Economic Theory **16**, 613–637 (2000)

Economic Theory

Intertemporal pricing in laboratory posted offer markets with differential information*

Aldo Rustichini[1] and Anne P. Villamil[2]

[1] Department of Economics, Boston University, Boston, MA 02215, USA
[2] Department of Economics, University of Illinois, Champaign, IL 61820, USA

Received: October 1, 1998; revised version: August 30, 1999

Summary. This paper analyzes intertemporal seller pricing and buyer purchasing behavior in a laboratory retail market with differential information. A seller posts one price each period that a buyer either accepts or rejects. Trade occurs over a sequence of "market periods" with a random termination date. The buyer and seller are differentially informed: The seller's cost of producing a unit of a fictitious good is known and constant in all periods, but the buyer's value for the good (demand) is a random variable governed by a Markov Process whose structure is common knowledge. At the beginning of each period the unit's value is determined by "nature" and is privately revealed only to the buyer. The market termination rule is a binary random variable. We conduct 32 laboratory experiments designed to study intertemporal pricing by human subjects in the Posted Offer Institution when demand follows a stochastic process. There are four series of experiments: 8 with simulated buyers, 8 with inexperienced subjects, 8 with once experienced subjects, and 8 with twice experienced subjects.

Keywords and Phrases: Intertemporal pricing, Differential information.

JEL Classification Numbers: C91, D82, D83, D40, C61, L16.

1 Introduction

In this paper we study seller intertemporal pricing strategies and buyer intertemporal purchasing strategies in a laboratory market where demand follows a Markov Process. The market is organized as a Posted Offer institution, where the seller posts a single "take-it" or "leave-it" price which a buyer either accepts or rejects. There is no bargaining or price revision.[1] This institution has

* We gratefully acknowledge financial support from the University of Illinois.
[1] See Davis and Holt (1998) for a recent analysis of price revision in the posted offer market.

been studied extensively in laboratory markets because it corresponds to trade in many naturally occurring retail markets. Indeed, the Posted Offer institution is the most common market structure in developed economies.[2] We extend experimental work on the Posted Offer institution by considering an explicitly intertemporal setting with differential information where demand follows a Markov Process. Most previous work has considered static, deterministic supply and demand conditions with private information about seller costs and buyer values for a fictitious good. In contrast, in our study the buyer's value for the good (i.e., the induced demand curve) follows a Markov Process whose structure is common knowledge, but the current realization is privately revealed only to the buyer at the beginning of each "market period." The precise date that the market ends is determined randomly and is independent of the stochastic demand process. This termination procedure corresponds to an intertemporal discount rate.

Intertemporal pricing and purchasing relationships are clearly important in many naturally occurring markets. In addition, many markets are subject to recurrent cyclical fluctuations (e.g., seasonal and/or business cycle movements). To our knowledge, the only experimental papers that consider fluctuating demand conditions in the Posted Offer institution are by Hoffman and Plott (1981) and Davis, Harrison, and Williams (1992). However, the fluctuations in these models are perfectly predictable (e.g., follow a high, low, high, low pattern, etc.). In contrast, in our experiments demand fluctuations are uncertain but follow a well defined stochastic process. Extending experimental work in this direction is important because a main contribution of experimental research has been to establish that certain market institutions drive sparsely informed agents to equilibrium in static, deterministic environments. Indeed, many economists have argued that this is why market economies have been so successful (relative to non-market economies).[3] A crucial unanswered question in this literature is — Are markets as effective in stochastic, intertemporal environments where agents problems' are more formidable? This paper analyzes this question.

The paper is organized as follows. In Section 2 we discuss the theoretical model we will test and specify the equilibria. In Section 3 we discuss the testable implications of the theory. In Section 4 we discuss our experimental design and the results of 32 laboratory experiments. Finally, Section 5 contains concluding remarks.

2 The model and equilibria

Consider a Posted Offer market with two risk neutral traders, a buyer and a seller, who may trade over a sequence of periods. In each period, the seller may produce one unit of an indivisible good. The seller's production cost is fixed and common knowledge. In contrast, the buyer has a reservation value for the

[2] See Plott (1989) for a survey of the large experimental literature on this institution. See also Davis and Holt (1996).

