

David Easley[†] and John O. Ledyard[‡]

[†]Cornell University, Ithaca, NY 14853, and [‡]California Institute of Technology, Pasadena, CA 91125

THREE

Theories of Price Formation and Exchange in Double Oral Auctions

We provide a theory to explain the data generated by experiments with double oral auctions. Our theory predicts convergence to the equilibrium implied by the law of demand and supply and provides an explanation of disequilibrium behavior. The predictions of our theory seem to fit the data better than do the predictions of Walrasian, Marshallian, or game theoretic models. Our theory also suggests that, in demand-supply environments, the double oral auction is remarkably robust in the sense that aggregate performance is similar for a very wide range of individual behaviors.

1. INTRODUCTION

One of the main justifications for the use of equilibrium models in economics is the argument that there are forces which tend to drive agents and their decisions towards an equilibrium if they are not at one already. Market equilibrium models have proven to be extremely powerful in the analysis of many situations; however, attempts to model and explain the forces that do drive an economy to equilibrium

have met with little success. Most of the literature on the stability of equilibrium uses the fiction of a disinterested auctioneer who adjusts a single known price for each good in response to stated excess demand resulting from agents' equilibrium plans. The limitations and defects of this approach are well known: for a survey of the literature, see Arrow and Hahn.¹ In addition, as far as we know, the *only* institutional arrangement that even approximates this idealized model of price formation is the London gold market (see Jarecki⁹).

Now, however, a body of data has been generated which provides detailed information on the disequilibrium behavior of traders in auction markets similar to those of organized commodity or stock exchanges. These data are difficult to ignore since they are generated experimentally under controlled conditions, and cannot be explained away by reference to measurement error, unobserved variables, or other fudge factors. In the experiment, a small number of traders, each with limited imperfect information, determine prices and quantities transacted through interactive bargains. There is neither a single price nor a single price quoter. Nonetheless, the quantities exchanged and the prices at which transactions take place typically converge to, or near to, the values predicted by the law of demand and supply. But, in spite of the fact that the traditional demand-supply model appears to yield reasonably accurate predictions of the long-run average prices and quantities in these markets, it fails to yield any insights into the process by which these prices and quantities are obtained.

In this paper we consider several positive theories of the price formation and exchange process for the class of experimental exchange markets called Double Oral Auctions.^[1] We examine three of these theories in detail and argue that one of them seems to be the most consistent with the data. The ability of this theory also to explain price formation and exchange in other markets such as the New York Stock Exchange depends, of course, on the degree of parallelism that exists between the two (see Smith²⁰). An astronomer's maintained hypothesis is that the physics of the lab is the same as that of the sun; our working hypothesis is that behavior in experimental markets is similar to that in other markets, and that insights discovered in the evidence generated in the lab are potentially transferable to nonexperimental markets with similar institutional structures. Thus, we view the theory in this paper as a first step towards constructing a positive theory of the process of exchange and price formation in many other markets.

2. THE EXPERIMENTAL MARKET

In a double oral auction (DOA) experiment, a pool of subjects (usually eight to twelve) is divided at random into a group of buyers and a group of sellers. The

^[1]In fact, some of the auctions are computerized rather than oral. All that matters is that participants can make bids or offers and acceptances, and are informed of others' bids or offers and acceptances.

buyers are given value schedules telling them the amount in cents that they will receive from the experimenter for each unit of the good they purchase. Buyers keep the difference between their value and the price they pay for that unit. The sellers are given cost schedules telling them their cost in cents for each unit of the good they sell. Sellers keep the difference between their selling price and their cost on each unit they sell. Each subject knows his own payoff schedule but is given no information about the others' payoffs. Smith²⁰ shows how these payoff schedules induce demand and supply schedules. An example of payoff schedules and the induced supply and demand schedules for one experiment is provided in Appendix A.

After they receive their payoff schedules, subjects are allowed to trade during a market period of some fixed length, usually called a market day. Buyers can make bids to buy a unit of the good and sellers can make offers to sell a unit. If a bid or offer is accepted, a binding trade occurs and all traders are informed of the contract price. Once a trade is completed, bids and offers can be made for another unit of the good. No information other than bids, offers, acceptances, and contract prices is transmitted or known by the participants.

When a market period ends, the subjects are given new payoff schedules, identical to their schedules for the previous period, and the experiment is repeated.^[2] Market demand and supply conditions are typically held constant across periods so that any equilibrating process that exists has a chance to establish an equilibrium. For a more detailed explanation of auction experiments and the usual results, see Williams²² and Smith and Williams.²¹

These experiments provide a unique opportunity to examine price formation for two reasons. The first is that, unlike nonexperimental markets, the actual prices and quantities predicted by the law of demand and supply are known to the experimenter. Secondly, complete data on bids, offers, contracts, and their timing is available. An example of a typical design and the data generated is provided in Appendix A. Demand and supply functions can be calculated from the subjects' valuations, and equilibrium prices and quantities can then be computed. The first obvious fact from these experiments is that *actual exchange prices are not equal to those predicted by the law of demand and supply*. In a strict sense, demand-supply theory is rejected by these data. The second obvious fact, however, is that after a very few replications, *transaction prices and quantities converge to near those predicted by the law of demand and supply*. These observations have been replicated many times. The only conclusion one can draw is that the traditional theory needs refining before one has a compelling explanation of the observed behavior in these markets. Not only must "equilibrium" be explained, but we must also explain the "disequilibrium" transactions, the sequence in which they occur, and the process by which participants are "learning."

In our search for a better theory of price formation, we have used several criteria. Firstly, we wanted the theory to predict convergence to the predictions of the law of demand and supply for those experiments in which convergence occurs and to predict nonconvergence in those experiments in which convergence does not occur.

[2] Other designs are also used; see Smith¹⁹ and Smith and Williams²¹ for some of these.

Secondly, we wanted the theory to be useful in understanding the dynamics of the adjustment process by making falsifiable predictions about the entire process. A theory which predicts eventual convergence at $T = \infty$, and nothing else, is consistent with the data but not very illuminating. Third, we wanted the predictions of disequilibrium behavior not to be at odds with the data. How one weighs these criteria against one's prior belief in any particular theory is a matter of judgement. Our choice will be evident from the theories we reject and the candidate we offer.

3. THREE POSSIBLE THEORIES

Our goal is to understand how the actual dynamics of these markets work, not how they should work. We recognize that there are a variety of models which purport to explain price adjustments, but we view the existence of experimental data as an opportunity to reject a subset of those theories which seem obviously inappropriate. The set of reasonable theories for these markets can now be constrained by the data in a way that has been unusual for economics. To see what this means, let us consider three candidates for a theory of market dynamics.

Since both the institutional description and the data from the experimental DOA markets reject the Walrasian tâtonnement auctioneer as the appropriate model of price formation, a natural alternative might be a Marshallian theory. In a naive version of this theory, the trading sequence depends on the differences in buyers' values (willingness to pay) and sellers' values. In particular, this theory predicts that trade will occur in the efficient order, i.e., the first trade will occur between the buyer with the highest induced value (Buyer 1 in the example in Appendix A) and the seller with the lowest induced cost (Seller 1 in the example in Appendix A). The second trade is predicted to occur between the buyer and seller with the second values and costs, and so on. This theory does not predict which prices will occur, but it does predict that the total quantity transacted will be the competitive equilibrium quantity. Unfortunately, this theory has little to do with reality. When we look closely at the microdata, we see that the theory is soundly rejected. A cursory glance at the summary data of Appendix A should convince even the most skeptical reader that the predictions of the naive Marshallian theory are not at all consistent with the data. (In IPDA14, the rank correlation coefficient between the order of the true values and the order of the transactions is .369 in week 1 and .273 in week 2.) This is an excellent example of a case in which the experimental setup allows us to test more hypotheses than would be possible if we only had access to nonexperimental market data. Testing the prediction concerning the order in which participants are involved in transactions would be impossible without explicit knowledge of the individual valuations.

A second candidate for a theory is found in Friedman³ who takes an alternative approach to the problem by redefining the experiment. He studies one day of a DOA with traders who are allowed to resell or repurchase the good being traded. Under a no-congestion condition which requires that at the day's end no trader wants to

reset the closing bid or ask prices or accept the outstanding bid or ask, he shows that the final allocation will be at most one transaction away from being Pareto-optimal. No congestion implies that the final ask be no more than the second-lowest cost of selling a unit, the final bid be no less than the second-highest value of buying a unit, and that no one wants to accept the final prices. With resale and repurchase allowed, this insures that all but perhaps one infra-marginal unit has traded and that no more than one extra-marginal unit has traded. Beyond the question of the appropriateness of the no-congestion assumption, the difficulty in applying this theory to the DOA experiments is that the theoretical conclusion relies heavily on the agents' ability to retrade, while retrading is not allowed in many of the experiments. The theory also finesses the issues of learning and dynamics. How no congestion occurs is left unexplained.^[3]

A third candidate for a theory would be a model based on game-theoretic considerations. For most of the DOA experiments, there is a complete-information Nash equilibrium (with price-quantity offers or bids as strategies) in which all trades take place at the competitive equilibrium price. However, the use of a Nash equilibrium concept to describe the experimental market has two difficulties. Firstly, the data are not consistent with this equilibrium (not all trades occur at the competitive price). Secondly, the participants in the experiments do not have enough information to calculate the strategies required to support this equilibrium. (They would have to be able to calculate the competitive equilibrium price.) Thus, one must turn to models with asymmetric information.

In the experiments which have been run, details on others' payoffs (and thus on the competitive equilibrium) may only be inferred by the subjects from the public data on bids, offers, and contracts. Thus, the structure in which subjects find themselves is a dynamic game with incomplete information. If an equilibrium were calculated for this game, its predictions could then be compared with the data. We feel that there are at least three difficulties with using this approach to construct a positive theory of double oral auctions. Firstly, as common knowledge about the distribution of valuations and the strategies selected are not controlled in the experiments, it is not clear how to apply game theory, as it currently exists, to the experiments. These are games of incomplete information: they are not games of imperfect information.^[4] One could try to ignore this problem and assume that there is, at some level, common knowledge. However, this leads to the second difficulty.

