

Price Formation in Double Auctions

by

Steven Gjerstad and John Dickhaut

Discussion Paper No. 284, November 1995

Center for Economic Research
Department of Economics
University of Minnesota
Minneapolis, MN 55455

Price Formation in Double Auctions

by

Steven Gjerstad* and John Dickhaut**

*Department of Economics
University of Arizona
Tucson, AZ 85721

**Department of Accounting
University of Minnesota
Minneapolis, MN 55455

Abstract

Markets have the capacity to resolve complex coordination problems. Hayek [1945] asked how privately held market information is organized through the trading process to arrive at competitive equilibrium. We propose strategies for sellers and buyers in a double auction (DA) market that result in transaction prices at or near competitive equilibrium in a variety of market environments.

A large experimental literature documents convergence in many market environments to competitive equilibrium, but the theoretical literature treating the bargaining behavior in this institution is relatively small, and the models presented to date do not account for the many regularities observed in the data from experiments. We provide a model that accounts for several important regularities of double auction data.

We model an informationally decentralized decision making procedure for sellers and buyers. Sellers form beliefs that an ask will be taken by some buyer. Similarly, buyers form beliefs that a bid will be taken by a seller. These beliefs are formed on the basis of observed market data, including frequencies of asks, bids, accepted asks, and accepted bids. Then traders choose an action that maximizes their own expected surplus. While traders in this model form beliefs about the probability that a given action they choose will result in a transaction, they have no beliefs about the types (costs or valuations) or strategies of other traders. The trading activity resulting from these beliefs is sufficient to achieve transaction prices near competitive equilibrium and complete market efficiency after several periods of trading.

We would like to thank Vernon Smith and Arlington Williams for providing data for comparison with the theory in this paper. We thank Dan Friedman, Leo Hurwicz, John Ledyard, Kevin McCabe, Charlie Plott, Vernon Smith, and Arlington Williams for helpful discussions. We would like especially to thank Jim Jordan for his help with the formulation of this model.

I Introduction

A variety of financial instruments and commodities trade in organized markets which are variants of an “open outcry” market. In these markets prices form via interactions of several traders. Many variations of the trade execution procedures used in these markets are possible. For example, in a clearinghouse (CH) market, sellers submit a quantity and the minimum price at which they are willing to trade this quantity, and buyers submit a quantity and the maximum price at which they are willing to buy that quantity. The trade execution algorithm then specifies that a price be chosen which maximizes quantity traded. All buyers with bids above the specified price then purchase units at that price from sellers with asks at or below that price. Results of theoretical investigation of the CH can be found in Rustichini, Satterthwaite, and Williams [1994] and references therein. Another common institution is the continuous double auction (DA) market. In this auction sellers submit ask prices and buyers submit bid prices until there is an ask less than or equal to the current high bid, or until a bid meets or exceeds the current low ask, at which point a trade is executed. The clearinghouse market and double auction market are two among a wide variety of trading institutions, each of which can be thought of as a composite mapping from the economic environment to outcomes or allocations, as in Hurwicz [1972] or Smith [1982]. The first mapping in this composition is the behavioral function from traders’ characteristics and current market history to messages. The second is the map from traders’ messages to outcomes or allocations.

While many specifications of trading institutions are possible, many security and commodity markets – such as the New York Stock Exchange or the Chicago Mercantile Exchange – are organized as variants of the DA. Since the environments of markets such as the NYSE involve assets yielding uncertain dividends, and since the valuations, costs, and information of traders cannot be observed, the data of laboratory double auctions – where these aspects of the environment can be controlled – provide a natural place to begin developing a theory of the price formation process.

Although competitive equilibrium theory has been a powerful tool in the investigation of economic phenomena for over a century, the interpretation of the theory as a description of the price formation process is problematic. Equilibrium theory is silent on the issue of how information is coordinated and market clearing prices are attained. The usual interpretation of the price formation process in the ‘tatonnement’ literature is that an ‘auctioneer’ calls out prices and traders submit net trades desired at these prices. The auctioneer then adjusts prices so that excess demand in the markets is reduced, and desired net trades at the new prices are submitted by traders. Hahn [1987] discusses several features of the tatonnement process which do not correspond closely to features of actual markets. The following exemplify these features.

Tatonnement processes typically involve movement toward equilibrium in many markets simultaneously. In markets such as the NYSE, prices form in each market separately. In addition, price formation in the tatonnement process involves repeated reporting of excess demand to a central auctioneer who then responds with adjustments to the price vector. In actual markets, the adjustment process takes place in real time with out of equilibrium trades being common. Hahn states that equilibrium in the tatonnement framework is not based on a theory of the interactions of rational market participants, and goes on to claim

that “Tatonnement ... sidesteps the important question of the coordinating power of the price mechanism.”

The double auction is a non-tatonnement price formation mechanism. In this mechanism, sellers place ask prices (also called offers) and buyers post bids. There is a specified time period for trading, and bids and asks arrive continuously during this time period. Trade takes place any time there is a ‘crossing’ of asks and bids. Empirical investigation of trader behavior and of market performance within this mechanism began with Smith [1962]. Smith induced supply and demand conditions by giving buyers a redemption value for each unit of an abstract commodity purchased, and by giving sellers a cost for each unit of this abstract commodity sold. Buyers receive surplus equal to the difference between this redemption value and the purchase price negotiated with a seller, and sellers receive surplus equal to the difference between the purchase price paid by the buyer and the unit cost. Since reservation prices – and therefore supply and demand conditions – are known when this procedure is employed, the procedure makes comparison between experimental outcomes and theoretical predictions possible. The basic result found by Smith [1962] in these experiments is that prices do converge quickly to within a few cents of competitive equilibrium prices in markets with stationary supply and demand. Smith and many other economists in the 30 years since his pioneering work have also documented features of the path of convergence to equilibrium in a variety of market environments. For summaries of these results see Plott [1982] and Smith [1982].

In this paper we model individual decision making in the DA which leads to market equilibrium prices. In the decision procedure developed here, proposed prices come from market participants rather than from an external agent such as an auctioneer responsible for determining market prices. In this respect, the behavioral rules and institution we describe corresponds more closely to actual markets than the tatonnement literature does.

Gode and Sunder [1993] show that in markets with a single commodity, random behavior within the DA is sufficient to achieve efficient outcomes. We demonstrate that it is possible to achieve competitive equilibrium outcomes (prices and allocations) in a market when individual choices are made myopically using heuristic beliefs, and that these beliefs are flexible enough that traders will respond to unanticipated changes in market conditions by moving to the new market equilibrium.

The organization of the paper is as follows. Four models of the double auction are discussed in Section II. The model is formulated in Section III. Simulations of the model are shown and some important statistical properties of these simulations are summarized in Section IV. Section V concludes.

II Theories of DA Price Formation

Several theories of the price formation process in the double auction have been proposed. These theories are listed below in increasing order of the degree of rationality attributed to market participants. We provide a model that accounts for a set of empirical regularities of behavior and market outcomes that no other model predicts or replicates in simulations. The first theory models traders who are neither adaptive nor strategic. The next two model behavior which is adaptive but not strategic, and in the last theory behavior is both adaptive and strategic.

Zero intelligence traders

Gode and Sunder [1993] model Zero Intelligence (ZI) traders who have no memory of past market activity, and choose their actions randomly. Each trader has an equal probability of being the next trader to make a bid or an ask. Sellers make asks which are random and uniformly distributed on the interval $[c, M]$, where c is the seller's cost, and M is some upper bound on their set of possible choices. Buyers make bids that are uniformly distributed on $[0, v]$, where v is the buyer's valuation. No learning takes place in this model. The actions of sellers and buyers do not change on the basis of previous experience in the market.