[3] See Smith (1982) for a discussion of markets as "economizers of information."

good each period v which follows a Markov Process that is common knowledge. The unit's value v takes on one of two possible values, h (high) or l (low). The Markov Process describes the serial correlation in the unit's value over the sequence of market periods. The structure of the process is common knowledge, but the current realization of v is privately observed by the buyer each period. In particular, all agents know that given the unit's value in the previous period, the probability that the unit's value is the same in the current period is $1 - \alpha$, and the probability that it changes (i.e., either from high to low or from low to high, given the previous state) is α. The unit's value to the buyer in the first period (i.e., the initial state) is a random variable, drawn from a known uniform distribution. We restrict attention to the case of positive serial correlation (i.e., $0 < \alpha < 1/2$).[4]

Each period trade occurs according to the following sequence of events. At the beginning of the period the unit's value is determined by "nature." The value of the first unit is drawn randomly from a known distribution. The value of each subsequent unit is determined by the stationary Markov process. A complete description of the process is common knowledge of both agents, but the unit's current value is privately observed by the buyer only. The seller posts a single price offer that the buyer can either accept or reject. If the buyer accepts, the unit is traded at the price posted by the seller. The buyer's profit is the difference between the unit's value and the posted price $v - p$, and the seller's profit is the difference between the posted price and the cost $p - c$. If the buyer rejects the offer the unit is not produced and is not traded, and both traders make zero profit. This concludes the market period. The probability that the market continues in each subsequent period is given by $\delta \in (0, 1)$. For risk neutral traders, this is equivalent to assuming a discount factor of δ. These procedures are repeated in each period until the market ends.[5]

We focus on equilibria in *stationary Markov* strategies. Given a distribution of values for the initial unit, a history of past seller price offers, and buyer answers, the seller forms a belief (for any equilibrium of the game) about the buyer's value for the current unit of the good. Each period the seller considers only this belief about the current unit's value to the buyer, and the buyer considers only his/her own (known) value and the seller's belief about the value. The sequence of previous price offers and answers is reflected in the seller's current belief. To derive equilibrium strategies, we assume that the seller's initial belief is common knowledge.[6] A stationary strategy for the seller is a map from the seller's belief to the space of price offers, where the seller's belief is a number $w \in [0, 1]$ which denotes the seller's subjective probability that the buyer's value for the current unit of the good is *high*. A stationary strategy for the buyer is a map from the

[4] Three other cases are possible: if $\alpha = 0$ the initial draw determines the unit's value for all periods; if $\alpha = 1/2$ the unit's values are independent; and if $1/2 < \alpha \leq 1$ the unit's values exhibit negative serial correlation.

[5] See Villamil (1999) for a discussion of alternative termination procedures.

[6] This assumption is induced in the experiments by drawing the first unit's value randomly from a fixed distribution that is known by subjects.

unit's value for that period, the seller's belief, and the seller's price offer to a binary decision variable which indicates the buyer's answer.

Rustichini and Villamil (1996) prove that equilibria in stationary strategies are characterized by a triple (w^*, p_l, p_h), which indicates the *seller's critical belief*, the *low price offer* and the *high price offer*, respectively. Equilibrium strategies for each agent are a pair $(p(w), A(p, v))$, where

$$p(w) = \begin{cases} p_h & \text{if } w \geq w^*; \text{ or} \\ p_l & \text{otherwise;} \end{cases} \tag{P_s}$$

and

$$A(p, v) = \begin{cases} Y & \text{if either } p \leq p_l \text{ or } p \leq p_h \text{ and } v = h; \\ N & \text{otherwise.} \end{cases} \tag{P_b}$$

(P_s) is a state dependent seller price offer, and (P_b) is a state dependent buyer answer where Y denotes "yes" and N denotes "no."

Rustichini and Villamil (1996) show that in equilibrium, strategies (P_s) and (P_b) solve the following dynamic programming problems:[7]

The Seller's Problem: A stationary strategy for the seller is a pricing function p. The seller knows his/her own beliefs, w, and value function, $V_s(\cdot)$. Let $\mathbf{1}_{\{A=Y\}}$ denote the indicator function of the set $\{A = Y\}$, which is a random event from the seller's perspective, and c denote the seller's (known) cost of producing each unit sold.[8] The seller's discounted dynamic programming problem is to choose a pricing function to maximize the following functional equation:

$$V_s(w) = \max_{p \geq 0} E_w\{p\mathbf{1}_{\{A=Y\}} + \delta V_s(w'(p, A, w))\},$$

where E_w denotes the expectation over w and δ is the discount factor.