^[3]In an important recent paper, Friedman⁴ has filled in this gap with a model based on search-theoretic principles. We discuss this interesting model in more detail below in Section 4.D.

^[4]In his seminal articles, Harsanyi⁷ was very careful to differentiate between incomplete and imperfect information. Beginning with a game of incomplete information, he converted it to a game of imperfect information *from player i's point of view*. There was no guarantee, absent an assumption of objective common knowledge, that the game from *i's* point of view would be the same as the game from *j's* point of view. Therefore, without the common knowledge hypothesis, players can be surprised *on the equilibrium path*: they discover that they are in an entirely different game tree than they thought they were. At this point Bayes' rule provides no guidance and players can do anything. With this freedom, one can make any outcome of the experiment a Bayes-Nash (cont'd.)

With an assumption of common knowledge, the natural model is the Bayes-Nash equilibrium. If the subjects are risk neutral, we know from Gresik and Satterthwaite⁵ and the revelation principle that any Bayes-Nash equilibrium has the property that no extra-marginal units are traded when subjects only own one unit of the commodity. Yet in the experiments, extra-marginal units are often traded (see, e.g., the data from IPDA14 in Appendix A). If the subjects are risk averse, then we know from Ledyard¹⁰ that virtually anything can be an equilibrium. If risk attitudes are not controlled for (see Roth and Malouf¹⁷), then the game-theoretic model explains everything.

Our third difficulty with the game-theoretic approach is that, as far as we know, no one has solved for an equilibrium of the appropriate game. Wilson²³ has found strategies for a one-shot version of the DOA which satisfy the necessary conditions for a Bayes-Nash equilibrium. However, Wilson's model predicts that the rank correlation coefficient between the order of true values and trades is one which, as we noted above, is strongly at odds with the data. But the experimental DOAs are repeated, common knowledge is not controlled for and subjects may not be risk neutral as Wilson assumes. Under these circumstances, it is not fair to compare Wilson's predictions with the data. It is also unfair to expect much from the general approach.

Since neither the Marshallian model nor the Friedman model, nor any currently available game-theoretic model appears to be appropriate as a positive theory, and since the goal of building an appropriate game-theoretic model has eluded us and others, development of an alternative model seems warranted. We turn to that next.

4. A POSITIVE THEORY

A. PRELIMINARIES

A participant in a double oral auction experiment has a complex decision problem. He must decide when to bid, how much to bid, and whether or not to accept the trades offered by other subjects. Further, all of these decisions must be made with very imperfect information. The subject does not know the payoffs or expectations of other agents, he does not know the terms of trade that will be available to him in the future, and he does not know the effect of his actions on the actions of others. This is a very complex interactive decision problem with incomplete information in which individuals must choose bidding and acceptance strategies. To place some structure on this problem, we first introduce some notation and definitions concerning the data known to both the experimenter and us.

The *true payoffs or values given to buyers* are integers and are ranked as $V^1 \geq V^2 \geq \dots \geq V^n \geq 0$, where V^i is the i th highest value and there are n units. A buyer

equilibrium of the *incomplete* information game without objective common knowledge. We show how to do this in footnote 12.

b will be assigned a subset of these units $V^{b1} \geq V^{b2} \geq \dots \geq V^{bB}$ and will trade them one at a time in the sequence $b1, b2, \dots, bB$. No recontracting is allowed. The *true costs given to sellers* are integers and are ranked as $0 \leq M^1 \leq M^2 \leq \dots \leq M^m$, where M^j is the cost of the j th unit and there are m units. A seller will be assigned a subset of these units and will trade them one at a time. No recontracting is allowed. It should be noted that typically the values, V^i , and costs, M^i , are assigned once and remain fixed. Each buyer (seller) knows only his own values (costs) and no participant is given any information as to how these values and costs were chosen. There is no basis for common knowledge assumptions about independence of values or their distributions. Consequently, we neither make such assumptions nor use these concepts in our theory.

Market period or days for an experiment are indexed by $d = 1, 2, \dots$. The *time remaining in any given day* is indexed by $t = 0, 1, \dots, T$. Contract prices, bids, and offers are in integer units in the interval $[0, H]$, where $H < \infty$ is some arbitrarily selected upper bound above V^1 and M^m , and during any particular day, d , each participant observes all contract prices, bids, and offers.^[5]

To summarize, each buyer knows the rules of the auction, the value of his own units, and the sequence, timing, amount, and identity of all past bids, offers, and contracts. It is these data alone on which the buyer can base his decisions to bid and to accept. A symmetric remark applies to each seller.

B. AN INTUITIVE LOOK

We adopt the spirit of both revealed preference theory and demand-supply analysis by placing assumptions on individual behavior which, we believe, are consistent not only with optimal behavior but also with a vast range of "boundedly rational" rules-of-thumb. We do not model how agents should make their decisions. Instead, we provide criteria which, we believe, sensible individuals in these markets act as if they satisfy. This allows us to construct a theory which is robust to a wide variety of individual behaviors and yet which is reasonably sharp in its predictions about the data. We model "reduced form" behavior by decomposing the decision problem into three main elements: expectations, reservation prices, and bidding strategies. These are most easily explained in reverse order.

Assume that at each instant of time, there is for each buyer (seller) a *reservation price*, possibly different from his true value, which summarizes his willingness to bid up (offer down) to that price or to accept any offer up (bid down) to that

^[5] Each participant also observes the timing of each contract, bid, and offer. It is highly probable that the timing of these events is an important piece of information which affects the actions of the buyers and sellers. However, the level of complexity required to incorporate timing into the model seems to outweigh the gains to be achieved. Thus, we ignore it throughout the paper.

price.^[6] If each participant has such a reservation price as a function of time, the buyers' side of the double auction can then be thought of as proceeding like an ascending bid (English) auction, *with these reservation prices substituting for the true values*. After some period of time, the outstanding bid will always be held by the buyer with the highest reservation price (not necessarily the highest untraded value), and that bid will be at least as high as the second-highest reservation price. Otherwise, the holder of the second-highest reservation price will bid, causing the holder of the highest reservation price to rebid, and so on. We find it unnecessary to explicitly model this process, and we assume that it occurs instantaneously. Thus, all observed bids will be the reduced form results of the above English auction. Since we have also assumed that the buyer is willing to accept any price lower than the reservation price, an acceptance of an offer will occur whenever that offer is lower than the highest reservation price of a buyer. To make the theory as simple as possible, we assume the buyer with the highest reservation price moves instantaneously faster than any other buyer. (This is only a restriction if offers jump down in large discrete increments and several buyers have similar reservation prices—a situation likely to occur only in the opening minutes of any trading day.) This intuitive view of the bidding is formalized in Assumption 1 below. Sellers' offers are viewed symmetrically in Assumption 1'.

Since bids and offers depend on reservation prices and not directly on the induced values, bids and offers ultimately depend on the relationship between reservation prices and the data observed by each agent. This relationship is assumed to depend on two principles of learning. Firstly, it is true under Assumption 1 that whenever the bids and acceptance prices of a buyer are higher than were necessary to complete a transaction, the buyer completes a trade but overpays. We assume that a buyer will realize that he overpaid and will, during the next auction, lower his reservation price. If it is not lowered too much, the buyer should still be able to complete a transaction but at a better price. Secondly, it is true that if a buyer waits too long to bid or, what is the same thing, maintains too low a reservation price during the day, then that buyer may not complete a transaction even though profitable ones are available. We assume that if a buyer could have purchased a unit at less than its value to him, V^i , but did not, then that buyer will realize he underbid and will, either that day or during the next auction, raise his reservation price at each time of day. It is the delicate balance between "paying too much" and "not offering to pay enough" which the buyers must learn in order to be successful in the auction. We do not explicitly model this learning process; instead, we provide assumptions about reservation price behavior which, if satisfied, reflect these learning principles. We summarize this rather simple intuition in Assumption 2 below.

[6] A more sophisticated theory might distinguish between the amount a buyer is willing to bid and the lowest offer he would accept. In particular, buyers may not be willing to bid up to their reservation price (see Wilson²³). This distinction could be easily incorporated into our model, but it is not apparent that it would add to the explanatory power of the model.

C. BIDDING BEHAVIOR

We start our description of the formal theory with the introduction of a hypothesis concerning the existence of the key unobservable of our model. It is important to realize that we treat reservation prices in this paper in the way that preferences are generally treated in economics. We cannot observe whether subjects really compute reservation prices; we can only assume they act as if they do. For a coherent theory, the reservation prices may need to be related in a systematic way to the true values but, *a priori*, do not need to be.

ASSUMPTION 0: RESERVATION PRICES. For each buyer unit and seller unit, there is an (unobservable) reservation price at each day d and time t , denoted $r_d^i(t) \in R^1$ for buyers and $s_d^j(t) \in R^1$ for sellers.

Assumption 0 only contains notation. To link the unobservables to the data, we need to tie the bids and acceptances to the reservation prices, and then to tie the reservation prices to the true values and costs. As we indicated in the previous section, this is done by assuming that, given reservation prices, bids and acceptances are the reduced form of English auction behavior.

ASSUMPTION 1: BUYERS' BIDS AND ACCEPTANCES.

- i. $b_d(t)$, the current outstanding bid in day d , with time t left, is held by buyer i^* where $r_d^{i^*}(t) \geq r_d^i(t)$, for all $i = 1, \dots, n$.
- ii. $b_d(t) \leq r_d^{i^*}(t)$.
- iii. $b_d(t) \geq r_d^i(t)$, for all $i \neq i^*$.
- iv. Buyer i^* accepts the current outstanding offer, $o_d(t)$, if and only if $o_d(t) \leq r_d^{i^*}(t)$. No other i accepts $o_d(t)$.

Simply stated, at each point in time, the current bid is held by the buyer with the highest reservation price—*not necessarily the buyer with the highest true value*. This bid lies below that reservation price and above the second-highest reservation price. Under Assumption 1, and 1' below, trades always occur between the buyer with the highest reservation price and the seller with the lowest reservation price. We emphasize that *these need not be the buyer with the highest value and the seller with the lowest cost* since the English auction is based on reservation prices and not on the "true values," V^i and M^i .