Easley-Ledyard model

Easley and Ledyard [1993] model traders who form strategies based directly on observed market data, rather than on choices determined from beliefs about characteristics and strategies of other traders. Buyers bid up as in a reduced form English auction to a reservation price which is formed on the basis of observed market data. Traders observe \underline{P}_n , the minimum transaction price or ask price in period n , and \overline{P}_n , the maximum transaction price or bid in that period. In period $n + 1$, reservation prices are formed by buyers on the interval $[\min\{v, \underline{P}_n\}, \min\{v, \overline{P}_n\}]$. In effect, this results in trading activity which narrows across periods to intervals determined by the prices (and offers and bids) observed in the previous period, allowing for convergence to equilibrium across several periods of trading. The clearest difference between their model and the model presented in Section III of this paper is the case of markets in which traders experience changes in supply or demand conditions, as shown in figure 7. In the model of Section III, prices will adjust within a few transactions to the new equilibrium in markets of this type, whereas in the Easley-Ledyard model, prices in period $n + 1$ remain within the range $[\underline{P}_n, \overline{P}_n]$. The traders modeled in Section III also make use of a larger set of information from past trading activity in forming beliefs and strategies than do traders in the Easley-Ledyard model.

Bayesian game against nature

Friedman [1991] models the agents' decision problem as a Bayesian game against nature (BGAN). The traders in his model "neglect strategic feedback effects" and play as though asks, bids, and acceptances come from distributions which are unaffected by the trader's own strategy choice. This assumption is analogous to price taking behavior in the theory of competitive markets. The analogy is that traders take the distributions of asks, bids, and acceptances as given, and do not behave as though their own actions will affect the beliefs held or strategies chosen by other players. The model in Section III shares this feature with the BGAN model.

One conclusion that can be drawn from the performance of the model in this paper is that the mathematical consistency of Bayesian updating is not necessary to achieve equilibrium price formation. Also, as shown in Gjerstad [1995], our model produces negative autocorrelation of price changes. Friedman's model predicts positive autocorrelation of price changes, whereas there is substantial evidence that in markets with stationary supply and demand conditions, successive price changes in double auctions are negatively autocorrelated.

Bayes–Nash equilibrium model

Wilson [1986] formulates a model in which prices form as the Bayes–Nash equilibrium of a game of incomplete information. This approach to modeling the double auction extends earlier work (e.g., Chatterjee and Samuelson [1984]) on bargaining under asymmetric information. Wilson formulates the problem as a sequence of bilateral matchings of sellers and buyers. Each seller and each buyer has prior beliefs on the distribution of types (cost and valuation parameters) and on the equilibrium strategies of other traders. The collection of these prior beliefs for the traders is common knowledge. After observing ask, bid and transaction data each trader updates beliefs about the distribution of types optimally (according to Bayes' rule).

Since the structure of Bayes–Nash equilibrium depends on the assumption of common knowledge of the beliefs of traders about the distribution of types, and common knowledge of the equilibrium strategies employed by other traders, and since in practice subjects come into experiments with very little information about the environment or the strategies likely to be employed by other traders, Wilson's approach overstates the extent of rationality required to produce a model of the price formation process. Moreover, Ledyard [1986] shows that the Bayes–Nash equilibrium concept is capable of accounting for any outcome in a game that does not involve a player choosing a dominated strategy, as long as traders' cardinal utility functions and prior beliefs about the distribution of players types and strategies are not controlled.

These four models and the model of Section III can be classified according to two criteria. All the models with the exception of the Gode and Sunder model are adaptive. Traders acquire information through the trading process; use of this information allows traders to make decisions that result in some stability of transaction prices. While the ZI traders of Gode and Sunder are able to extract most of the rents in markets populated with other ZI traders, these strategies do not lead to the formation of an equilibrium market price or a narrow range of transaction prices. A second criterion is strategic behavior. The Wilson model is the only strategic model, in the sense that traders anticipate the effects of their actions on the strategies of the other traders. As noted above, Wilson's approach does not address how beliefs are formed. In addition, the computational complexity inherent in that formulation makes it difficult to implement the model and evaluate outcomes for various belief specifications. Moreover, Wilson models the trading process within a single trading period and makes no prediction about the price paths across periods. The motivation for our model is to strike a balance between the approach taken by Wilson and that of Gode and Sunder, and at the same time consider beliefs which avoid the rigidity of prices found in the Easley–Ledyard model, and avoid the positive autocorrelation of price changes found in Friedman's model. In formulating the trading in a double auction as the Bayes–Nash equilibrium of a game of incomplete information, Wilson's model depends on the assumption of common knowledge of beliefs and strategies, but we do not know what beliefs participants in a double auction hold about the distribution of types and strategies, and we certainly don't know that the collection of these beliefs is common knowledge. On the other hand, in the ZI model traders make choices randomly, making no use of the information from trading that we know participants do have. The Easley–Ledyard model

and Friedman's model circumvent these problems: nevertheless, our model is consistent with a set of empirical features that is not found in any one of the four alternative models (see Gjerstad [1995]). Perhaps most importantly, though, the model in Section III is the only existing model of DA price formation which has been implemented in simulations *and* converges to (approximate) competitive equilibrium prices. This allows investigation of the path of convergence to equilibrium, and provides a basis for extension to more complex environments.

III The Model

The double auction is an example of a *microeconomic system* as in Hurwicz [1972] and Smith [1982]. The primary features of an economic system are the *environment* \mathbf{E} , consisting of the characteristics of the economic agents, and the *institution* \mathbf{I} , which includes the messages that traders may send to one another, the allocation rules, and the adjustment process rules. A microeconomy is an economic system $\mathbf{S} = (\mathbf{E}, \mathbf{I})$ – together with behavioral actions β^i for market participants – as shown in figure 1.

The environment \mathbf{E} consists of a set $\mathcal{A} = \{1, 2, \dots, n\}$ of agents. Each agent i has characteristics $e^i = (P^i, Y^i, \omega^i)$ where P^i , Y^i , and ω^i are that agents' preferences, technology, and endowment, respectively. The environment is then $\mathbf{E} = \prod_{i \in \mathcal{A}} e^i$. The institution \mathbf{I} consists of a message space M^i for each agent, an adjustment process rule specifying the sequence of agent messages, and an outcome function or allocation function $h(m_t) = (h^1(m_t), h^2(m_t), \dots, h^n(m_t))$, where $m_t = (m_t^1, m_t^2, \dots, m_t^n) \in M_t = \prod_{i \in \mathcal{J}} M_t^i$ is the vector of agents' messages.

According to Smith ([1982], p. 930)

“We want to measure messages because we want to be able to identify the behavioral modes, $\beta^i(e^i, I)$, revealed by the agents and test hypotheses derived from theories about agent behavior.”

When an environment \mathbf{E} and an institution \mathbf{I} are specified in a market experiment, and an outcome \mathbf{X} is observed, the only elements remaining to be specified are the behavioral actions $\{\beta^i(H_t | e^i, \mathbf{I})\}_{i \in \mathcal{A}}$, where H_t is the history of activity observed by agents through time t . The focus of the research in this paper is to specify forms of behavior $\{\beta^i(H_t | e^i, \mathbf{I})\}_{i \in \mathcal{A}}$ that are consistent with observations \mathbf{X} for a variety of environments \mathbf{E} when the institution \mathbf{I} is the double auction.




Figure 1: A dynamic microeconomic system.

Note 1 This notation for a dynamic adjustment process differs from the usual notation. Here the adjustment $m_{t+1}^i = \beta^i(H_t | e^i, \mathbf{I})$ depends on the history H_t of previous messages, rather than on only the most recent message m_t . (A complete description of the history vector H_t observed by market participants in this institution is given in Section B.3.)

This framework can be used to describe a double auction experiment. The environment of a DA experiment is described in Section A. The DA institution is described in Section B. The model of behavior is presented in Section C.

A Environment

The theory developed in this paper addresses the case of markets with the set \mathcal{A} of traders partitioned into a group I of sellers and a group J of buyers. We interpret a market experiment as an example of an economic system. Both types of agents (sellers and buyers) are assumed to have preferences over monetary rewards that are monotonically increasing.

