed at 100, we restrict attention to comparing transaction prices between the second and third rows in the 3×3 and 3×2 networks.

Result 4 (spreads) The average transaction prices in a given network are increasing in the row index for each network and in each tercile, with the exception of the 3×3 baseline treatment, where in the third tercile the average transaction price is considerably lower in the bottom row than in the middle row.

[Table 3 here]

4.6 Efficiency

Efficiency is one of the main concerns in the study of trading in networks. Trade is efficient if and only if the asset reaches the CGB, so that the surplus is realized. Transaction prices affect the distribution of the surplus but have no impact on efficiency. In this section we assess the efficiency of trade overall and its sensitivity to the various treatments.

Incompleteness of networks is a potentially important market friction. The greater the incompleteness of the network, the more intermediation is required to capture gains from trade and achieve an efficient outcome. Recall that trade between two rows requires that at least one bid is higher than or equal to the seller's ask. Thus, strategic uncertainty (about what other subjects will do) inevitably requires a period of learning and during this period trades may not be completed. Furthermore, even after a considerable amount of trade has occurred, trade may break down if subjects are too aggressive and misjudge the prices that are likely to be bid or asked by their opponents.

The completion rates, that is, the fraction of trades completed within a given network and tercile, are reported in Table 4. As usual, the data are taken from Appendix II above. Recall that the first row is excluded because there is no possibility of incomplete trades in that row.

[Table 4 here]

Other things being equal, the level of efficiency is generally highest in the 2×3 treatments (where the degree of intermediation is lower than in the 3×3 network) and lowest in the 3×2 treatments (where the degree of competition is lower than in the 3×3 network). Thus, we conclude that more intermediation or less competition can lower the efficiency of trade. In the 3×3 baseline treatment, the level of efficiency increases markedly through the three terciles. It increases through the first two terciles in the 3×2 baseline treatment, but is quite high and essentially flat in the 2×3 baseline treatment. This suggests that it takes subjects longer to learn to coordinate when there is more intermediation.

Comparing corresponding networks in the baseline and bid-price treatments shows that changing the transaction pricing rule from the average of

the winning bid and ask to the winning bid price reduces efficiency in all networks. In the 3×2 bid-price treatment, efficiency increases sharply from the first to the second tercile and increases modestly between the second and third terciles, and in the 2×3 bid-price treatment, efficiency is again quite high and essentially flat. The most interesting feature of the bid-price treatments is that efficiency declines sharply from the first to the third tercile in the 3×3 network, but it is the aberrant behavior of a few individuals in a single session that accounts for most of this drop in efficiency. The effect of even a few anomalous subjects can propagate through the network as a change in prices in one part of the network affects what traders are prepared to bid and ask elsewhere.

In the loss treatments, efficiency is overall lower than in the baseline treatments, as one would expect. Also note that in the 3×2 loss treatment, efficiency increases and then decreases, whereas in the 3×2 baseline treatment, efficiency is steadily increasing over time. This discussion is summarized in our last result.

Result 5 (efficiency) The levels of *efficiency* appear to be lower when there is more intermediation or less competition. Further, the trading rules are important for *efficiency*: when the transaction price is equal to the bid price or subjects experience trading losses, the level of *efficiency* appears to be lower.

5 Concluding remarks

In this paper we have examined the robustness of market behavior to changes in network architecture, payoff functions, and pricing rules. We restricted our attention to rectangular arrays with symmetric structures. The advantage of this network architecture is that all nodes in a row are essentially the same, allowing us to randomize the assignment of subjects across nodes and networks and to pool the data that we collected. Nevertheless, there are many other network architectures that would be interesting to study, particularly asymmetric networks.

Among the many phenomena that could be studied using this platform, the impact of uncertainty on trade is one of the most interesting. Here we mention three possibilities.