For completeness, we make an assumption on the offers and acceptances of sellers that is symmetric with that made for buyers. The only difference is that we have arbitrarily assumed that if seller j^* is willing to accept $b_d(t)$ and buyer i^* willing to accept $o_d(t)$, then the buyer accepts first. We could reverse this without affecting the conclusions to come.

ASSUMPTION 1: SELLERS' OFFERS AND ACCEPTANCES.

- i. $o_d(t)$, the current outstanding offer in day d , with time t left, is held by seller j^* where $s_d^{j^*}(t) \leq s_d^j(t)$, for all $j = 1, \dots, m$.
- ii. $o_d(t) \geq s_d^{j^*}(t)$.
- iii. $o_d(t) \leq s_d^j(t)$, for all $j \neq j^*$.
- iv. Seller j^* accepts the outstanding bid, $b_d(t)$, if and only if $b_d(t) \geq s_d^{j^*}(t)$ and buyer i^* does not accept $o_d(t)$. No other j accepts $b_d(t)$.

We do not yet have a testable theory since, given any sequence of bids and contracts, it is possible to construct a sequence of reservation prices which, under Assumption 1, would imply the given data precisely. Unless we place some restrictions on the reservation prices, we can explain anything, and therefore nothing.

D. RESERVATION PRICE FORMATION

We now tie the theory down by restricting reservation price behavior in a way which relates it to observable data. This is the way in which we connect bids, contract prices, and the sequence of trades to the initial data known by the experimenter and, thus, provide testable propositions about these auctions.

Reservation prices are assumed to be formed in accordance with the intuitive principles outlined in Section 4.B. We begin by assuming that a buyer's expectations in any period are based on the prices of the previous period. In particular, we assume that the support of the buyer's expectations is the set of prices bounded by the maximum of last period's highest contract price or highest bid, and the minimum of last period's lowest contract price or lowest offer. Based on these expectations, reservation prices are formed over time as follows: (a) for most of a trading day, one's reservation price lies below the true value, V^i , and within the support of the expectations (when this is feasible), (b) if possible, the reservation price is actually below the maximum price in the support since the buyer does not want to "overpay," (c) eventually, if no contract is agreed to, buyers will cave in and let the reservation price rise and approach the maximum price in the support, and (d) if still no contract is completed, the reservation price will rise higher than even the maximum in the support of the expectations.

The sequence of actions (a), (b), and (c) are consistent with optimal behavior in finite-time, nonstrategic search models.^[7] If a buyer believes that offers are identically and independently distributed on $[\underline{P}, \bar{P}]$, no matter what he does, and that he will receive a finite number of draws of offers with replacement, it is really

^[7] For examples of this literature see Gronau,⁶ Lippman and McCall,^{11,12} Mortensen,¹³ and Cox and Oaxaca.²

easy to show that his reservation price satisfies (a), (b), and (c).^[8] A buyer who is certain that he can complete a trade at \bar{P} will only move his reservation price to \bar{P} at $t = 0$, the end of the day. If he could not complete a trade at this point, he would presumably be willing to pay more than \bar{P} as he now knows that his beliefs are incorrect. In (d) we assume that he reaches \bar{P} before $t = 0$, perhaps because he is not certain that a trade can be completed at \bar{P} .

Before formalizing our assumption on reservation prices, we need to introduce some notation. If a trade occurs at time t of day d , we let $c_d(t)$ be the contract price. Then for each day $d > 1$, let $\underline{P}_d = \text{Min}\{o_{d-1}(t), c_{d-1}(t) : t = 0, \dots, T\}$ and $\bar{P}_d = \text{Max}\{b_{d-1}(t), c_{d-1}(t) : t = 0, \dots, T\}$. We assign $[\underline{P}_1, \bar{P}_1] = [0, H]$. The interval $[\underline{P}_d, \bar{P}_d]$ is interpreted as the support of traders' price expectations in day d . Let $\Delta P_d = \bar{P}_d - \underline{P}_d$.

ASSUMPTION 2: BUYER'S RESERVATION PRICE FORMATION. For all buyers $i = 1, \dots, n$:

- i. If i has traded (accepted an offer or had a bid accepted) in day d before time t , then $r_d^i(t) = 0$.
- ii. For each day d there is time $\hat{t}_d^i > 1$ such that, if i has not traded in d before t , then:
 - a. For all $t > \hat{t}_d^i$;

$$\text{Min}\{V^i, \bar{P}_d\} > r_d^i(t) \geq \underline{P}_d \text{ if } \Delta P_d > 1 \text{ and } V^i > \underline{P}_d;$$

$$\text{Min}\{V^i, \bar{P}_d\} \geq r_d^i(t) \geq \text{Min}\{V^i, \underline{P}_d\} \text{ otherwise.}$$
 - b. $r_d^i(\hat{t}_d^i) = \text{Min}\{V^i, \bar{P}_d - 1\}$.
 - c. For all $t < \hat{t}_d^i$;

$$r_d^i(t) = \text{Min}\{V^i, b_d(t+1) + 1\} \text{ if}$$

$$b_d(t+1) \in \{\bar{P}_d, \bar{P}_d - 1\} \text{ and } b_d(t+1) \text{ unaccepted;}$$

$$r_d^i(t) \in \{r_d^i(t+1), \text{Min}\{V^i, b_d(t+1) + 1\}\} \text{ if}$$

$$b_d(t+1) > \bar{P}_d \text{ and } b_d(t+1) \text{ unaccepted;}$$

$$r_d^i(t) = r_d^i(t+1) \text{ otherwise.}$$

Assumption 2(i) sets the reservation price for traded units to zero to indicate that they have left the market. The conditions in Assumption 2(ii)(a) embody the intuition that, as a result of learning, reservation prices will not be "too high" early in the trading day. The conditions in Assumption 2(ii)(b,c) embody the intuition that, towards the end of the day, if the buyer has not completed a transaction, then

[8] Recently, Friedman⁴ has provided a model of the double oral auction in which agents are Bayesian and expected utility maximizers who ignore the strategic feedback effects of their own actions. He derives reservation strategies which provides a choice-theoretic underpinning to our model. (See Friedman,⁴ p. 57-58.) The main difference between his model and ours lies in our use of bounds \bar{P} and \underline{P} on the support of possible prices (he assumes priors are positive over all prices $[0, H]$) and property (d) which describes what happens when, as a Bayesian, the buyer is surprised that no offer is in $[\underline{P}, \bar{P}]$. (Friedman's agents are never surprised.)

that buyer will learn to raise his reservation price slowly. Towards the end of the day, reservation prices will not be "too low."^[9]

To complete the model we make a symmetric assumption about sellers' reservation prices which we call Assumption 2'.

ASSUMPTION 2': SELLERS RESERVATION PRICE FORMATION. For all sellers $j = 1, \dots, m$:

- i. If j has traded (accepted a bid or had an offer accepted) in day d before time t , then $s_d^j(t) = \bar{P}$.
- ii. For each day d there is a time $\bar{t}_d^j > 1$ such that, if j has not traded in day d before time t , then:
 - a. For all $t > \bar{t}_d^j$:
 $\text{Max}\{M^j, \bar{P}_d\} \geq s_d^j(t) > \text{Max}\{M^j, \underline{P}_d\}$ if $\Delta P_d > 1$ and $M^j < \bar{P}_d$;
 $\text{Max}\{M^j, \bar{P}_d\} \geq s_d^j(t) \geq \text{Max}\{M^j, \underline{P}_d\}$ otherwise.

^[9]Our assumptions on reservation prices can also be stated with a simple, Markov information structure. The only information used from past days is $(\underline{P}_d, \bar{P}_d)$. The only information used from the current day is the previous bid, offer, and contract price, and an indicator of whether the individual has traded. Let the trade indicator for individual i at time t of day d be $h_d^i(t)$,

$$\text{where } h_d^i(T) = 0 \text{ and } h_d^i(t-1) = h_d^i(t) + \begin{cases} 1 & \text{if } i \text{ trades at } t; \\ 0 & \text{otherwise.} \end{cases}$$

The information from the current day is then $I_d^i(t) = \{b_d(t), o_d(t), c_d(t), h_d^i(t)\}$ for any $t < T$ and $I_d^i(T) = \phi$. Any trader's reservation price evolves according to a transition probability which is time and information dependent and parameterized by the trader's value and $(\underline{P}_d, \bar{P}_d)$. To simplify the notation we drop the indices i and d , and we consider the problem from a typical buyer's point of view. The distribution of the buyer's reservation price at time t , $r(t)$, is given by $R(I(t+1), t, r(t+1))$. Let $A = \{r \in [0, \bar{P}] : \text{Min}(V, \bar{P}_d - 1) \geq r \geq \text{Min}(V, \underline{P}_d) \text{ if } \Delta P_d > 1 \text{ and } \text{Min}(V, \bar{P}_d) \geq r \geq \text{Min}(V, \underline{P}_d) \text{ otherwise}\}$. Then we can write Assumption 2 equivalently as

Assumption 2*: Buyers' Reservation Price Formation.