Each seller $i \in I$ has a technology $Y^{s,i}$ for acquisition of units of an abstract commodity at costs $c_i^1, c_i^2, \dots, c_i^{m_i}$, where c_i^1 is the cost to seller i of the first unit sold, c_i^2 the cost of the second unit, and so on. The gains to seller i on the k^{th} unit sold are determined by the difference $\pi_{s,i}^k(p_k, c_i^k) = p_k - c_i^k$ between the price p_k received from a buyer for that unit, and the cost c_i^k at which the unit is aquired. If seller i sells $\mu_i \leq m_i$ units at prices $p_1, p_2, \dots, p_{\mu_i}$, then the utility to this seller is $u^i\left(\sum_{k=1}^{\mu_i} (p_k - c_i^k)\right)$, where $u^i(\cdot)$ is assumed monotonically increasing.

Each buyer $j \in J$ has a technology $Y^{b,j}$ used for redemption of units of the abstract commodity at unit valuations $v_j^1, v_j^2, \dots, v_j^{n_j}$, where v_j^1 is the redemption value for the first unit aquired, v_j^2 is the redemption value for the second, and so forth. Buyer j has an endowment ω_j of a trading currency that is sufficient to purchase each unit at a price up to the redemption value of the unit, i.e., $\omega_j \geq \sum_{l=1}^{n_j} v_j^l$. Monetary rewards for buyers are the difference between the redemption values of units purchased and the price p_l paid to a seller:

$$\pi_{b,j}^l(p_l, v_j^l) = v_j^l - p_l.$$

If buyer j purchases $\nu_j \leq n_j$ units at prices $p_1, p_2, \dots, p_{\nu_j}$ the monetary gain from trading for buyer j is $\sum_{l=1}^{\nu_j} (v_j^l - p_l)$, and the utility of this monetary gain is $w^j\left(\sum_{l=1}^{\nu_j} (v_j^l - p_l)\right)$, where $w^j(\cdot)$ is monotonically increasing.

The environment is therefore

$$\begin{aligned} e &= \left\{ (u^i, Y^{s,i}, 0) \right\}_{i \in I} \cup \left\{ (u^j, Y^{b,j}, \omega_j) \right\}_{j \in J} \\ &= \left\{ (u^i, c_i, 0) \right\}_{i \in I} \cup \left\{ (u^j, v_j, \omega_j) \right\}_{j \in J}. \end{aligned}$$

Example 1 Figure 2 shows supply and demand conditions for market trading experiment 3pda01 run in the experimental lab at the University of Arizona by Vernon Smith and Arlington Williams. In this market there are four buyers, each with positive valuations for three units, and four sellers, each with finite costs for three units. The vector of buyers' valuations is

$$\begin{aligned} v &= \{v_1, v_2, v_3, v_4\} \\ &= \{\{\$3.30, \$2.25, \$2.10\}, \{\$2.80, \$2.35, \$2.20\}, \\ &\quad \{\$2.60, \$2.40, \$2.15\}, \{\$3.05, \$2.35, \$2.30\}\}. \end{aligned}$$




Figure 2: Supply and demand conditions for market experiment 3pda01.

The vector of sellers' costs is

$$\begin{aligned}
 c &= \{c_1, c_2, c_3, c_4\} \\
 &= \{\{\$1.90, \$2.35, \$2.50\}, \{\$1.40, \$2.45, \$2.60\}, \\
 &\quad \{\$2.10, \$2.30, \$2.55\}, \{\$1.65, \$2.35, \$2.40\}\}.
 \end{aligned}$$

Since buyer j with redemption value v_j^l makes a monetary gain at any purchase price $p < v_j^l$, and since the buyers preferences are assumed monotonically increasing in monetary gain, this buyer is willing to pay any price up to v_j^l for the l^{th} unit purchased. Therefore, the demand shown in figure 2 is determined by arraying the buyers' redemption value vectors. Supply is obtained analogously. Between five and seven units are exchanged at the competitive equilibrium price. Each side of the market receives the same rents at the equilibrium price of \$2.35. The total surplus available is \$4.80 per trading period.

In what follows, this market will be referred to as the symmetric market design.

Trading periods

A typical laboratory market experiment involves trading over several periods. Each seller has costs induced for the trading period, and each buyer has valuations induced. A buyer's valuation for a unit remains in effect throughout the trading period or until the buyer transacts that unit. After a unit is transacted, the seller's cost and the buyer's valuation for the unit just transacted are no longer in the supply and demand schedules, trading continues, and this process proceeds until there are no more surplus enhancing

trades remaining, or until time expires in the trading period. At the conclusion of a trading period, the costs and valuations are reinitialized – possibly at different amounts – in the subsequent period. In market experiment 3pda01 – using the symmetric market design of figure 2 – there were nine trading periods, each lasting 300 seconds.

B Institution

B.1 Messages and adjustment process rules (asks and bids)

Since traders receive surplus based on the transaction price, there must be a means of determining the price. In the double auction sellers post ask prices, and buyers post bids. In an economic mechanism, adjustment process rules specify the time when exchange of messages begins, a transition rule governing the sequencing and exchange of messages, and a stopping rule. The DA imposes no restrictions on the sequencing of messages. Any trader can send a message at any time during the period that the market is open. Allocation of units is by mutual consent between any two traders. If a seller's ask is acceptable to a buyer then a transaction is completed when the buyer takes (accepts) the seller's ask. Similarly, a buyer's bid may be accepted by a seller.

Definition 1 (The message space) Seller i at time t has a message space $M_t^{s,i}$ where

$$M_t^{s,i} \subset \{i\} \times \{0\} \times \tilde{N}.$$

Buyer j has a message space $M_t^{b,j}$ where

$$M_t^{b,j} \subset \{0\} \times \{j\} \times \tilde{N},$$

where $\tilde{N} = \{x : x = \frac{n}{100} \text{ for some } n \in N\}$.

Definition 2 (Asks) An ask a by seller i is an amount which seller i is willing to accept from a buyer as payment for a unit of the commodity being traded. A take of ask a by a buyer j is an agreement by j to purchase a unit from seller i at the transaction price $p = a$.

Definition 3 (Bids) A bid b by buyer j is an amount that buyer j is willing to pay to some seller for a unit. A take of bid b by a seller i is an agreement by i to sell a unit to buyer j at the transaction price $p = b$.

The lowest ask in the market at any time is called the outstanding ask oa and the highest bid is called the outstanding bid ob . These define the bid-ask spread $[ob, oa]$.

One commonly used market rule requires a bid-ask spread reduction in the posting of new bids and asks. (The DA with the spread reduction rule is the institution considered in this paper, although the strategies considered could be applied if asks and bids are queued.) At any time buyers may place a bid $b \in \tilde{N}$ with $b > ob$ or they may accept the outstanding ask. Similarly, sellers may accept the outstanding bid, or they may make an ask $a \in \tilde{N}$ below the outstanding ask.

Since messages are sent in an iterative process, they are subscripted in the order received, as in example 2.

Example 2 (Messages) Suppose the market depicted in figure 2 is in progress, and the actions have been a bid of \$2.00 by buyer 3, and an ask of \$2.65 by seller 4. These are messages m_1 and m_2 . They are written

$$m_1 = (0, 3, \$2.00)$$

and

$$m_2 = (4, 0, \$2.65).$$

With m_1 and m_2 as above, the message

$$m_3 = (0, 1, \$2.25)$$

is permissible, but

$$m_4 = (3, 0, \$2.75)$$

does not meet the condition of the improvement rule, since the outstanding ask is \$2.65, and seller 3's proposed ask of \$2.75 does not improve on the ask seller 4 submitted with message m_2 .