The Buyer's Problem: A stationary strategy for the buyer is an answer function $A \in \{Y, N\}$. The buyer knows the realization of the unit's current value, $v = \{h$ or $l\}$, takes the seller's price as given, and knows his/her own value function, $V_b(\cdot)$. The buyer's discounted dynamic programming problem is to choose an answer function to maximize the following functional equation:

$$V_b(w, p, v) = \max_{A \in \{Y,N\}} E_v \begin{cases} v - p \\ +\delta V_b(w'(p, Y, w), p'(w'(p, Y, w)), v') & \text{if } A = Y; \\ \delta V_b(w'(p, N, w), p'(w'(p, N, w)), v') & \text{if } A = N. \end{cases}$$

The expectation is taken over v, the stochastic process for the unit's value, where v is the current value.

Rustichini and Villamil (1996) show that the equilibria of the game are given by the triple (w^*, p_l, p_h). A crucial step in completing the analysis of the equilibria is to provide a rule that the seller follows to form beliefs, both on and off the equilibrium path. Let w denote the seller's belief that the buyer's value for the current unit is high, and $w'(\cdot)$ denote the seller's belief that the buyer's value

[7] "Primes" denote next period's value of a variable.

[8] The seller's cost plays no role in our analysis so we normalize $c = 0$.

for *next period's* unit is high. If the seller receives no additional information in a period, the seller's next belief is given by:

$$\hat{w}(\cdot) = w(1 - \alpha) + (1 - w)\alpha.$$

This equation indicates that the value for next period's unit may be high for two reasons: the current unit's value was high and did not change (the first term) or the unit's value was low but changed states (the second term).

The stationary strategy of each agent is binary. There are five possible belief situations on and off the equilibrium path:

(i) $w'(p, Y, w) = 1 - \alpha$ for any $p_h \geq p > p_l$: If the seller posts a price higher than the low price but less than the high price, and the buyer accepts, then the seller believes that the unit's value was high.

(ii) $w'(p, N, w) = \alpha$ for any $p > p_l$: If the seller posts a price higher than the low price and the buyer rejects, then the seller believes that the unit's value was low.

(iii) $w'(p, Y, w) = w(1 - 2\alpha) + \alpha \equiv \hat{w}$ for any $p \leq p_l$: If the seller posts a price lower than or equal to the low price and the buyer accepts, then the seller has no useful new information and revises his/her belief according to $\hat{w}$.

(iv) $w'(p, N, w) = \hat{w}$ for any $p < p_l$: We assume that if the seller posts a price less than the low price and the buyer rejects, then the seller updates as in the case where the buyer accepts.

(v) $w'(p, Y, w) = \hat{w}$ for any $p > h$: If the seller posts a price higher than the high reservation value and the buyer accepts, then we assume that the seller again updates via $\hat{w}$.

Beliefs (i), (ii), and (iii) are equilibrium path strategies that follow directly from $\hat{w} = w(1 - \alpha) + (1 - w)\alpha$. When the buyer accepts the seller's high price offer, this is a perfect signal that $v = h$; thus $w = 1$ and w' is given by (i). When the buyer rejects the seller's high price offer, this is a perfect signal (in equilibrium) that $v = l$; thus $w = 0$ and w' is given by (ii). When the buyer accepts the seller's low price offer, this action is *not* perfectly revealing; thus w' is given by (iii). Finally, (iv) and (v) are "off the equilibrium path" strategies, so we must attribute to the agent some belief to complete the belief specification rule. We assume that if the buyer rejects the seller's low price offer the seller believes $v = l$ so $w = 0$ and $\hat{w} = \alpha$. This is plausible because when $v = l$ the buyer loses nothing by rejecting p_l but if $v = h$ the buyer foregoes substantial profit. If the buyer accepts a price higher than his/her reservation value, the buyer loses profit on the trade but this gives the seller no new information. Thus the seller updates via $\hat{w}$.

3 Testable implications of the theory

Rustichini and Villamil (1996) prove that strategies (P_s) and (P_b) solve the respective seller and buyer problems.[9] In our intertemporal Posted Offer market the seller moves first by posting a price, but does not observe directly the actual realization of the unit's value in any period. The buyer moves second and responds to the seller's price offer—after the unit's current value has been privately revealed. When $0 < \alpha < 1/2$ and the seller is rational (i.e., uses all available information and Bayes rule), equilibrium strategy (P_s) indicates that the seller behaves as if he/she forms a belief, w, about the unit's value to the buyer, and then decides whether to offer a high or a low price by comparing the current belief with a critical belief, w^*. The seller posts a high price (p_h) if $w \geq w^*$, and a low price (p_l) otherwise. The seller uses this strategy both to maximize revenue *and* to acquire information. Equilibrium strategy (P_b) indicates that the buyer should accept the seller's price offer, regardless of whether it is high or low, if the unit's value is high; but should accept only the seller's low price offer if the value is low.