- (i) If $h(t+1) = 1$, then $R(I(t+1), t, r(t+1))(0) = 1$.
- (ii) If $r(t+1) \in A$, $h(t+1) = 0$, and $b(t+1) < \bar{P}_d - 1$,
or $b(t) = c(t)$, then $\int_A dR(I(t+1), t, r(t+1)) = 1$.
- (iii) For any $r \geq \bar{P}_d - 1$, $R(I(t+1), t, r(t+1))(r) = 0$
if $r > r(t+1) + 1$.
- (iv) If $r(t+1) < V$, $b(t+1) \in [\bar{P}_d - 1, \bar{P}_d]$ and $b(t+1) \neq c(t+1)$,
then $R(I(t+1), t, r(t+1))(r(t+1) + 1) = 1$.
- (v) There is a $t^* > 1$ such that if $h(t^* + 1) = 0$, then

$$\sum_{r \geq \text{Min}(\bar{P}_d - 1, V)} R(I(t^* + 1), t^*, r(t^* + 1))(r) = 1.$$

Either 2 or 2* can be used in the rest of the paper; we use 2.

$$b. \quad s_d^j(\bar{t}_d^j) = \text{Max}\{M^j, \underline{P}_d + 1\}.$$

c. For all $t < \bar{t}_d^j$;

$$s_d^j(t) = \text{Max}\{M^j, o_d(t+1) - 1\} \text{ if} \\ o_d(t+1) \in \{\underline{P}_d, \underline{P}_d + 1\} \text{ and } o_d(t+1) \text{ unaccepted;}$$

$$s_d^j(t) \in \{s_d^j(t+1), \text{Max}\{M^j, o_d(t+1) - 1\}\} \text{ if} \\ o_d(t+1) < \underline{P}_d \text{ and } o_d(t+1) \text{ unaccepted;}$$

$$s_d^j(t) = s_d^j(t+1) \text{ otherwise.}$$

Assumptions 1, 1', 2, and 2' constitute the full set of premises for our theory. One obvious and intended omission is any (direct) tie between reservation prices and valuations—other than the obvious constraint that buyer i 's reservation price be less than i 's value. In particular, we make no assumptions about the relative rankings of values and reservation prices.^[10] While such an assumption might tighten the predictions of the model, most of the theorems we are interested in and most of the implications consistent with the data do not require it. We think this is an attractive feature of the model in that decision making is decentralized. Agents need know nothing about each other but only need to look at observable data to decide what to bid and whether to trade. Convergence to competitive equilibrium under these conditions simply highlights the robustness of the double oral auction as a market institution. Any monotonic link between values and (even randomized) reservation prices would require some coordination between the agents. This in turn would seem to require some prior beliefs on values and a common strategy. This is unnecessary and inconsistent with the spirit of our analysis.

We have made two implicit assumptions which should be recognized. First of all, we assume that each buyer's and seller's behavior is independent of the total number of participants in the market. That is, a buyer's choices of bids and acceptances are the same whether he is a monopolist or one of 100 buyers. Although this runs counter to conventional economics, experimental evidence suggests that if the number of buyers and the number of sellers are both greater than two, then this assumption is satisfied. Further, even if there is a single seller, what little evidence there is suggests that the model we propose may still be appropriate. We leave as an open empirical question just how few participants, if any, are needed before our theory is not applicable.

The second implicit assumption is that buyers and sellers with multiple units to purchase or sell will decide on strategies for each unit separately. That is, the bids and acceptances a buyer makes for his, say, highest valued (first) unit are assumed to be independent of the total number of units he may want to buy. This is not "rational behavior," but the interaction effects are difficult to model (we know of

[10] An example of such an assumption is $[V^i > V^j] \Rightarrow [r_d^i(t) \geq r_d^j(t)] \forall d, t, i, j$. This particular hypothesis, which is closely related to the hypotheses of the Marshallian model and game-theory models, yields predictions seriously at odds with the data. See Section 7 for other possible connections between valuations and reservation prices.

no literature which does this^[11]). The simplicity this assumption gives the theory is, we feel, well worth the price.^[12]

A question that naturally arises is whether optimal behavior for a game-theoretic formulation of the DOA is consistent with our behavioral rules. Since this would be a dynamic incomplete-information game, and since the common knowledge that is an integral part of recent game theory is not controlled for in the experiments, current theory provides little guidance. However, as we indicated in section 3, footnote 4, if one recognizes the important difference between a game of *imperfect* information with *objective common knowledge* priors and a game of *incomplete* information, then one can understand how Bayesian game theory is consistent with our rules. In particular, even assuming risk neutrality, we can construct a vector of strategies, one for each agent, such that i 's component of a Bayes equilibrium for the DOA game *from i 's subjective point of view*, such that revisions in beliefs as the game is played using these strategies satisfy the Bayes' rule *on nonzero probability events*, and such that the trades, bids, and offers are consistent with Assumptions 1, 1', 2, and 2'. Because of the absence of objective common knowledge, agents may be surprised, even on the equilibrium path, but neither the Bayes rule nor the Bayes-Nash equilibrium prevents this possibility.^[13] Nevertheless, we do not believe that this game-theoretic behavior is what subjects in DOA experiments are really doing. Thus we prefer to analyze our more general model

[11] Nossair¹⁵ has recently solved this problem for a uniform-price sealed-bid mechanism, but it remains unsolved in general.

[12] Holt, Langan, and Villamil⁸ report a series of experiments in which traders had multiple units with payoffs structured to give some traders market power on some units. Their data are nonetheless reasonably consistent with the predictions of our theory. So our implicit assumption that traders decide on strategies for each unit separately seems not to be at odds with the facts.

[13] We describe the equilibrium from one agent's point of view. Suppose that every trader believes that values are drawn such that two or more buyers and sellers have p^* as their value and that $S(p^*) = D(p^*)$. (This is the way a typical experiment is set up. The only new feature here is having all traders believe the same p^* and believe it with certainty.) Further, suppose that this subjective belief is common knowledge. Consider a buyer's strategy which is to bid p^* if $V \geq p^*$, to bid V if $p^* > V$, to accept any offer $o \leq p^*$ if $o \leq V$, and to reject any offer $o > p^*$. The buyer is assumed to follow this strategy forever and sellers are assumed to follow symmetric strategies. To complete the description of an equilibrium, we have to describe the updating of beliefs. With the proposed strategies any bid above p^* or offer below p^* is a zero probability event so the Bayes rule has no implications in this case. Let p_d^* be the common believed price for day d with $p_0^* = p^*$. If no zero-probability events are observed in d , then $p_{d+1}^* = p_d^*$. If any zero-probability event is observed, let p_{d+1}^* be one of the zero-probability prices in $(\underline{P}_d, \bar{P}_d)$. The claim is that these strategies and updating rules define a Bayes-Nash equilibrium. Consider any buyer. He knows that his bids can affect beliefs only if the bids are above p_d^* . But this is undesirable, at least until the end of the day, as a high bid has him paying more than he believes is necessary and can only raise future prices. At the end of the day, a buyer who does not hold the outstanding bid of p_d^* knows that he cannot trade at p_d^* as he believed. (This will happen if and only if there are two or more buyers who have not yet traded at $T - 1$. Further, the buyer who will not be able to trade as expected knows who he is.) This is a zero-probability event (according to his belief) so he can plan to do anything in this contingency. To complete the description of the equilibrium strategy, we assume that he bids $p_d^* + 1$ if $p_d^* + 1 \leq V$ and V otherwise. (cont'd.)

recognizing that one set of behavioral rules satisfying Assumptions 1, 1', 2, and 2' are consistent with a game-theoretic treatment.

Now a final observation. We believe that there are traders in the experiments whose behavior is, at least for a few iterations, vastly different from behavior which would be consistent with our assumptions. In particular there are traders, who hold out for a highly profitable trade to the end of the day even though they never complete one. These traders usually modify their behavior after a few days. Those who do not lose a considerable amount of opportunity income. We do not attempt to explain their irrationality.

We turn now to the derivation of a number of testable implications of the theory. We then confront these with the data from a small number of representative experiments. At that point, the reader should be able to decide whether or not our model offers a realistic description of actual behavior in double auctions.

5. THEOREMS

In this section we trace through some of the implications of our theory. As will become apparent, most of the action will occur when there is an "excess demand or supply" of two or more units remaining in the auction, as there are then competitive pressures on bids and offers. Thus we are interested in the following concepts.

DEFINITION Let $D^c(P) = \#\{V^i \geq P\}$, $D^o(P) = \#\{V^i > P\}$, $S^c(P) = \#\{M^j \leq P\}$, and $S^o(P) = \#\{M^j < P\}$. Let $P_* = \min\{P : D^c(P) \leq S^o(P) - 2\}$ and $P^* = \max\{P : S^c(P) \leq D^o(P) - 2\}$.

P_* is the minimum price at which there is an excess supply of two units and P^* is the maximum price at which there is an excess demand of two units. For IPDA14 in Appendix A, $P_*=4.31$ and $P^*=3.99$. In Figure 2, Appendix B, $P_*=101$ and $P^*=99$. An excess of two is important to provide the competitive forces that will drive prices. To see this, consider the following propositions. All results are stated under Assumptions 0, 1, 1', 2, and 2'. Remember \underline{P}_d is the lowest contract price or offer observed during day $d-1$ and is a "lower bound" on the agent's support in day d . \underline{P}_{d+1} will be the lowest contract price or offer observed during day d .

Now consider a buyer who contemplates a defection from the proposed equilibrium by refusing to trade. This can only be valuable if it changes beliefs. But any one buyer knows he can cause only one seller to remain untraded so this strategy will result in no lower offers. This structure describes a Bayes-Nash equilibrium which produces behavior consistent with Assumptions 1, 1', 2, and 2'. We have not explored the possibilities for refinements.

LEMMA 1:

- a. If $\underline{P}_d \geq P_*$, then $\underline{P}_{d+1} < \underline{P}_d$. If $\underline{P}_d < P_*$, then $\underline{P}_{d+1} < P_*$.
 b. If $\overline{P}_d \leq P^*$, then $\overline{P}_{d+1} > \overline{P}_d$. If $\overline{P}_d > P^*$, then $\overline{P}_{d+1} > P^*$.

That is, there are competitive forces driving minimum contract prices below P_* and keeping them there. These same forces drive maximum contract prices above P^* and keep them above P^* .

PROOF: We prove (a); the proof of (b) is symmetric.

Suppose $\underline{P}_d \geq P_*$ and $\underline{P}_{d+1} \geq \underline{P}_d$. As $\underline{P}_d \geq P_*$, we have $D^c(\underline{P}_d) \leq S^o(\underline{P}_d) - 2$. The number of trades in day d is no more than $D^c(\underline{P}_d)$ as by hypothesis all trades have been at price \underline{P}_d or above. Thus at $t=2$ there are at least two sellers j and j' with $M^j, M^{j'} < \underline{P}_d$ who have not yet traded. Then by applying Assumptions 2'(ii)(c) and 1' repeatedly, we have $o_d(0) \leq \underline{P}_d - 1$. But then $\underline{P}_{d+1} < \underline{P}_d$ which contradicts $\underline{P}_{d+1} \geq \underline{P}_d$.