B.2 Allocations (takes)

At any time during the operation of a market, if there is an outstanding ask oa posted, then buyer j with valuation v_j^l has the option of accepting the outstanding ask (and therefore purchasing a unit at price $p = oa$), resulting in surplus

$$\pi_{b,j}^l(oa, v_j^l) = v_j^l - oa.$$

Similarly, if there is an outstanding bid ob then seller i with cost c_i^k has the option of accepting ob and receiving surplus

$$\pi_{s,i}^k(ob, c_i^k) = ob - c_i^k.$$

Example 3 (Acceptances) With the history of messages

$$m_1 = (0, 3, \$2.00)$$

$$m_2 = (4, 0, \$2.65)$$

$$m_3 = (0, 1, \$2.25),$$

the outstanding ask $oa = \$2.65$ could be accepted by buyer 1 at this point¹ by sending the message

$$m_4 = (0, 1, \$2.65).$$

At that point a trade would be executed between seller 4 and buyer 1 at the price \$2.65. This is the outcome of message 4. It is denoted

$$h_4 = (4, 1, \$2.65).$$

¹This example is the actual sequence of bids and asks leading up to the first transaction in period 2 of market experiment 3pda01. The sequence of transaction prices for this market is shown in figure 3.

B.3 Histories

After n messages have been sent there will be a history of length n which summarizes the activity in the market to that point. In example 3, the messages m_1, \dots, m_4 result in the history

$$\begin{aligned} H_4 &= \{h_1, h_2, h_3, h_4\} \\ &= \{m_1, m_2, m_3, h_4\} \\ &= \{(0, 3, \$2.00), (4, 0, \$2.65), \\ &\quad (0, 1, \$2.25), (4, 1, \$2.65)\}. \end{aligned}$$

Note that in H_4 the triple h_4 unambiguously denotes an accept by buyer 1 of the offer of \$2.65 made with message 2 by seller 4. (See ‘Accept of oa’ in Definition 4.)

Definition 4 (Histories) After n messages have been sent, there will be a history H_n consisting of n ordered triples. For any message m_{n+1} that is sent, one of six cases will hold.

Invalid ask or bid A message $m_{n+1} = \{i, 0, a\}$ is not valid if $a \geq oa$. An invalid ask will not be included in the history. In effect, the institution ignores messages that violate the spread reduction rule. Similarly, a message $m_{n+1} = \{0, j, b\}$ is not valid if $b \leq ob$.

No ask outstanding If no ask has been made since the last transaction, then there is no outstanding ask, and any ask $a \in \tilde{N}$ is valid. If, in addition, $m_{n+1,3} > ob$, then $h_{n+1} = m_{n+1}$.

No bid outstanding Similarly, if no bid has been made since the last transaction, then there is no outstanding bid, and any bid $b \in \tilde{N}$ is valid. If $m_{n+1,3} < oa$, then $h_{n+1} = m_{n+1}$.

Accept of ob If $m_{n+1,1} \neq 0$ and $m_{n+1,3} \leq ob$ then seller $m_{n+1,1}$ is making an offer at or below ob , so m_{n+1} is an accept of ob. The buyer’s identity is found by looking back in H_n and finding the last h_k with $h_{k,2} \neq 0$, that is $k^* = \max\{k : h_{k,2} \neq 0\}$. Then $(h_{n+1,1}, h_{n+1,2}, h_{n+1,3}) = (m_{n+1,1}, h_{k^*,2}, ob)$.

Accept of oa If $m_{n+1,2} \neq 0$ and $m_{n+1,3} \geq oa$ then m_{n+1} is an accept of oa (by buyer $m_{n+1,2}$). The seller’s identity is found by looking back in H_n and finding $k^* = \max\{k : h_{k,1} \neq 0\}$. Then $(h_{n+1,1}, h_{n+1,2}, h_{n+1,3}) = (h_{k^*,1}, m_{n+1,2}, oa)$.

Improving ask or bid If $m_{n+1,3} \in (ob, oa)$ then m_{n+1} is either an improving ask, or an improving bid, and $h_{n+1} = m_{n+1}$.

C Behavior

C.1 Frequencies of takes

As noted in Section A, sellers attempt to maximize $\pi_{s,i}^k(p_k, c_i^k)$ and buyers attempt to maximize $\pi_{b,j}^l(p_l, v_j^l)$. Since asks or bids must be accepted in order to result in a transaction, we take the point of view that sellers will maximize *expected* surplus myopically, where the

expectation is taken relative to beliefs $p(a)$ that an ask a will be accepted by some buyer. These beliefs are formed on the basis of observed market data (as described in Section C.2). Similarly, buyers are assumed to maximize expected surplus myopically, where the expectation is taken relative to beliefs $q(b)$ that a bid b will be accepted by some seller. In both cases the domain of these beliefs is the set D of all asks and bids that have occurred in the market within the memory length specified for the trader.

The history that traders consider in forming beliefs is restricted to those messages leading up to the last L transactions, where $L \in \{0, 1, 2, \dots\}$. The parameter L is therefore the memory length of the trader. The next definition provides a procedure for truncating the history, so that beliefs can be constructed using the data within the trader's memory. The procedure for constructing beliefs using this (truncated) history is described in Section C.2.

Note 2 We will work with the vector H_n , although traders do not have access to all the information in H_n . Specifically, traders know their own asks or bids, but do not necessarily know the identities of the traders making the other bids and asks. Information about identities is not used in the formation of beliefs or strategies, so use of H_n is made only to avoid complicating notation.

Definition 5 (Remembered history) Let H_n be the history vector for a market after n messages have been sent, as in definition 4. Given H_n , make the following definitions.

Trade function For a vector H_n , define a function $T : H_n \mapsto \{0, 1\}^n$ by setting $T_k(h_k) = I_{\{h_{k,1} \cdot h_{k,2} > 0\}}(h_k)$. Then each component T_k of T indicates whether a trade occurred in the k -th element of the history.

Number of trades Let $x^n = (x_1, x_2, \dots, x_n)$. For each n , define $S_n : \{0, 1\}^n \mapsto \mathbb{N}$ by

$$S_n(x^n) = \sum_{k=1}^n x_k.$$

Then $S_n(T(H_n))$ is the number of trades resulting from the first n messages.

Remembered history Let L be the memory length of a given trader. For fixed n and H_n , to simplify notation, let $S = S_n(T(H_n))$. Let n' be the position of trade $S - L$ if $S > L$, and let $n' = 0$ if $S \leq L$. Define $\tilde{H}_n^{(L)}$ by

$$\tilde{H}_n^{(L)} = \{h_{n'+1}, h_{n'+2}, \dots, h_n\}.$$

Deletion of oa and ob from history Let $n'' = \max\{k : T_k(h_k) = 1\}$, i.e., n'' is the index of the most recent trade. Let n^* be the index of the lowest (most recent) ask in the vector $\{h_{n''+1}, \dots, h_n\}$. Let n_* be the index of the highest bid in the vector $\{h_{n''+1}, \dots, h_n\}$. Note that if $T_n(h_n) = 0$, then $h_{n,3}$ is either the outstanding ask or the outstanding bid, as a consequence of the spread reduction rule, and if $T_n(h_n) = 1$ then there is no outstanding bid and no outstanding ask. If $T_n(h_n) = 0$ and $h_{n,1} = 0$, then $h_{n,3}$ is the outstanding bid. If $h_{k,1} \neq 0$ for some $k \in \{n'' + 1, \dots, n - 1\}$, then $n^* \neq \emptyset$ and we can define $H_n^{(L)}$ by

$$H_n^{(L)} \equiv \{h_{n'+1}, h_{n'+2}, \dots, h_{n^*-1}, h_{n^*+1}, \dots, h_{n-1}\}.$$

That is, $H_n^{(L)}$ is $\tilde{H}_n^{(L)}$ with h_{n^*} and h_{n_*} removed. This is done because it is not known at time n if the outstanding ask or bid will be accepted. The other case – where $h_{n,3}$ is the outstanding ask – is treated similarly. Then $H_n^{(L)}$ is the part of the history from which traders with memory length L form beliefs.