Recall that an equilibrium is identified with the triple (w^*, p_l, p_h). Rustichini and Villamil (1996) characterize each component of this triple. The following results from their analysis are relevant for experimental investigation of the theory. If (w^*, p_l, p_h) describes all equilibria, then $p_l = l$ and $l < p_h \leq h$. Intertemporal price paths are determined as follows. The probability that the unit's value is high (or low) is $1/2$, and this is the seller's limit belief when no new information about the unit's value to the buyer is acquired.[10] The relationship between critical belief w^* and $1/2$ determines the equilibrium price path.[11]

Rustichini and Villamil (1996) derive parametric restrictions on the model which lead to only two types of equilibrium price paths. First, when $\alpha < w^* < 1/2$, equilibrium prices follow a cyclical pattern for any realization of value v. Suppose $w < w^*$, and the seller's equilibrium price offer is p_l. The sequence of future beliefs is given by $\hat{w}^i$, until the first time i_0 that $\hat{w}^{i_0} > w^*$ (where i_0

[9] Appendix B in Rustichini and Villamil (1996) proves existence of incentive compatible equilibria in stationary strategies. The buyer's *equilibrium* strategy sometimes truthfully reveals information to the seller because (P_b) involves an essential trade off: If the buyer accepts price p_h this action reveals information to the seller, *but* gives the buyer an immediate reward for telling the truth (i.e., profit from the trade). If the buyer lies by rejecting p_h when $v = h$, this distorts the seller's belief but "costs" the profit foregone on the rejected trade. Thus, the buyer faces a trade off between current profit and manipulating the seller's beliefs (to obtain higher future profit)—in a sequential game with a random termination date and an oscillating value sequence. Seller equilibrium strategy (P_s) takes this trade off into account.

[10] The $\lim_{i\to\infty} \hat{w}^i = 1/2$ for every w, where $\hat{w}^i$ is the ith iterate of the equilibrium Bayesian belief formation rule $w' = w(1-\alpha) + (1-w)\alpha$. Iteration shows that w' converges to $1/2$ for every w when $0 < \alpha < 1/2$.

[11] Rustichini and Villamil (1996, Section 4) show that if $h > \frac{2-\delta}{1-\delta}$, then $w^* < 1/2$ and equilibrium prices have persistent cycles. This result is essential for experimental tests of the theory, and is consistent with the following intuition: When the seller believes the unit's value is high (and $h > \frac{2-\delta}{1-\delta}$), he/she believes there is a large amount of consumer surplus available for extraction. Thus, the seller posts a high price to increase revenue *and* to acquire information (i.e., learn) about the unit's value.

is finite) when the seller's price offer becomes p_h. If the unit's value is low, the buyer refuses the offer, the seller sets the new belief to α, and the process begins again. If the unit's value is high, the buyer accepts the offer, the seller sets his/her new belief to $1 - \alpha$ and maintains a price offer of p_h (and a belief of $1 - \alpha$) until the unit's value becomes low. When the unit's value becomes low, the buyer rejects the offer, and a new period of low price offers begins. The average length of periods in which the seller makes low price offers is constant and given by:[12]

$$L \equiv \min_i \{\hat{\alpha}^i \geq w^*\}.$$

Thus when the seller's beliefs oscillate about w^*, price cycles are optimal. That is, the model predicts that persistent oscillations between high and low prices will be observed. In equilibrium, the length of the low price phase of the cycle is given by L, and the high price phase of the cycle should persist until p_h is rejected.

Second, when $w^* \geq 1/2$, equilibrium price cycles converge to p_l. Intuitively, suppose that the seller's initial price offer is p_h. If the buyer accepts p_h the seller continues with this price until the first time the price is rejected. The seller then sets $\hat{w}^i$ to α. However, since w' converges to $1/2$ in the limit, it will always be the case that the seller's belief is less than the critical value w^*, and (P_s) indicates that when $w < w^*$ the seller should post p_l. Thus, when $w^* \geq 1/2$ the seller's belief $\hat{w}^i$ can never "build-up" enough to make it optimal to try the high price so the seller always posts p_l (except for perhaps an initial phase).