Suppose $\underline{P}_d < P_*$ and $\underline{P}_{d+1} \geq P_*$. Then all trades have been at prices at or above P_* . A minor modification of the argument above then yields a contradiction. QED

LEMMA 2: Suppose $\Delta P_d > 1$; then

- a. If $D^o(\underline{P}_d) \geq S^o(\overline{P}_d)$, then $\underline{P}_{d+1} > \underline{P}_d$.
 b. If $S^o(\overline{P}_d) \geq D^o(\underline{P}_d)$, then $\overline{P}_{d+1} < \overline{P}_d$.

That is, there are competitive pressures driving (the lowest) contract prices up if they are too low relative to the highest prices and driving maximum contract prices down if they are too high relative to the lowest prices.

PROOF: We prove (a); the proof of (b) is symmetric.

Suppose that $D^o(\underline{P}_d) \geq S^o(\overline{P}_d)$ and $\underline{P}_{d+1} \leq \underline{P}_d$. Then there exists a time t' such that either $o_d(t') = \underline{P}_{d+1}$ or $c_d(t') = \underline{P}_{d+1}$. Since $\underline{P}_{d+1} \leq \underline{P}_d$ and $\Delta P_d > 1$, it follows from Assumption 2'(ii)(a),(c) that there exists a time $\hat{t} > t'$ such that $o_d(\hat{t}) = \underline{P}_{d+1}$ was not accepted. Therefore, as $\Delta P_d > 1$, Assumption 1(ii)(a) implies that all units $V^i > \underline{P}_d$ have been traded before time \hat{t} . So the number of trades before time \hat{t} is at least $D^o(\underline{P}_d)$ which is $\geq S^o(\overline{P}_d)$. Then as Assumption 2'(ii)(a) implies that if $\Delta P_d > 1$, all $M^j < \overline{P}_d$ trade before any $M^j \geq \overline{P}_d$, we know that all $M^j < \overline{P}_d$ have been traded before time \hat{t} . So all $M^j \leq \underline{P}_d$ have been traded before time \hat{t} . Then by Assumption 2', $s_d^j(t)$ and

$o_d(t) > \underline{P}_d$ for all $t \leq \hat{t}$ and all j . This contradicts either $o_d(t') = \underline{P}_{d+1} \leq \underline{P}_d$ or $c_d(t') = \underline{P}_{d+1} \leq \underline{P}_d$. QED

Buyers' reluctance to pay too much and sellers' reluctance to accept too little eventually force minimum and maximum contract prices closer together. Of course, the difference between maximum and minimum contract prices does not necessarily decrease every day. In a day where there is excess demand at the upper bound \bar{P} , prices may rise, but they will not go above the cost of unit number $D^\circ(\underline{P})$. This occurs because units up to $D_o(\underline{P})$ trade first (if $\Delta P > 1$) and these can all be traded at prices no more than $M^{D^\circ(\underline{P})}$. Thus the statistic that falls, or at least does not rise, in every period is the maximum of \bar{P} and $M^{D^\circ(\underline{P})}$.

DEFINITION Let $u_d = \max\{\bar{P}_d, M^{D^\circ(\underline{P}_d)}\}$ and $\ell_d = \min\{\underline{P}_d, V^{S^\circ(\bar{P}_d)}\}$.

In IPDA14, Appendix A, suppose $\underline{P}_d=4.00$ and $\bar{P}_d=4.10$. Then $u_d=4.20$ since $D^\circ(\underline{P}_d)=6$ and $\ell_d = \underline{P}_d=4.00$ since $S^\circ(\bar{P}_d)=4$. The next lemma is useful in the proof of the results of main interest further on.

LEMMA 3: If $\Delta P_d > 1$, then $u_{d+1} \leq u_d$, $\ell_{d+1} \geq \ell_d$, and $|u_{d+1} - \ell_{d+1}| < |u_d - \ell_d|$.

PROOF: There are two cases to consider: (1) $S^\circ(\bar{P}_d) \geq D^\circ(\underline{P}_d)$ and (2) $D^\circ(\underline{P}_d) \geq S^\circ(\bar{P}_d)$. We prove the lemma under case 1; the proof under case 2 is symmetric. We first need to establish:

CLAIM 1: If $\Delta P_d > 1$ and $S^\circ(\bar{P}_d) \geq D^\circ(\underline{P}_d)$, then $\underline{P}_{d+1} \geq \ell_d$.

PROOF: Suppose that $\Delta P_d > 1$, $S^\circ(\bar{P}_d) \geq D^\circ(\underline{P}_d)$, and $\underline{P}_{d+1} < \ell_d$.

From the definition of \underline{P}_{d+1} we know that there is a time t' in day d such that $o_d(t') = \underline{P}_{d+1}$ or $c_d(t') = \underline{P}_{d+1}$. Then as $\underline{P}_{d+1} < \ell_d \leq \underline{P}_d$, there must be a time \hat{t} in day d such that $o_d(\hat{t}) = \ell_d$ was not accepted. This implies that all $V^i \geq V^{S^\circ(\bar{P}_d)}$ have traded before time \hat{t} . So the number of units traded before time \hat{t} is at least $S^\circ(\bar{P}_d)$. By Lemma 2(b) we have $\bar{P}_{d+1} < \bar{P}_d$. So the number of units traded in day d is no more than $S^\circ(\bar{P}_d)$. Thus the number of units traded in day d , before time \hat{t} , is $S^\circ(\bar{P}_d)$. So all $M^j < \bar{P}_d$ have traded before time \hat{t} . Then there does not exist a seller unit $M^j \leq \underline{P}_{d+1} < \ell_d$ to offer $o_d(t') = \underline{P}_{d+1}$ or accept a contract at $c_d(t') = \underline{P}_{d+1}$. This contradicts $\underline{P}_{d+1} < \ell_d$. The proof of Lemma 3 now follows directly from Claims 2 and 3.

CLAIM 2: If $\Delta P_d > 1$ and $S^o(\bar{P}_d) \geq D^o(\underline{P}_d)$, then $\ell_{d+1} \geq \ell_d$.

PROOF: By Lemma 2, $\bar{P}_{d+1} < \bar{P}_d$. So $S^o(\bar{P}_{d+1}) \leq S^o(\bar{P}_d)$.

This implies that $V^{S^o(\bar{P}_{d+1})} \geq V^{S^o(\bar{P}_d)}$. By claim 1, $\underline{P}_{d+1} \geq \ell_d$. Now $\ell_{d+1} = \text{Min}\{\underline{P}_{d+1}, V^{S^o(\bar{P}_{d+1})}\} \geq \text{Min}\{\underline{P}_{d+1}, V^{S^o(\bar{P}_d)}\} \geq \ell_d$.

CLAIM 3: If $\Delta P_d > 1$ and $S^o(\bar{P}_d) \geq D^o(\underline{P}_d)$, then $u_{d+1} < u_d$.

PROOF: By Lemma 2(b), $\bar{P}_{d+1} < \bar{P}_d$ and by Claim 1, $\underline{P}_{d+1} \geq \ell_d = \text{Min}\{\underline{P}_d, V^{S^o(\bar{P}_d)}\}$. Thus $D^o(\underline{P}_{d+1}) \leq \text{Max}\{D^o(\underline{P}_d), S^o(\bar{P}_d)\} = S^o(\bar{P}_d)$. So $M^{D^o(\underline{P}_{d+1})} \leq M^{S^o(\bar{P}_d)} < \bar{P}_d$. Then $u_{d+1} = \text{Max}\{\bar{P}_{d+1}, M^{D^o(\underline{P}_{d+1})}\} < \bar{P}_d \leq u_d$. So $u_{d+1} < u_d$. QED

The forces embodied in Lemmas 1, 2, and 3 serve to drive contract prices together and into the interval $[P^*, P_*]$. If supply and demand balance at this point, prices will stay in this interval. Before proceeding to Theorem 1, we need to show that the interval is well defined.

CLAIM 4: $P_* \geq P^*$.

PROOF: Suppose $P^* \geq P_*$. Then $S^c(P_*) \leq D^o(P_*) - 2$ and $D^c(P_*) \leq S^o(P_*) - 2$. So $D^c(P_*) + 2 \leq S^o(P_*) \leq S^c(P_*) \leq D^o(P_*) - 2$. This implies $D^c(P_*) < D^o(P_*)$ which is false.

THEOREM 1: If $D^c(P^*) = S^c(P_*)$, then there exists a day $d^* < \infty$ such that $P^* \leq \underline{P}_d < P_*$ and $P^* < \bar{P}_d \leq P_*$ for all $d \geq d^*$.

PROOF: As the price set—the integers in $[0, \bar{P}]$ —is finite, Lemma 1 implies that there is finite day \bar{d} such that $\underline{P}_d < P_*$ and $\bar{P}_d > P^*$ for all $d \geq \bar{d}$. Then by Lemma 3 there is a finite day $d^* \geq \bar{d}$ such that $[\underline{P}_{d^*}, \bar{P}_{d^*}] \subseteq [P^*, P_*]$ and $\Delta P_{d^*} \leq 1$.