Set of asks and bids Let $D_n^{(L)}$ be the set of all asks and bids that have been made in $H_n^{(L)}$, i.e., $D_n^{(L)} \equiv \bigcup_{k \in \{n'+1, \dots, n\} \setminus \{n_*, n^*\}} \{h_{k,3}\}$.

Definition 6 (Ask frequencies) For each $d \in D_n^{(L)}$, let $A(d)$ be the total number of asks that have been made at d , and let $TA(d)$ be the total number of these that have been accepted. Let $RA(d) \equiv A(d) - TA(d)$ be the rejected asks at d .

For $A(d)$, the counting procedure is as follows. For each

$$k \in \{n' + 1, \dots, n\} \setminus \{n_*, n^*\},$$

if $h_{k,3} = d$, $h_{k,1} \neq 0$ and $h_{k,2} = 0$, then $A(d)$ is incremented by one. If $h_{k,3} = d$ and $T_k(h_k) = 1$, then h_k is either a taken ask or a taken bid. To determine which is the case, find $m^* = \min\{m \geq 1 : h_{k-m,3} = h_{k,3}\}$. If $h_{k-m^*,1} \neq 0$, then $A(d)$ and $TA(d)$ are incremented by one. The rejected asks at d are given by $RA(d) \equiv A(d) - TA(d)$.

At each time during a market, the proportion of asks at $a \in D$ that have been accepted is

$$\check{p}(a) = \frac{TA(a)}{A(a)}$$

whenever $A(a) > 0$.

Definition 7 (Bid frequencies) For each $d \in D_n^{(L)}$, let $B(d)$ be the total number of bids that have been made at d , and let $TB(d)$ be the total number of these that have been accepted. Let $RB(d) \equiv B(d) - TB(d)$. The interpretations and counting procedures for $B(d)$, $TB(d)$, and $RB(d)$ are analogous to those described in definition 6 for asks.

At each time during a market, the proportion of bids at $b \in D$ that have been accepted is

$$\check{q}(b) = \frac{TB(b)}{B(b)}$$

whenever $B(b) > 0$.

Note 3 In what follows, the sets of asks and bids is frequently denoted D , with the subscripts and superscripts omitted. When traders have finite memory, that will be noted. After n messages have been sent, the relevant set of asks and bids is $D_n^{(L)}$ and the relevant history is $H_n^{(L)}$.

In stationary market environments these empirical frequencies show substantial regularity: $\check{p}(a)$ tends to be a decreasing function of a and $\check{q}(b)$ tends to be an increasing function of b .

C.2 Beliefs

While the frequencies $\check{p}(a)$ and $\check{q}(b)$ from definitions 6 and 7 tend to be monotonic when the number of asks and bids is large, there is more variability in small samples. For this reason, it is useful to work with a modification of these summary statistics.

Modification of $\check{p}(a)$ is made by taking the point of view that if an ask $a' < a$ is rejected then had that ask been made at a it would also have been rejected. This assumption is made because $a > a'$ and is therefore less appealing to buyers than a' , which was rejected. Similarly, if ask $a' > a$ was made and taken, then that ask would also have been taken if it were made at a . Also, if a bid $b' > a$ is accepted, then an ask $a' = b'$ would have been taken if it had been made (the assumption being that this ask of a' would be acceptable to the buyer who bid b'). This heuristic – and an analogous one for buyers' beliefs – are formalized in the next two definitions.

Definition 8 (Sellers' beliefs) For each potential ask $a \in D$, define

$$\hat{p}(a) = \frac{\sum_{d \geq a} TA(d) + \sum_{d \geq a} TB(d)}{\sum_{d \geq a} TA(d) + \sum_{d \geq a} TB(d) + \sum_{d < a} RA(d)}.$$

Then $\hat{p}(a)$ is the seller's belief that an ask amount a will be acceptable to some buyer.

The notation is simplified by the following definitions. Let $TAG(a) = \sum_{d \geq a} TA(d)$, $TBG(a) = \sum_{d \geq a} TB(d)$, and $RAL(a) = \sum_{d < a} RA(d)$. These are the taken asks greater than or equal to a , the taken bids greater than or equal to a , and the rejected asks less than or equal to a , respectively.

Then

$$\hat{p}(a) = \frac{TAG(a) + TBG(a)}{TAG(a) + TBG(a) + RAL(a)}.$$

Definition 9 (Buyers' beliefs) For each possible bid $b \in D$, define

$$\hat{q}(b) = \frac{\sum_{d < b} TB(d) + \sum_{d < b} TA(d)}{\sum_{d < b} TB(d) + \sum_{d < b} TA(d) + \sum_{d \geq b} RB(d)}.$$

As in definition 8, to simplify notation we introduce functions $TBL(b) = \sum_{d < b} TB(d)$, $TAL(b) = \sum_{d < b} TA(d)$, and $RBG(b) = \sum_{d \geq b} RB(d)$. These are the taken bids less than or equal to b , the taken asks less than or equal to b , and the rejected bids greater than or equal to b .

Then

$$\hat{q}(b) = \frac{TBL(b) + TAL(b)}{TBL(b) + TAL(b) + RBG(b)}.$$

With the specification of beliefs in definitions 8 and 9, the belief functions are monotonic. (See propositions 1 and 2 in Section C.4.) This captures the intuition that a trader who has seen an ask of a rejected should decrease the belief that a will be accepted later, and decrease the belief that an ask at any value greater than a will be accepted.

C.3 Spread reduction rule and beliefs

The spread reduction rule has the effect of making the probability of a take for an ask $a \geq oa$ equal to 0. Similarly, $\tilde{q}(b) = 0$ for all b with $b \leq ob$. These facts are incorporated into traders' beliefs in the first parts of definitions 10 and 11 below.

In order that the bidding and asking are capable of beginning, there must be some action that traders believe will result in a positive expected surplus to themselves. This is accomplished by placing a lower bound of η on the belief of a take for each action being considered by the traders, and by always having 0 and M for some $M > \max_{j,l} \{v_j^l\}$ as potential actions, i.e., $\{0, M\} \subseteq D$ for all n , and for all L . In the simulations of section IV, the value of η has been fixed at $\eta = 0.001$, and the value of M has been fixed at $M = 10$. (These two parameters simply insure that traders' prior beliefs are diffuse.)

Definition 10 Let

$$\tilde{p}(a) = \hat{p}(a) \cdot I_{[0, oa)}(a)$$

for all $a \in D$. That is $\tilde{p}(a) = \hat{p}(a)$ if $a < oa$ and $\tilde{p}(a) = 0$ if $a \geq oa$. For fixed $\eta > 0$, let $p(a) = \max\{\tilde{p}(a), \eta\}$ for each $a \in D$.

Definition 11 Let

$$\tilde{q}(b) = \hat{q}(b) \cdot I_{(ob, M]}(b)$$

for all $b \in D$. Let $q(b) = \max\{\tilde{q}(b), \eta\}$.

C.4 Monotonicity of beliefs

The function $\hat{p}(a)$ defined in Section C.2 is monotonically non-increasing. That is, as the ask a is increased, $\hat{p}(a)$ – the belief that an ask a will be accepted – is non-increasing in a . Similarly, $\hat{q}(b)$ is non-decreasing in b : higher bids are believed to be more likely to be taken. These results are proven in propositions 1 and 2.

Proposition 1 For all $a_1 \in D$, $a_2 \in D$, with $a_1 < a_2$, $\hat{p}(a_1) \geq \hat{p}(a_2)$.

Proof Let $TG(a) = TAG(a) + TBG(a)$. Note that

$$TG(a_2) \leq TG(a_1)$$

and

$$RAL(a_1) \leq RAL(a_2)$$

because $a_1 < a_2$. Multiplying these two inequalities results in

$$TG(a_2) RAL(a_1) \leq TG(a_1) RAL(a_2). \quad (1)$$

Now add

$$TG(a_2) TG(a_1)$$

to both sides of inequality (1) and from this sum factor out $TG(a_2)$ from the left side of the equation and factor $TG(a_1)$ out of the right side, then divide both sides of the resulting inequality by

$$[TG(a_1) + RAL(a_1)][TG(a_2) + RAL(a_2)]$$

to get $\hat{p}(a_1) \geq \hat{p}(a_2)$. ■

Proposition 2 For all $b_1 \in D$, $b_2 \in D$, with $b_1 < b_2$, $\hat{q}(b_1) \leq \hat{q}(b_2)$.