- Random endowments. A trader's endowment places an upper limit on

what he can bid for an asset and serves as a “liquidity constraint.” Random endowments introduce liquidity shocks that will change pricing both directly, by constraining bids, and indirectly by reducing competition for bidders and lowering resale prices for intermediaries.

- Random graphs. Random graphs are intrinsically interesting because they introduce uncertainty about the availability of counterparties to trade with. They also give rise to interesting strategic phenomena. For example, if the number of bidders in an auction is uncertain and with positive probability the number of bidders is one, the only equilibrium involves mixed strategies. Further, the effects of randomness can propagate through the network.
- Random values. Uncertainty about the values assigned to the asset by the CGS and the CGB introduces uncertainty about the probability of trade and the possibility of learning the value of the asset over time. Our framework can provide insight into how these important phenomena will be affected by network architectures.

While the small networks we studied are insightful, especially in experimental contexts, the development of the theory depends on properties of networks that can be generalized. In order to determine which factors are important in explaining market behavior, it will be necessary to investigate a large class of networks in the laboratory. Fortunately, our experimental design enables us to do this systematically and efficiently.

References

- [1] Charness, G., M. Corominas-Bosch and G. Fréchette (2007) “Bargaining and Network Structure: An Experiment.” *Journal of Economic Theory* **136**, 28-65.
- [2] Bosch-Domènech, A. and S. Sunder (2000). “Tracking the Invisible Hand: Convergence of Double Auctions to Competitive Equilibrium.” *Computational Economics* **16**, 257-284.
- [3] Corominas-Bosch, M. (2004) “Bargaining in a Network of Buyers and Sellers.” *Journal of Economic Theory* **115**, 35-77.

- [4] Forsythe, R., T. Palfrey and C. Plott (1982) "Asset Valuation in an Experimental Market." *Econometrica* **50**, 537-567.
- [5] Forsythe, R., T. Palfrey and C. Plott (1984) "Futures Markets and Informational Efficiency: A Laboratory Examination." *Journal of Finance* **39**, 955-981.
- [6] Gale, D. and S. Kariv (2007) "Financial Networks." *American Economic Review, Papers & Proceedings* **97**, 99-103.
- [7] Kosfeld, M. (2004) "Economic Networks in the Laboratory: A Survey." *Review of Network Economics* **3**, 20-41.
- [8] Kranton, R. and D. Minehart (2001) "A Theory of Buyer-Seller Networks." *American Economic Review* **91**, 485-508.
- [9] Plott, C. and S. Sunder (1982) "Efficiency of Experimental Security Markets with Insider Information: An Application of Rational-Expectations Models." *Journal of Political Economy* **90**, 663-698.
- [10] Plott, C. and S. Sunder (1988) "Rational Expectations and the Aggregation of Diverse Information in Laboratory Security Markets." *Econometrica* **56**, 1085-1118.
- [11] Smith, V. (1962) "An Experimental Study of Competitive Market Behavior." *Journal of Political Economy* **70**, 111-137.
- [12] Smith, V. (1965) "Experimental Auction Markets and the Walrasian Hypothesis." *Journal of Political Economy* **73**, 387-393.
- [13] Sunder, S. (1995) "Experimental Asset Markets: A Survey." in *Handbook of Experimental Economics*, ed. A. Roth and J. Kagel, Princeton University Press.

Table 1: The effect of intermediation
(Average winning bids and seller's asks)

A: Baseline

	Treatment/ row	Trading periods		
		1-10	11-20	21-30
Bids	3×3/3-2×3/2	97.1 - 90.0	99.9 - 99.9	100.0 - 100.0
	3×3/2-2×3/1	88.2 - 82.5	99.9 - 100.0	100.0 - 100.0
Asks	3×3/2-2×3/1	82.8 - 70.7	91.6 - 87.8	81.0 - 94.3

B: Bid-price

	Treatment/ row	Trading periods		
		1-10	11-20	21-30
Bids	3×3/3-2×3/2	66.9 - 95.0	94.9 - 98.7	99.4 - 99.4
	3×3/2-2×3/1	63.6 - 90.0	92.8 - 98.4	98.4 - 99.4