We now prove the theorem by an induction argument. Suppose $[\underline{P}_d, \overline{P}_d] \subseteq [P^*, P_*]$ for some day $d \geq d^*$. We need to show that this implies $[\underline{P}_{d+1}, \overline{P}_{d+1}] \subseteq [P^*, P_*]$. Suppose not, say $\underline{P}_{d+1} < P^*$. Then there is a time t' in day d such that $o_d(t') < P^*$ or $c_d(t') < P^*$. As $P^* \leq \underline{P}_d$, there is a time $\hat{t} \geq t'$ such that $o_d(\hat{t}) = P^*$ was not accepted. As $P^* \leq \underline{P}_d$, this implies that all units $V^i \geq P^*$ have traded before time \hat{t} . So the number of units traded is at least $D^c(P^*) = S^c(P_*)$. As $P_* \geq \overline{P}_d$, this implies that all units $M^j \leq P_*$ have traded before time \hat{t} . Then there is no seller with a unit $M^j < P^* < P_*$ to offer $o_d(t') < P^*$ or to accept $c_d(t') < P^*$. The proof that $\overline{P}_{d+1} \leq P_*$ is symmetric. By the induction argument above and Lemma 1, we have a day $d^* < \infty$ such that $P^* \leq \underline{P}_d < P_*$ and $P^* < \overline{P}_d \leq P_*$ for all $d \geq d^*$. QED

Theorem 1 applies to experiments which have a Walrasian equilibrium price P^e and quantity Q^e . These experiments fall into three groups. Firstly, if there are multiple units at the Walrasian equilibrium price (and if $D^c(P^e - 1) = S^c(P^e + 1)$), then Theorem 1 predicts that prices will eventually remain within one cent of P^e (as $P_* = P^e + 1$, $P^* = P^e - 1$) and that quantity traded will be at least $Q^e - 1$ and no more than the maximum of $S^c(P_* + 1)$ and $D^c(P^* - 1)$. Secondly, if there is only one unit at the Walrasian equilibrium price and $D^c(P^*) = S^c(P_*)$, the situation in IPDA14 in Appendix A, then Theorem 1 predicts that eventually the maximum price will be no more than one cent above the minimum of the value of the first infra-marginal buyer (V^{Q^e-1}) and the cost of the first extra-marginal seller (M^{Q^e+1}). For IPDA14, this is 4.30. The prediction for the minimum price is symmetric. In the limit, prices tend to keep out extra-marginal units and to keep in infra-marginal units. For this class of experiments, the prediction is again that the quantity traded will eventually remain in the interval $[Q^e - 1, \max\{S^c(P_* + 1), D^c(P^* - 1)\}]$. Finally, if the experiment presents an interval of prices, any of which can be a Walrasian equilibrium, with no units at any of these prices and if $D^c(P^*) = S^c(P_*)$, the predictions of Theorem 1 are again that prices eventually remain in $[P^*, P_*]$. However, it is possible to design payoff schedules with one unit at the Walrasian equilibrium or with no units at any Walrasian equilibrium so that $D^c(P^*) \neq S^c(P_*)$. This case and cases where there is no Walrasian equilibrium are addressed by the following theorem.

THEOREM 2:

- a. If $D^c(P^*) > S^c(P_*)$, then there exists a day $d^* < \infty$ such that $P^* \leq \underline{P}_d < P_*$ and $P^* < \overline{P}_d \leq M^{D^c(P^*)}$ for all $d \geq d^*$.
- b. If $S^c(P_*) > D^c(P^*)$, then there exists a day $d^* < \infty$ such that $V^{S^c(P_*)} \leq \underline{P}_d < P_*$ and $P^* < \overline{P}_d \leq P_*$ for all $d \geq d^*$.

PROOF: We prove part (a); the proof for (b) is symmetric. We first need to establish the relationship between P^* , P_* , $M^{D^c(P^*)}$, and $V^{S^c(P_*)}$.

CLAIM 5: If $D^c(P^*) > S^c(P_*)$, then $M^{D^c(P^*)} > P_* > V^{S^c(P_*)} \geq P^*$.

PROOF:

- i. Suppose $P_* \geq M^{D^c(P^*)}$. Then $S^c(P_*) \geq D^c(P^*)$. A contradiction.
- ii. Suppose $V^{S^c(P_*)} \geq P_*$. Then $D^c(P_*) \geq S^c(P_*) \geq S^o(P_*)$. But by definition, $D^c(P_*) + 2 \leq S^o(P_*)$.
- iii. Suppose $P^* > V^{S^c(P_*)}$. Then $D^c(P^*) < S^c(P_*)$. A contradiction.

By the argument in Theorem 1 we know that there is a day $d^* < \infty$ such that $[\underline{P}_d, \bar{P}_d] \subseteq [P^*, P_*]$. By Claim 5, $M^{D^c(P^*)} > P_*$. So $[\underline{P}_d, \bar{P}_d] \subset [P^*, M^{D^c(P^*)}]$. The proof now proceeds by induction. We need to show that if $[\underline{P}_d, \bar{P}_d] \subseteq [P^*, M^{D^c(P^*)}]$, then $[\underline{P}_{d+1}, \bar{P}_{d+1}] \subseteq [P^*, M^{D^c(P^*)}]$. There are two cases to consider: (1) $\bar{P}_d \leq P_*$ and (2) $\bar{P}_d > P_*$.

Case 1: $\bar{P}_d \leq P_*$. As $M^{D^c(P^*)} > P_* \geq \bar{P}_d$ and $V^{S^c(P_*)} \geq P^*$ by Claim 5, an argument similar to the proof of Theorem 1 shows that $\bar{P}_{d+1} \leq M^{D^c(P^*)}$ and $\underline{P}_{d+1} \geq P^*$.

Case 2: $\bar{P}_d > P_*$. We know that $\underline{P}_d < P_*$ for all $d \geq d^*$, so $\bar{P}_d > P_*$ implies that $\Delta P_d > 1$. So by Lemma 3, $u_{d+1} \leq u_d$. By definition $u_d = \text{Max}\{\bar{P}_d, M^{D^o(\underline{P}_d)}\}$ and by hypothesis $\underline{P}_d \geq P^*$. So $D^o(\underline{P}_d) \leq D^o(P^*)$. Thus, $M^{D^o(\underline{P}_d)} \leq M^{D^o(P^*)} \leq M^{D^c(P^*)}$. By hypothesis $\bar{P}_d \leq M^{D^c(P^*)}$. So $u_d \leq M^{D^c(P^*)}$. By definition $u_{d+1} = \text{Max}\{\bar{P}_{d+1}, M^{D^o(\underline{P}_{d+1})}\}$. Now $u_{d+1} \leq u_d \leq M^{D^c(P^*)}$. So $\bar{P}_{d+1} \leq M^{D^c(P^*)}$.

We also need to show that $\underline{P}_{d+1} \geq P^*$. Suppose not. Then $\underline{P}_{d+1} < P^*$. This requires $S^o(\underline{P}_d) > D^c(\underline{P}_d)$. Thus, $S^c(P_*) \geq S^o(P_*) \geq S^o(\underline{P}_d) > D^c(\underline{P}_d) \geq D^c(P^*)$. This contradicts $D^c(P^*) > S^c(P_*)$. So $\underline{P}_{d+1} \geq P^*$.

Theorem 2 now follows from the induction argument above and Lemma 1. QED

Although Lemmas 1, 2, and 3 imply that prices are eventually contained in the interval $[P^*, P_*]$, they need not stay in this interval if supply and demand are not equal there. For example, if $D^c(P^*) > S^c(P_*)$ and low-value buyers (those with $P^* - 1 \leq V^i < P_*$) trade first, the remaining high-value buyers may bid prices up. However, they need not, and so will not, bid more than $M^{D^c(P^*)}$ in order to complete a trade. So the range of prices could expand to be $[P^*, M^{D^c(P^*)}]$.

In subsequent days it will shrink until it is again contained in $[P^*, P_*]$. It seems unlikely that this process would continue, and our theory does not predict that it will, only that it might. In fact, Lemma 3 implies that for all supply and demand configurations if all extra-marginal units are excluded by $[\underline{P}_d, \bar{P}_d]$, then the interval will shrink to at most one cent and then remain fixed.

6. COMPARISONS OF THE PREDICTIONS WITH THE DATA

Our prediction of convergence seems consistent with the experimental data, but it is not directly testable with these data as the number of repetitions necessary for convergence is not specified. In any case, obtaining the competitive equilibrium in the limit is only a first test of a theory of price formation in double oral auctions. We have rejected the models considered in Section 3, at least in part, on the basis of their incorrect predictions about dynamics. In this section we compare the predictions of our model with experimental data. There are three categories of data for which our theory has implications: The sequence of minimum and maximum prices, the sequence of trading partners, and the number of units traded.

The three lemmas in Section 5 directly yield predictions about the dynamics of minimum and maximum prices. Lemma 1 implies that these prices move to bracket the competitive equilibrium price and that once this is accomplished, the theoretical equilibrium price remains in the interval $[\underline{P}, \bar{P}]$. Lemmas 2 and 3 imply that minimum and maximum prices respond to the forces of demand and supply. The prediction is that the minimum price will rise if demand at \underline{P} exceeds supply at \bar{P} and that the maximum price will fall if supply at \bar{P} exceeds demand at \underline{P} . In the excess demand case ($D^\circ(\underline{P}) \geq S^\circ(\bar{P})$), the maximum price may rise, but the prediction is that it will go no higher than the level necessary to allow the $D^\circ(\underline{P})$ th unit to trade. For the excess supply case, the prediction is that although the minimum price may fall it will not go below $V^{S^\circ(\bar{P})}$.

Remember that \underline{P}_d is the lowest contract price or offer that occurred during the day *before* d . \bar{P}_d is the highest contract price or bid observed during the day *before* d .

PREDICTION 1. Prices (See Lemmas 1, 2, and 3):

- i. If $\underline{P}_d \geq P_*$, then $\underline{P}_{d+1} < \underline{P}_d$. If $\underline{P}_d < P_*$, then $\underline{P}_{d+1} < P_*$.
- ii. If $\bar{P}_d \leq P^*$, then $\bar{P}_{d+1} > \bar{P}_d$. If $\bar{P}_d > P^*$, then $\bar{P}_{d+1} > P^*$.
- iii. If $\Delta P_d > 1$ and $D^\circ(\underline{P}_d) \geq S^\circ(\bar{P}_d)$, then $\underline{P}_{d+1} > \underline{P}_d$ and $\bar{P}_{d+1} \leq u_d$.
- iv. If $\Delta P_d > 1$ and $S^\circ(\bar{P}_d) \geq D^\circ(\underline{P}_d)$, then $\bar{P}_{d+1} < \bar{P}_d$ and $\underline{P}_{d+1} \geq \ell_d$.