Proof The proof is similar to the proof of proposition 1. ■

Proposition 3 The functions $p(a)$ and $\tilde{p}(a)$ are non-increasing.

Proof Since $\tilde{p}(a) = \hat{p}(a) I_{[0, oa)}(a)$, and $\hat{p}(a)$ is non-increasing, $\tilde{p}(a)$ is also non-increasing. Since $p(a) = \max\{\tilde{p}(a), \eta\}$, $p(a)$ is also non-increasing. ■

Proposition 4 The functions $q(b)$ and $\tilde{q}(b)$ are non-decreasing.

Proof The proof is similar to the proof of proposition 3. ■

C.5 Expected surplus maximization

Seller i with cost $c_i^k < oa$ may make an offer $a \in (c_i^k, oa)$ and obtain the expected surplus

$$E[\pi_{s,i}^k(a, c_i^k)] = (a - c_i^k) \cdot p(a).$$

The maximum expected surplus of seller i for the sale of unit k is then

$$S_{s,i}^k = \max\{\max_{a \in D} E[\pi_{s,i}^k(a, c_i^k)], 0\}.$$
²

Similarly, buyer j with valuation v_j^l may make a bid $b \in (ob, oa)$ and obtain the expected surplus

$$E[\pi_{b,j}^l(b, v_j^l)] = (v_j^l - b) \cdot q(b).$$

The maximum expected surplus for buyer j purchasing unit l is therefore

$$S_{b,j}^l = \max\{\max_{a \in D} E[\pi_{b,j}^l(b, v_j^l)], 0\}.$$

C.6 Timing of messages

Let t be the parameter for time within a trading period. Let T be the length of the trading period and let $t_\kappa \in [0, T)$ be the time of the κ^{th} offer, bid, or acceptance of an offer or bid. At time t_κ let $T_{s,i}^\kappa$ be the random variable that specifies the time which seller i would allow to elapse before sending a message; let $T_{b,j}^\kappa$ be the random variable that specifies the time which buyer j would allow to elapse before sending a message. If $T_{s,i}^\kappa$ and $T_{b,j}^\kappa$ are exponentially distributed with parameters $\alpha_{s,i} = f_{s,i}(S_{s,i}^k; t_\kappa, T)$ and $\beta_{b,j} = f_{b,j}(S_{b,j}^l; t_\kappa, T)$, then the probability that seller i' will be the next trader to send a message in the market is

$$p_{s,i'} = \frac{f_{s,i'}(S_{s,i'}^k; t_\kappa, T)}{\sum_{i \in I} f_{s,i}(S_{s,i}^k; t_\kappa, T) + \sum_{j \in J} f_{b,j}(S_{b,j}^l; t_\kappa, T)}.$$

This is shown in the following proposition.

²Seller i with cost vector $c_i = \{c_i^1, c_i^2, \dots, c_i^{m_i}\}$ faces the problem of choosing a sequence of asks or accepts to maximize $\sum_{k=1}^{m_i} (p_i^k - c_i^k)$, where p_i^k is the purchase price received for unit k . We assume that the seller will attempt to maximize the surplus of each unit in sequence, independently of other units. In addition to simplifying the strategy choice, this is consistent with the myopic formulation of strategy choice. A similar remark applies to buyers.

Proposition 5 If $T_{s,i}^\kappa$ and $T_{b,j}^\kappa$ are independent exponentially distributed random variables on $[0, \infty)$ with parameters $\alpha_{s,i} = f_{s,i}(S_{s,i}^k; t_\kappa, T)$ and $\beta_{b,j} = f_{b,j}(S_{b,j}^l; t_\kappa, T)$ (where $S_{s,i}^k$ and $S_{b,j}^l$ are as defined in Section C.4), i.e.,

$$Pr\{T_{s,i} < t\} = 1 - e^{-\alpha_{s,i} \cdot t},$$

then the probability that seller i' will be the next trader to send a message is

$$p_{s,i'} = \frac{f_{s,i'}(S_{s,i'}^k; t_\kappa, T)}{\sum_{i \in I} f_{s,i}(S_{s,i}^k; t_\kappa, T) + \sum_{j \in J} f_{b,j}(S_{b,j}^l; t_\kappa, T)}.$$

Proof Consider, for example, seller 1. Let

$$T_{s,-1} = \min_{i>1, j \geq 1} \{T_{s,i}, T_{b,j}\}.$$

This random variable is exponentially distributed with parameter

$$\alpha_{s,-1} = \sum_{i>1} \alpha_{s,i} + \sum_{j \geq 1} \beta_{b,j}.$$

Then the probability that seller 1 is the next seller to move will be the probability that $T_{s,1} < T_{s,-1}$, i.e.,

$$\begin{aligned} p_{s,1} &= Pr\{T_{s,1} < T_{s,-1}\} \\ &= \int_0^\infty \int_0^u \alpha_{s,1} \cdot e^{-\alpha_{s,1} \cdot t} \cdot \alpha_{s,-1} \cdot e^{-\alpha_{s,-1} \cdot u} dt du \\ &= \frac{\alpha_{s,1}}{\alpha_{s,1} + \alpha_{s,-1}}. \blacksquare \end{aligned}$$

There are two reasons for defining the timing of each trader's actions as exponential random variables. The first is an important conceptual issue: with this formulation, the mechanism is informationally decentralized, in that the information about each trader's surplus is not held by any agent. The decision by an individual trader to send a message is independent of any (unobserved) characteristics – such as costs or valuations – of other traders. The second issue is empirical. With this formulation, the timing of bids and asks is testable within the model, and it is possible to compare the timing data for various specifications of the functions $f_{s,i}(S_{s,i}^k; t_\kappa, T)$ and $f_{b,j}(S_{b,j}^l; t_\kappa, T)$ with the timing data from experiments.

The specifications of $f_{s,i}(S_{s,i}^k; t_\kappa, T)$ and $f_{b,j}(S_{b,j}^l; t_\kappa, T)$ that have been used in the simulations of Section IV are the functions

$$f_{s,i}(S_{s,i}^k; t_\kappa, T) = S_{s,i}^k \cdot \frac{T}{(T - t_\kappa)}$$

and

$$f_{b,j}(S_{b,j}^l; t_\kappa, T) = S_{b,j}^l \cdot \frac{T}{(T - t_\kappa)}.$$

C.7 Randomization of asks and bids

In formulating the model, two specifications of the belief functions $p(a)$ and $q(b)$ have been considered, and simulations of the model for each formulation have been run and examined. The beliefs in Section C.2 are defined on the set D of bids and asks that have previously occurred during the course of (recent) trading.

One version of the belief functions that was considered involved forming piecewise linear functions. For ask values $a \in [0, M]$ between successive values $d_k \in D$ and $d_{k+1} \in D$, the beliefs were given by

$$p(a) = \frac{(a - d_k) \cdot p(d_{k+1}) + (d_{k+1} - a) \cdot p(d_k)}{d_{k+1} - d_k}.$$

With this piecewise linear specification of beliefs, the expected surplus function is piecewise quadratic and continuous. If for seller i there is an interior maximum d^* of the expected surplus function on some segment (d_k, d_{k+1}) , and if this is also a global maximum of the expected surplus function, then d^* will be the optimal bid for seller i . If seller i is the next trader to send a message, that message will be $(i, 0, d^*)$. When d^* is subsequently accepted or rejected, it will be included in D . While this possibility suggests that the piecewise linear formulation of beliefs could result in a sufficient amount of exploration of new prices, in simulations interior maxima occurred infrequently. As a consequence, the set of possible bids and asks remained too small, and did not concentrate near the equilibrium. This resulted in instability of the price sequence.