4 Experimental design and results

The trading rules of the Posted Offer Institution are reported in Ketcham, Smith, and Williams (1984). Our market has a seller, a buyer, a sequence of "trading periods," and one unit of a fictitious good each period. The trading rules specify a "two-step" decision procedure: First, the seller privately makes a price decision and posts a "take-it" or "leave-it" offer. Second, the buyer either accepts or rejects the offer. We assign costs and values for the good each period in accordance with the procedures described in Smith (1976). The seller's cost of producing each unit is known and equal to zero (i.e., $c = 0$). The buyer's value for each unit (v) is either 1 or h in any period, where h is a treatment variable in the experiments. The unit's value is determined randomly each period as follows:

(i) The initial v is drawn from a known equal distribution: the probability it is high (i.e., h) is $1/2$ and the probability it is low (i.e., 1) is $1/2$.

[12] Acceptance of p_h by the buyer perfectly reveals that $v = h$. Further, if the buyer follows *incentive compatible* equilibrium strategy (P_b) the seller also implicitly learns when p_h is rejected that the state has switched from h to l. However, learning is imperfect because there is no way to signal a state change when the seller offers p_l. That is why the seller uses rule L: it tells him/her how many periods to use the (certain to be accepted but uninformative) low price before switching to the (informative but costly–in terms of lost trades) price p_h. Our parameter choices imply $L = 6$ when $h = \$2.50$.

(ii) All subsequent v's are determined by the first order Markov Process: $P(v = 1|v = 1) = P(v = h|v = h) = 1 - \alpha = 0.9$, and $P(v = h|v = 1) = P(v = 1|v = h) = \alpha = 0.1$, where $P(\cdot|\cdot)$ is a conditional probability. Thus, the probability that any subsequent unit's value is the same as last period's is 90 percent (so the probability v has changed is 10 percent).[13]

Cyclical pricing is optimal whenever $h > \frac{2-\delta}{1-\delta}$. In all the experiments we set l = \$1, $\delta = 0.05$, and $\alpha = 0.10$. The level of subject experience and h are treatment variables. It follows immediately that cyclical pricing is optimal for $h >$ \$2.05 and a "flat" pricing strategy is optimal otherwise.

The *value determination rule* was publicly announced, and was induced in the experiment by the following procedure: The experimenter rolled a 20-sided die at the beginning of each market period. In period 1, if the outcome was an 11 through 20 the *first unit's* value was high; otherwise it was low. In periods $2, \ldots, end$, if the outcome was a 1 through 18, the *current unit's* value was the same as last period's value; otherwise it changed. Examples of stochastic processes were shown to subjects in the Instructions. In addition, subjects were told that (i) there would be period-to-period dependence in the unit's value over the course of the experiment; and (ii) on average about half the values would be high and half would be low, if the experiment lasted for many market periods.

Both the seller and the buyer knew the seller's cost (i.e., $c = 0$), and both knew the value determination rule. The unit's current value was privately revealed only to the buyer at the beginning of the period, but was never revealed directly to the seller at any time during the experiment. The fact that agents' had differential information was public knowledge. The termination rule used in all experiments was stochastic: Subjects were told the experiment would last between twenty minutes and three hours because the final market period would be determined by the roll of a 20-sided die at the end of every period. If the outcome was a 1 the experiment would end; otherwise it would continue. This termination procedure corresponds to a discount factor of $\delta = 0.05$. Subjects were undergraduate and graduate students at the University of Illinois at Urbana-Champaign.

We conducted the following series of experiments: Eight experiments with inexperienced subjects and simulated buyers denoted as series (a) experiments. Eight experiments with inexperienced buyers and sellers denoted as series (i) experiments. Eight experiments with once experienced buyers and sellers denoted as series (ii) experiments. Eight experiments with twice experienced buyers and sellers denoted as series (iii) experiments. In the series (a) experiments h was a treatment variable. Six experiments were conducted in which $h > 2.05$ thus cyclical pricing was predicted by the theory and two experiments were conducted in which h was less than this amount and thus "flat" pricing was predicted. In the series (i), (ii) and (iii) experiments experience was a treatment variable. Thus, each subject participated in a sequence of three experiments with the same parameters but different counterparts (and of course, different realizations of the

[13] In general, when α is small there is persistence in the process so information about v is valuable to the seller.

Markov process and termination period). The experience profiles of all subjects are reported in Appendix A.