TABLE 1 Violation Percentages¹

Experiment	Exp	Marg	Descriptors					Analysis		
			Com	Units	Q^e	NYSE	Que	Price	Seq	Quant
IPDA8	No	1	5	8	6	Yes	No	13.9	0	11.1
IPDA9	No	1	5	8	6	Yes	No	9.4	1.6	0
IPDA10	Yes	1	5	10	8,6	Yes	No	6.3	0	0
IPDA11	No	1	5	10	8	Yes	No	18.8	1.3	0
IPDA14	Yes	1	5	8	6	Yes	No	9.4	0	0
IIPDA14	No	3	10	21	15	Yes	No	12.5	0	0
IIPDA22	No	0	10	16,11	11	No	No	3.1	NA	12.5
IIPDA25	Yes	2	10	12	7	No	Yes	21.9	0.8	12.5
IIPDA57	No	3	10	21	15	Yes	Yes	16.7	1.0	8.3
Average								12.7	0.7	5.3

1

PDA =Plato Double Auction

Exp =experienced subjects (have participated in another PDA)

Marg =number of extra marginal units

Com =commission in cents per unit traded

Units =number of units on each side of the market

Q^e =competitive equilibrium quantity

NYSE =New York Stock Exchange rules (new bids and offers must improve on outstanding bids and offers)

QUE =electronic queuing of bids and offers (see Smith and Williams²¹)

Price =% violation of price predictions

Seq =% violation of trading sequence predictions

Quant =% violation of quantity traded predictions

Our model yields no predictions for day 1 of any experiment. In any experiment in which supply or demand was shifted, we treat the first day after the shift as day 1 of a new experiment. In each non-initial day of an experiment, we have four possible violations of price predictions: violations of 1(i) and 1(ii), and the two predictions of either 1(iii) or 1(iv). So, in an experiment running for ten days with no shifts, there are 36 possible price violations. The entry for price violations is the number of violations divided by the number of possible violations. In each non-initial day the number of possible trading sequence violations is the number of buyer plus seller units ($n + m$). The number of actual violations is the number of units traded out of order. In each non-initial day there is one quantity prediction, and thus one possible quantity violation.

TABLE 2 Price Violations

Price violations of $x\%$ or less not counted	Percentage of price violations over nine DOAs
$x = 0.01$	6.3
$x = 0.05$	3.3
$x = 0.10$	1.3

Our prediction about the sequence of trading partners follows from the proofs of the Lemmas. It is essentially that sellers with costs below \bar{P} trade before those with costs above \bar{P} and buyers with values above \underline{P} trade before those with values below \underline{P} .

PREDICTION 2. Trading Sequence:

- i. If $\Delta P_d > 1$:
All sellers with cost below \bar{P}_d trade before any sellers with cost at or above \bar{P}_d .
All buyers with value above \underline{P}_d trade before any buyers with value at or below \underline{P}_d .
- ii. If $\Delta P_d \leq 1$:
All sellers with cost at or below \bar{P}_d trade before any sellers with cost above \bar{P}_d .
All buyers with value at or above \underline{P}_d trade before any buyers with value below \underline{P}_d .

Our prediction about the number of units traded is that it will be at least the competitive equilibrium for demand and supply curves truncated at \bar{P} and \underline{P} , respectively, less one unit.

PREDICTION 3. Quantity Traded:

The quantity traded in day d will be $Q_d \geq \text{Max}\{K : \hat{V}_d^K \geq \hat{M}_d^K\} - 1$ where $\hat{V}_d^K = \text{Min}\{\underline{P}_d, V^K\}$ for $K = 1, \dots, n$ and $\hat{M}_d^K = \text{Max}\{\underline{P}_d, M^K\}$ for $K = 1, \dots, m$.

Table 1 summarizes, for nine DOA experiments, violations of our predictions about prices, sequence of trades, and number of units traded as a percentage of total possible violations. This table is based on data from Williams,²² and on unpublished data which were made available by Vernon Smith. We realize this table is neither easy to understand nor conclusive empirical verification of our theory. We offer it

only as supporting evidence of plausibility. We have yet to see a DOA experiment which is much different in its violations of Predictions 1 to 3 than those in Table 1.

To see the total number of price violations in perspective, Table 2 illustrates the margins of error. This table reports the total number of violations of our price predictions (over all nine experiments) which were more than x cents, as a percentage of the total number of possible violations of our price predictions.

To put sequence and quantity violations in context, it is useful to compare them with the violations of the sequence and quantity predictions of the Marshallian theory and the sequence predictions of the game theory approach. The Marshallian theory predicts that units will trade in the order of value and that all profitable trades will occur. The violations of this prediction as a percentage of possible violations in IPDA8 is 42.5%. The game theory approach²³ predicts that units will trade in the order of value but yields no further prediction on the number of trades. The violations of this prediction as a percentage of possible violations in IPDA8 is 29.4%.

7. FURTHER EXPERIMENTS

There is now a role for further interaction between theory and experiments. The class of experiments described in Section 6 motivated our theory, and it in turn suggests several experiments which could lead to refinements or rejection of the theory. There are several aspects of our theory which could be tested. Firstly, we do not assume that traders' reservation prices and bids or offers converge to their true values at the end of each day. The data that we have seems to reject such an assumption. However, without this assumption we can establish convergence only to an interval determined by P_* and P^* . Our theory admits as an equilibrium a situation in which one extra-marginal unit is included or in which one infra-marginal unit is excluded. For example, Theorem 1 applies to the demand and supply configuration in Figure 1 of Appendix B to predict equilibrium in the interval [114,148]. The placement of the first extra-marginal units in that figure has no effect on our equilibrium prediction. Charles Plott and Chris Worrell have run a DOA experiment using the configuration of Figure 1. Their data suggest that prices converge into the interval [133,139] determined by the first extra-marginal units. This conclusion does not reject our theory, but it does suggest that the theory might be refined to produce sharper convergence results.

Secondly, we have refrained from placing any direct assumption on the relative (between agent) rankings of true values and reservation prices. Possible ranking hypotheses on buyers' reservation prices include (1) if $V^i > V^j$, then $r_d^i(t) > r_d^j(t)$ and (2) if $V^i > \bar{P}_d$ and $V_j \leq \underline{P}_d$, then $r_d^i(t) > r_d^j(t)$. Hypothesis 1 is clearly rejected

by the data, but whether hypothesis 2 is rejected depends on one's standard of acceptance. The absence of a ranking hypothesis is responsible for the relatively weak prediction of Theorem 2. In DOAs where $D^c(P^*) \neq S^c(P_*)$, our theory predicts convergence into the interval $[P^*, P_*]$, but it then admits the possibility of cycles between prices in this interval and prices as low as $V^{S^c(P_*)}$ if $S^c(P_*) > D^c(P^*)$ or prices as high as $M^{D^c(P^*)}$ if $D^c(P^*) > S^c(P_*)$. In the presence of either ranking hypothesis (and a symmetric hypothesis on sellers' reservation prices) cycles would not occur and prices would remain in $[P^*, P_*]$. We have some data about experiments where our theory admits the possibility of cycles. In both IPDA8 and IPDA9 (reported in Section 6), $S^c(P_*) > D^c(P^*)$. In neither of these experiments do we see cycles; prices seem to remain approximately in $[P^*, P_*]$. However, the extra-marginal seller unit at $P_* - 1$ is occasionally traded, so it is possible that cycles would have arisen had the experiments continued beyond ten days.¹⁴ This suggests two possible further experiments. Firstly, IPDA8 could be run for more days to decide whether cycles will appear. Secondly, an experiment with a design more likely to produce cycles could be run. The supply and demand configuration of Figure 2 in Appendix B is such a design. The prediction of Theorem 2 for this configuration is that prices will remain in the interval $[V^{S^c(P_*)}, P_*] = [70, 101]$. Our conjecture is that in the experiments any cycles would eventually disappear, with prices remaining in $[P^*, P_*]$ and perhaps following a time path during each day starting at P^* , and then rising during the day. By offering a small discount (to P^*) early in the trading day, infra-marginal sellers could insure that they complete a trade. Prices would then rise by one or two cents as marginal traders complete their trades. If this occurs, it suggests that the theory might be further refined.

There are several other experiments which could lead to refinements or rejection of our theory. Firstly, our theory is silent about the fine details of organizing a DOA. All that counts is that traders can make bids or offers and acceptances, and that they are informed of others' bids, offers, and acceptances. Thus the predictions of the theory are unchanged by the use of New York Stock Exchange rules, electronic queues, or other details. However, the data are not unchanged by these details (see Smith and Williams²¹). It may be that sharper predictions would result if these details were taken into account. Secondly, the theory does not apply to experiments in which one side of the market is not allowed to bid or offer. See Plott and Smith¹⁶ for some of these experiments. Third, the theory does not yield predictions about the effect of shifts in supply and demand curves. There is now data from experiments in which supply and demand curves are shifted systematically. The theory would need to be refined to yield useful predictions about the effect of such changes in market conditions.

¹⁴In IPDA8 the extra-marginal unit at $P_* - 1$ is seller 2's second unit. This unit is traded in both days 9 and 10 of the experiment.

The methodology of using experiments to test the predictions of theory can also be applied to the alternative theories that we have described. For instance, Wilson's game theoretic model of DOAs does not directly apply to the existing DOA experiments, but a DOA experiment could be designed to test the theory. Trader's values and costs could be drawn independently across days from distributions which the traders know, risk attitudes could be controlled for, and the experiment could be repeated for a number of days to allow for learning about the game and about strategies. Wilson's predictions could then be compared to data from the final day of the experiment.

8. CONCLUSION

The theory presented here is deterministic and, although it does not completely describe precise paths of bids, offers, and contracts, it does place fairly tight bounds on these data. One observation not in accord with these bounds is grounds for rejection of the theory, and in fact there are a number of such observations. However, the low percentage of observations which violate the obvious implications of the theory seems acceptable for a theory which yields fairly precise predictions about prices and trades.

The potential importance of this theory is not just that it seems to describe what happens in DOA experiments, but also that it is the beginning of a positive theory of how market prices are formed and of how they adjust to changes in demand and supply conditions. The question of price formation has a long history of ad hoc and unsuccessful attempts at an answer. Our theory is also ad hoc in the sense that we make assumptions on individual behavior which are not derived from an optimizing model. However, our assumptions seem sensible, are consistent with the behavior implied by at least one optimizing, game-theoretic model and, more importantly, they seem to do a reasonable job of describing actual bids, offers, and contracts. There is now a target for experimentalists to reject with data or for theorists to improve on by obtaining a better fit to the data.