To avoid this problem, the following specification of beliefs and strategy choices is adopted. Beliefs are defined only on the set D , and traders choose asks and bids randomly in an interval around their expected surplus maximizing ask or bid. This randomization of bids and asks allows for exploration of possible prices and focusing on the equilibrium price.

The procedure for this randomization is as follows. Assume that the domain of the belief functions $D_n^{(L)}$ is ordered in increasing values. Suppose $d_{k^*} \in D_n^{(L)}$ is the expected surplus maximizing ask for seller i . Let

$$m_{k^*} \equiv \min \{d_{k^*} - d_{k^*-1}, d_{k^*+1} - d_{k^*}\}.$$

Then the randomization will be made on

$$(\max\{d_{k^*} - m_{k^*}, c_i\}, \min\{d_{k^*} + m_{k^*}, oa\}) \cap \tilde{N}$$

according to a uniform distribution. For buyer j , if $d_{k^*} \in D_n^{(L)}$ is the expected surplus maximizing bid, the bid will be made on

$$(\max\{d_{k^*} - m_{k^*}, ob\}, \min\{d_{k^*} + m_{k^*}, v_j\}) \cap \tilde{N}$$

according to a uniform distribution.

IV Simulations

Much of the evidence presented on performance of the model is from simulations. This is because the dynamics of the model are complex, making analytic characterization of its properties difficult to obtain. To determine the distribution of the next ask, bid or acceptance value at a given time, it would first be necessary to determine the distribution that determines which seller or buyer is next to submit an ask or bid. This distribution depends on the cost and valuation schedules and on the traders' beliefs at that time. Whenever a trade occurs, the cost schedule for one of the sellers changes, and the valuation schedule for one of the buyers changes. Hence the distribution of the next seller or buyer to submit an ask or bid changes. Moreover, whenever an ask or bid is submitted or accepted, that may change the belief functions for the traders, which will change both the distribution of the next trader to submit a message, and the optimal ask for the trader who does send a message.

Considerations of this type make analysis of the distributions difficult to obtain. It may be possible to obtain results on asymptotic convergence to an approximate equilibrium, but asymptotic convergence alone does not provide information about the path. For the important question of the path of convergence to equilibrium, these difficulties suggest that simulations are a useful tool for investigating properties of this model. Comparison is made between outcomes of laboratory trading and 100 simulations of the model of Section III. In figure 3, laboratory market 3pda01 is shown. A simulation of the model of Section III is shown in figure 4. A simulation of the ZI model³ is shown in figure 6. Figure 7 shows a simulation of the model of Section III in a symmetric market with a shift in the supply and demand schedules after five periods of trading. In addition to these graphs of simulations, a summary of convergence and efficiency statistics for the lab markets, for the 100 simulations of the model of Section III,⁴ and for the 100 simulations of the ZI model is provided in table 1.

A Symmetric market design

Several experiments with the symmetric market design were run in laboratory experiments by Arlington Williams and Vernon Smith. Examples of outcomes from six markets with this design are published in Ketcham, Smith and Williams [1984]. In addition, market 3pda01 in figure 3 was also run with this design. These experiments consisted of nine identical trading periods, each lasting 300 seconds. The transaction price sequences in laboratory markets under these or similar conditions typically exhibit early volatility, and then settle down to transaction prices close to competitive equilibrium prices. On the left of figure 3 is a graph of the supply and demand conditions for the symmetric market 3pda01. On the right side of that figure is a graph of the sequence of transaction prices through the 9 periods of trading. The equilibrium price is shown as a solid line across the diagram. Each trading period is separated by a vertical line, and the number of transactions per period is indicated at the

³When the memory length in the model of Section III is set to $L = 0$, that model reduces to the ZI model. Simulations of the ZI model were run using programs for the model of Section III with $L = 0$.

⁴The simulations of this model reported here were run with $L = 5$. Simulations with other values of L have been run and examined. Model performance is not substantially different for values of L between 3 and 8.




Figure 3: Supply and demand conditions (left) and transaction prices (right) for market experiment 3pda01.




Figure 4: Supply and demand conditions and transaction prices for a simulation of market 3pda01.

bottom of the diagram below the vertical line that indicates the end of the trading period. At the equilibrium price $p^* = \$2.35$ in this market, the equilibrium quantity of trade is $q^* \in \{5, 6, 7\}$. A simulation of our model under the same supply and demand conditions is shown in figure 4. Price sequences from lab experiments with this design (as in figure 3) and from simulations with this market design (as in figure 4) both converge quickly to prices near the competitive equilibrium price and an equilibrium quantity of trade typically occurs in each period.

Table 1 summarizes important statistical features of the seven lab markets, of 100 simulations of the model of Section III, and of 100 simulations of the ZI model. In this table, the mean absolute deviation of transaction price from equilibrium price is shown for the first two periods of trading, for the entire 9 periods of trading, and for the last two periods of trading. This statistic is evaluated for each of the three data sets described above. Market

efficiency is likewise examined at the beginning of trading, throughout trading, and at the end of trading for each of these data sets. While the ZI traders attain high efficiency (in markets with a single commodity), that model does not result in the formation of equilibrium prices. The data on mean absolute deviation in table 1 show not only that the model of Section III converges to within a few cents of the equilibrium price, but that the rate of convergence is similar to that found in laboratory experiments.

	Symmetric Markets	Model Simulations	ZI Model Simulations
Mean Absolute Deviation (First two periods)	0.101	0.138	0.276
Mean Absolute Deviation (Entire experiment)	0.050	0.052	0.237
Mean Absolute Deviation (Last two periods)	0.022	0.031	0.209
Market efficiency (First two periods)	0.907	0.979	0.968
Market efficiency (Entire experiment)	0.959	0.988	0.968
Market efficiency (Last two periods)	0.970	0.999	0.967

Table 1: Summary of statistics from simulations of models and from lab data.

The stability of prices in simulations can be explained in terms of the sellers' and buyers' belief functions in figure 5. The belief functions $\hat{p}(a)$ and $\hat{q}(b)$ shown in that figure are from the the end of the fifth period of the simulation of figure 4 using definitions 8 and 9 of Section III.C.2. In this graph, a seller's belief that ask a will be accepted by a buyer is shown for each ask from \$2.32 to \$2.38; buyers' beliefs are shown for bids from \$2.32 to \$2.38.⁵ These belief functions are monotonic (see propositions 1 and 2), so the value of the sellers' belief is $\hat{p}(a) = 1$ for all $a < \$2.32$ and it is 0 for all $a > \$2.38$. With this belief function, and with myopic surplus maximization, the optimal ask is \$2.34 for any seller with unit cost of \$2.30 or less. The buyers' belief functions in this case have a similar property: the optimal bid for a buyer with valuation of \$2.40 or greater is \$2.37. At the beginning of the sixth period, the first action will be either an ask of \$2.34 or a bid at \$2.37. Suppose that the first action is an ask at \$2.34. As a result of the spread reduction rule, sellers' asks must

⁵Note that the range from the lowest cost to the highest valuation in this market is \$1.40 to \$3.30, with an equilibrium price of \$2.35; beliefs are focused in a narrow range around the equilibrium price.

be less than \$2.34. An ask of \$2.33 would result in expected surplus $(\$2.33 - c) \cdot 1$ for any seller with cost less than or equal to \$2.30. Since the distribution of costs and valuations at the beginning of each period in this market is symmetric, and since the probability of each trader being the next to send a message is equal to that traders' proportion of total surplus (see proposition 1), the probability that a seller will send the next message is 0.50, so that the probability of two consecutive asks is approximately 0.25. In general, the distribution of absolute deviations from equilibrium is approximately the distribution of run lengths of a sequence of Bernoulli trials, since a low price results from a sequence of asks and a high price results from a sequence of bids.