C: Loss

	Treatment/ row	Trading periods		
		1-10	11-20	21-30
Bids	3×3/3-2×3/2	96.4 - 88.2	98.1 - 99.0	99.6 - 100.0
	3×3/2-2×3/1	79.8 - 65.3	97.1 - 90.2	98.2 - 99.4
Asks	3×3/2-2×3/1	64.8 - 74.0	95.0 - 93.7	96.1 - 96.9

Table 2: The effect of competition
(Average winning bids and seller's asks)

A: Baseline

	Treatment/ row	Trading periods		
		1-10	11-20	21-30
Bids	3×3/1-3×2/1	77.6 - 36.6	99.9 - 65.0	100.0 - 91.3
	3×3/2-2×3/2	88.2 - 58.6	99.9 - 78.5	100.0 - 91.4
	3×3/3-3×2/3	97.1 - 84.4	99.9 - 96.6	100.0 - 98.8
Asks	3×3/2-3×2/2	73.9 - 50.3	85.9 - 71.5	95.4 - 87.4
	3×3/2-2×3/2	82.8 - 70.9	91.6 - 89.2	81.0 - 96.4

B: Bid-price

	Treatment/ row	Trading periods		
		1-10	11-20	21-30
Bids	3×3/1-3×2/1	51.0 - 39.9	84.8 - 58.9	95.4 - 70.2
	3×3/2-2×3/2	63.6 - 65.0	92.8 - 77.4	98.4 - 85.1
	3×3/3-3×2/3	66.9 - 79.1	94.9 - 86.6	99.4 - 91.4

C: Loss

	Treatment/ row	Trading periods		
		1-10	11-20	21-30
Bids	3×3/1-3×2/1	65.0 - 28.6	92.1 - 35.6	94.0 - 49.8
	3×3/2-2×3/2	79.8 - 53.3	97.1 - 62.2	98.2 - 69.8
	3×3/3-3×2/3	96.4 - 70.0	98.1 - 73.9	99.6 - 80.0
Asks	3×3/2-3×2/2	53.2 - 43.7	83.0 - 53.1	87.4 - 60.8
	3×3/2-2×3/2	79.4 - 64.8	95.0 - 69.0	96.1 - 75.7

Table 3: Price spreads
(Average transaction prices)

A: Baseline

Treatment/ row	Trading periods		
	1-10	11-20	21-30
3×3/2-3×3/3	81.0 - 88.6	92.9 - 95.3	97.7 - 90.2
3×2/2-3×2/3	54.4 - 77.6	75.0 - 92.7	89.4 - 97.5

B: Bid-price

Treatment/ row	Trading periods		
	1-10	11-20	21-30
3×3/2-3×3/3	63.1 - 66.7	91.8 - 93.5	97.3 - 96.7
3×2/2-3×2/3	65.0 - 79.1	77.4 - 86.6	85.1 - 91.4

C: Loss

Treatment/ row	Trading periods		
	1-10	11-20	21-30
3×3/2-3×3/3	66.5 - 87.6	90.1 - 96.4	92.8 - 97.9
3×2/2-3×2/3	48.5 - 66.5	57.6 - 71.0	65.3 - 76.5

Table 4: The efficiency of trade
(by treatment and tercile)

Treatment		Trading periods		
		1-10	11-20	21-30
Baseline	3×3	0.72	0.88	0.97
	3×2	0.67	0.87	0.88
	2×3	0.93	0.99	0.96
Bid-price	3×3	0.67	0.58	0.48
	3×2	0.55	0.75	0.78
	2×3	0.90	0.87	0.85
Loss	3×3	0.72	0.82	0.87
	3×2	0.64	0.78	0.70
	2×3	0.92	0.98	0.94

Figure 1: The 3×3 network




Figure 2: The experimental design




Figure 3: The payoff functions