9. SOME ACKNOWLEDGMENTS AND A HISTORICAL NOTE

This paper benefited from discussions in seminars at Cornell, Northwestern, and Stony Brook, and an NSF Conference on Experimental Economics at the University of Arizona. It has also been affected by several referees, some of whom have been helpful in forcing us to more clearly present our ideas. This version is significantly

different from earlier versions. We would like to thank Vernon Smith and Arlington Williams for making data on their Plata DOA experiment data available to us.

We would like to thank the editors for inviting us to participate in this volume. We thank John Rust for encouraging us to expose our theory in the testbed of the computerized Double Auction Market at the Santa Fe Institute. We feel reasonably happy with our fifth place finish and like to think that even though our theory was designed for continuous-time open-outcry systems rather than the synchronized double auction, it (the theory) held up pretty well. We thank Dan Friedman for not forgetting our paper and for forcing us, once again, to drag it out of retirement. Finally we thank Bob Wilson for his insights and Vernon Smith and Arlington Williams for providing data and support in the early days of 1980-81.

Because the first version of this paper appeared in 1981 and the last significant revision was made in 1988, some of the references and most of the data may seem somewhat outdated. The task of constructing a new type of theory to explain a new type of data derived from nascent experimental markets was actually begun by us in 1978. Since 1981, game theory arose as a new paradigm in economics and many editors were loath to publish a non-optimizing base theory such as that we had proposed. Our last 1988 revision was written to respond to these concerns. Since then, game theorists analyzing dynamics have turned from models with common knowledge and Bayes equilibrium to models of Bayesian learning, fictitious play, and other non-optimizing models of behavior which do not require, for example, common knowledge of rationality. We are glad others are pursuing the path we took in 1980. We believe that many of these "learning" models are consistent with our assumptions, but that is another paper.

Not only has theory changed since 1981 but experimental technology has significantly advanced. As one can see, from articles in this volume, the simple DOA experimental market we have described in this paper has become more sophisticated and interesting. In our sporadic efforts at casual empiricism, we have yet to see data in the more recent experimental markets that significantly differs from the predictions of our theory. For example, DOA's where commissions are not paid look like the DOA's we analyze. Of course what is needed is a serious empirical study to see whether our impressions are valid. As far as we know, our challenges to theorists and experimentalists in Sections 7 and 8 remain unaccepted. We hope someone responds.

APPENDIX A**TABLE 3 Values and Costs for DOA #IPDA14**

	Week 1		Week 2	
	Unit 1	Unit 2	Unit 1	Unit 2
BYR1	5.20	3.80	3.70	3.60
BYR2	5.00	4.00	3.80	3.50
BYR3	4.80	4.20	3.90	3.40
BYR4	4.60	4.40	4.00	3.30
SLR1	3.70	4.40	3.10	3.30
SLR2	3.80	4.30	2.90	3.50
SLR3	3.90	4.20	2.70	3.70
SLR4	4.00	4.10	2.50	3.90

TABLE 4 Period 9

	MKR	TM	BIDS	OFFERS	TKR	TM
1	B2	297	3.30			
2	B1	295	3.34			
3	B2	292	3.35			
4	S4	289		3.45	B3	285
5	B2	281	3.30			
6	B4	278	3.35			
7	B2	275	3.36			
8	S3	274		3.50		
9	B1	268	3.39			
10	B4	258	3.40		S2	252
11	S1	254		3.45	*	
12	B2	247	3.30			
13	B1	245	3.40			
14	S3	244		3.50		
15	S1	236		3.45		
16	B2	231	3.41			
17	S3	191		3.44		
18	S1	181		3.43	B2	154
19	B4	151	3.30			
20	B2	147	3.31			
21	S2	146		4.50		
22	B1	146	3.40			
23	B2	140	3.41			
24	S1	136		3.45	B1	68
25	B2	122	3.42		*	
26	B1	112	3.43		*	
27	B2	106	3.44		*	
28	S2	63		4.50		
29	B4	62	3.30			
30	B1	61	3.40			
31	S3	57		3.49		
32	B2	55	3.41			
33	B1	45	3.43			
34	B2	41	3.44			
35	B2	23	3.45		S3	17
36	B4	11	3.30			
37	S2	10		3.50	B1	0
38	B1	8	3.40		*	

TABLE 5 Summary Data from IPDA14

	Contract Price	Buyer	Seller
Day 1	4.25	2	3
	4.20	4	1
	4.50	3	2
	4.40	4	4
	4.30	1	2
Day 2	4.35	3	1
	4.30	4	2
	4.30	4	2
	4.30	1	4
	4.25	2	4
Day 3	4.25	4	1
	4.35	3	3
	4.30	4	2
	4.27	2	4
	4.25	1	3
Day 4	4.30	4	1
	4.39	3	4
	4.30	4	4
	4.26	2	2
	4.25	1	3
	4.20	3	2
Day 5	4.30	4	1
	4.26	2	2
	4.35	3	4
	4.26	1	3
	4.25	4	4
Day 6	3.35	4	1
	3.30	3	2
	3.35	3	3
	3.32	2	1
	3.35	1	4
Day 7	3.31	2	2
	3.35	3	1
	3.40	3	4
	3.38	2	1
	3.45	1	3

TABLE 5 Summary Data from IPDA14
(cont'd.)

	Contract Price	Buyer	Seller
Day 8	3.40	4	2
	3.40	1	1
	3.40	1	1
	3.40	3	4
	3.40	2	3
Day 9	3.45	3	4
	3.40	4	2
	3.43	2	1
	3.45	1	1
	3.45	2	3
Day 10	3.50	1	2
	3.41	2	2
	3.44	1	4
	3.45	4	1
	3.45	1	3
	3.50	3	1

APPENDIX B




FIGURE 1 An induced demand-supply schedule.




FIGURE 2 Supply and demand schedules for further experiments. Data is from an unpublished experiment by Charles Plott and Chris Worrell; reprinted by permission of the authors.

REFERENCES

1. Arrow, K. J., and F. H. Hahn. *General Competitive Analysis*. San Francisco: Holden-Day, 1971.
2. Cox, J. C., and R. L. Oaxaca. "Laboratory Experiments with a Finite Horizon: Job Search Model." *J. Risk & Uncertainty* 2 (1989): 301-329.
3. Friedman, D. "On the Efficiency of Double Auction Markets." *Am. Econ. Rev.* 74 (1984): 60-72.
4. Friedman, D. "A Simple Testable Model of Double Auction Markets." *J. Econ. Behav. & Organization* 15 (1991): 47-70.
5. Gresik, T. A., and M. Satterthwaite. "The Rate at Which a Simple Market Converges to Efficiency as the Number of Traders Increases: An Asymptotic Result for Optimal Trading Mechanisms." *J. Econ. Theory* 48 (1989): 304-332.
6. Gronau, R. "Information and Frictional Unemployment." *Am. Econ. Rev.* 61 (1971): 290-301.
7. Harsanyi, J. "Games with Incomplete Information Played by Bayesian Players." Parts I, II, and III. *Mgmt. Sci.* 14 (1967-68): 159-182, 320-334, 486-502.
8. Holt, C. A, L. W. Langan, and A. P. Villamil. "Market Power in Oral Double Auctions: Convergence to Competitive Equilibrium Prices Reconsidered." *Econ. Inquiry* 24 (1986): 107-123.
9. Jarecki, H. G. "Bullion Dealing, Commodity Exchange Trading and the London Gold Fixing: Three Forms of Commodity Auctions." In *Bidding and Auctioning for Procurement and Allocation*, edited by Y. Amihud, 146-154. New York: New York University Press, 1976.
10. Ledyard, J. "The Scope of the Hypothesis of Bayesian Equilibrium." *J. Econ. Theory* 39 (1986): 59-82.
11. Lippman, S. A., and J. J. McCall. "The Economics of Job Search: A Survey." *Econ. Inquiry* 14 (1976): 155-189.
12. Lippman, S. A., and J. J. McCall. "The Economics of Job Search: A Survey." *Econ. Inquiry* 14 (1976): 347-368.
13. Mortensen, D. T. "Job Search, the Duration of Unemployment, and the Phillips Curve." *Am. Econ. Rev.* 60 (1970): 847-862.
14. Myerson, R., and M. Satterthwaite. "Efficient Mechanisms for Bilateral Trade." *J. Econ. Theory* 29: 265-281.
15. Noussair, C. "Auctions with Multi-Unit Demands." Mimeo, California Institute of Technology, 1991.
16. Plott, C. R., and V. L. Smith. "An Experimental Examination of Two Exchange Institutions." *Rev. Econ. Studies* 45 (1978): 133-153.

17. Roth, A., and M. Malouf. "Game Theoretic Models and the Role of Information in Bargaining." *Psych. Rev.* **86** (1979): 574-594.
18. Smith, V. L. "Experimental Economics: Induced Value Theory." *Am. Econ. Rev.* **66** (1976): 274-279.
19. Smith, V. L. "Reflections on Some Experimental Market Mechanisms for Classical Environments," In *Research in Marketing, Choice Models for Buyer Behavior*, edited by Leigh McAlister, 13-47. Greenwich: JAI Press, 1982.
20. Smith, V. L. "Microeconomic Systems as an Experimental Science." *Am. Econ. Rev.* **72** (1982): 923-955.
21. Smith, V. L., and A. W. Williams. "An Experimental Comparison of Alternative Rules for Competitive Market Exchange." In *Auctions, Bidding and Contracting: Uses and Theory*, edited by M. Shubik. New York: New York University Press, 1983.
22. Williams, A. "Computerized Double-Auction Markets: Some Initial Experimental Results." *J. Business* **53** (1980): 235-258.
23. Wilson, R. "Equilibrium in Bid-Ask Markets." In *Arrow and the Ascent of Economic Theory: Essays in Honor of Kenneth J. Arrow*, edited by G. Feiwel. London: McMillan, 1987.