Figure 5: Sellers' beliefs (left) and buyers' beliefs (right) after five periods of trading in simulation of market 3pda01.

Figure 6 shows a simulation of the ZI model in the symmetric market environment. In the ZI model, there is no belief formation process. As a result there is no movement of transaction price to the equilibrium. In table 1, the mean absolute deviation of transaction price from equilibrium price is shown for 100 simulations of the ZI model in the symmetric market design.⁶ This statistic is also shown for 100 simulations of the model of Section III and for the seven lab markets. These data (and the graph in figure 6) show that the ZI model does not result in convergence to competitive equilibrium, but that the behavior in the model of Section III does facilitate convergence to approximate equilibrium prices after several periods of trading.

B Market with shifting conditions

The diagram on the left of figure 7 shows two sets of supply and demand conditions. The lower set – shown with thinner lines labeled S and D – is identical to the supply and demand conditions in figures 2, 3, 4, and 6. If after one period or several periods of trading, buyers have each valuation increased by \$0.50 and sellers have the cost of each unit increased

⁶Although the mean absolute deviation in the last two periods of the ZI model simulations is less than in the first two periods, this is not the result of convergence. The price sequence in each period constitutes a draw from the same distribution.




Figure 6: Supply and demand conditions and transaction prices for a simulation of market 3pda01 using the ZI model.




Figure 7: Supply and demand conditions and transaction prices for a simulation of a market with a supply and demand shift.

by \$0.50, then the new supply and demand are those shown with heavier lines labeled S' and D' . The equilibrium quantity of trade and the total surplus are unaffected by this shift, but the equilibrium price increases from \$2.35 to \$2.85. Since expectations focus near the original equilibrium after several periods of trading (see figure 5), the dynamics of movement to the new equilibrium can be examined by considering this type of market. The sequence of transaction prices from a simulation of the model of Section III in this type of market – with the shift occurring after 5 periods of trading – is shown in figure 7. From periods 6 through 10 in this market, supply is S' and demand is D' : the equilibrium price is then \$2.85. In the process of adjusting to the new equilibrium, several trades typically occur in the range \$2.35 and \$2.85 before prices become established near the new equilibrium. In the simulation shown, convergence to the original equilibrium occurs by the end of period 2. Beginning in period 6, the equilibrium price shifts up \$0.50. By the end of

period 7 transaction prices establish near the new equilibrium price. This simulation shows that in the model developed here, traders respond to shifting market parameters, and prices quickly adjust to a new equilibrium.

V Conclusions

The simulations in Section IV demonstrate that the model in Section III results in competitive equilibrium prices after several periods of trading, and that behavior is responsive to altered market conditions. The data from simulations summarized in table 1 demonstrate that the model of Section III results in market efficiency of 99.9% in the symmetric market environment after several periods of trading *and* in convergence to approximate equilibrium prices. The ZI model, which demonstrates that random behavior results in market efficiency, does not result in the formation of equilibrium prices. Markets have the capacity to resolve complex coordination problems. In markets with multiple goods and production, optimal choices of input levels depend on expectations of relative prices of inputs, and on expectations of output prices. In a general equilibrium environment, consumers' optimal consumption choices depend on relative prices. A model of firm and individual behavior in exchange and production environments must therefore capture the formation of price expectations and of equilibrium prices.

Several market experiments have tested the competitive model in more complex environments. Gjerstad, Williams, Smith, and Ledyard [1995] have examined the formation of equilibrium prices in a market with two commodities, A and B . In their design there is a set of m buyers and a set of n sellers. Sellers have independent marginal cost arrays in both market A and market B , with costs induced as in Section III.A. Buyers have induced quasi-concave preferences over commodities A and B . For example, if the experimenter wants a buyer to have the utility function $u_i(A_i, B_i) = (a \cdot A_i^r + b \cdot B_i^r)^{1/r}$, this is accomplished by giving the buyer the same monetary reward for any final commodity holding lying on an indifference curve $u_i(A_i, B_i) = c$. Gjerstad, Williams, Smith, and Ledyard find that in these environments the bargaining process quickly leads to the formation of equilibrium prices. Since demand in each market depends on the prevailing price in the other market, and since the price in each market must form through the bargaining process, the problem that consumers solve in this environment is complex. It is clear that any model of the formation of price expectations and of bargaining in this environment must be based on a model that results in the formation of equilibrium prices in a market with a single commodity. Lian and Plott [1993] is an example of a complex market experiment involving production and consumption. In their market, consumers have a labor endowment which is sold in a labor market to producers. Producers have concave production functions which transform labor into a consumption good. This consumption good is then sold to consumers who have quasi-concave preferences over leisure and the consumption good. In both these market designs, consumers and producers will be able to make optimal decisions if they are able to accurately forecast prices in the market for each commodity. Extensions of the model of Section III to these environments would significantly enhance understanding of the process of market equilibration.

References

- [1] Chatterjee, K. and L. Samuelson. "Infinite Horizon Bargaining with Alternating Offers and Two-Sided Incomplete Information," Technical Report, Pennsylvania State University, 1984.
- [2] Easley, David, and John Ledyard. "Theories of Price Formation and Exchange in Double Oral Auctions," in *The Double Auction Market: Institutions, Theories, and Evidence*, Daniel Friedman and John Rust, eds., Addison-Wesley, 1993.
- [3] Friedman, Daniel. "A Simple Testable Model of Price Formation in the Double Auction Market," *Journal of Economic Behavior and Organization*, 1991.
- [4] Gjerstad, Steven. "Price Formation in Double Auctions," Ph.D. thesis, University of Minnesota, January 1995.
- [5] Gjerstad, Steven, Arlington Williams, Vernon Smith, and John Ledyard. "Simultaneous Trading in Two Competitive Markets: An Experimental Examination," mimeo, University of Arizona, October 1995.
- [6] Gode, Dhananjay, and S. Sunder. "Allocative Efficiency of Markets with Zero Intelligence Traders: Market as a Partial Substitute for Individual Rationality," *Journal of Political Economy*, vol. 101, 1993.
- [7] Hahn, Frank. "Auctioneer," in *The New Palgrave: A Dictionary of Economics*, John Eatwell, Murray Milgate, and Peter Newman, eds., MacMillan Press, 1987.
- [8] Hayek, F.A. "The Use of Knowledge in Society," *American Economic Review*, vol. 35, 1945.
- [9] Hurwicz, Leonid. "On Informationally Decentralized Systems," in *Decision and Organization*, C.B. McGuire and Roy Radner, eds., University of Minnesota Press, 1972.
- [10] Ketcham, Jon, Vernon L. Smith, and Arlington W. Williams. "A Comparison of Posted-Offer and Double Auction Pricing Institutions," *Review of Economic Studies*, vol. LI, 1984.
- [11] Ledyard, John. "The Scope of the Hypothesis of Bayesian Equilibrium," *Journal of Economic Theory*, vol. 39, 1986.
- [12] Lian, Peng, and Charles R. Plott. "General Equilibrium, Macroeconomics, and Money in a Laboratory Experimental Environment," Social Science Working Paper 842, California Institute of Technology, March 1993.
- [13] Plott, Charles. "Industrial Organization Theory and Experimental Economics," *Journal of Economic Literature*, vol. XX, December 1982.
- [14] Rustichini, A., M. Satterthwaite, and S. Williams. "Convergence to Efficiency in a Simple Market with Incomplete Information," *Econometrica*, vol. 62, 1994.

- [15] Smith, Vernon L. "An Experimental Study of Competitive Market Behavior," *Journal of Political Economy*, vol. **LXX**, 1962.
- [16] Smith, Vernon L. "Microeconomic Systems as an Experimental Science," *American Economic Review*, vol. **72**, 1982.
- [17] Wilson, Robert B. "On Equilibria of Bid-Ask Markets," in *Arrow and the Ascent of Modern Economic Theory*, G.W. Feiwel, ed., New York University Press, 1987.