The series (a) experiments were designed to investigate the seller's pricing strategy when the seller knew the buyer would always follow equilibrium strategy (P_b). Thus, in these experiments the seller was told that the buyer would accept the seller's price offer (regardless of whether it was high or low) if the value was high, but would accept only the low price if the value was low. These experiments are designed to establish baseline empirical support for pricing policy (P_S) when there is no strategic uncertainty about the buyer's behavior. In the first 6 experiments in this series parameters were chosen so that the theory predicts equilibrium price cycles.[14] In the remaining 2 experiments in this series parameters are chosen so that the theory predicts equilibrium prices which converge to p_l.[15] These latter two experiments are an important check on the model's theoretical consistency, though this parametric case is unlikely to be relevant in most naturally occurring markets.

The series (i), (ii), and (iii) experiments focus on cyclical pricing when there is strategic uncertainty about the behavior of one's counterpart. That is, both the seller and the buyer were human subjects with no information about how their counterpart would behave other than the information contained in the Instructions and their experience in the current and previous (if any) markets. In all of these experiments, h = \$2.50.

Tables 1 and 2 report summary statistics for each series. Specifically, the Tables report *measures of predictive success* (cf., Selten (1991)) for testable implications of the theory. The theory predicts a "high or low" seller pricing strategy (P_s), and a "yes" or "no" buyer answer strategy (P_b). The measure of predictive success assesses the statistical usefulness of an area theory.[16] Let m denote the measure of success, r denote the *hit rate* of the theory (the relative frequency of correct predictions), and a denote the area (relative size) of the predicted subset compared with the set of all possible outcomes. Define $m = r - a$, where $m \in [-1, +1]$ in general.

Figures 1a–8a and Figures 1-24 report price and answer data for the 32 experiments. Figures 1a–8a report the results of the eight series (a) experiments with simulated buyers; Figures 1-8 report the results of the eight series (i) experiments with inexperienced buyers and sellers; Figures 9-16 report the results of the eight series (ii) experiments with once experienced buyers and sellers; and Figures 17-24 report the results of the eight series (iii) experiments with twice experienced buyers and sellers.

Table 1 reports measures of predictive success (i.e., m's) for the eight series (a) experiments with simulated buyers. In Experiment 1a measure m for (P_s) is computed as follows: $r = 32/37$, where 32 is the number of successful predictions

[14] Price cycles occur when $h > \frac{2-\delta}{1-\delta}$, or $h >$ \$2.05 given $\delta = 0.05$. We set h = \$2.50 in Experiment 1a and h = \$2.20 in Experiments 2a–6a (see Figures 1a–6a).

[15] p_l is optimal when $h < 2 - \delta$. We set h = \$1.50 in these experiments (see Figures 7a and 8a).

[16] Area theories for the prediction of experimental results delineate regions of predicted outcomes within the set of all possible outcomes.

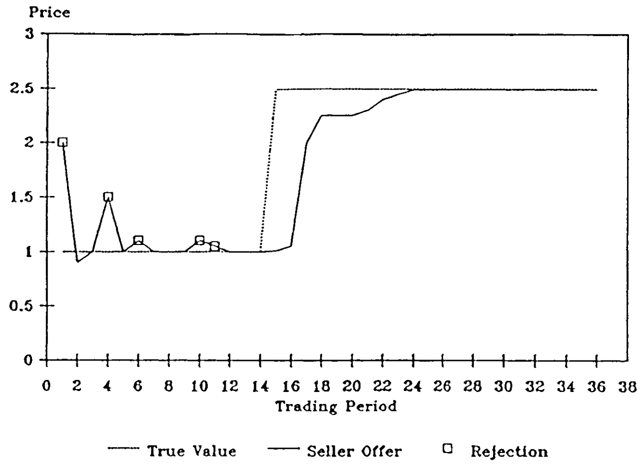

Figure 1a

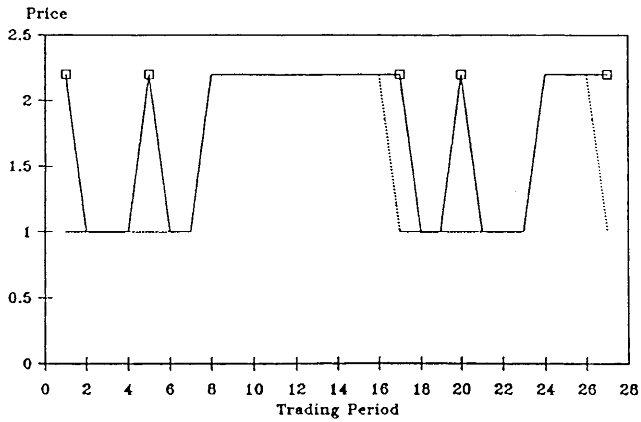

Figure 2a

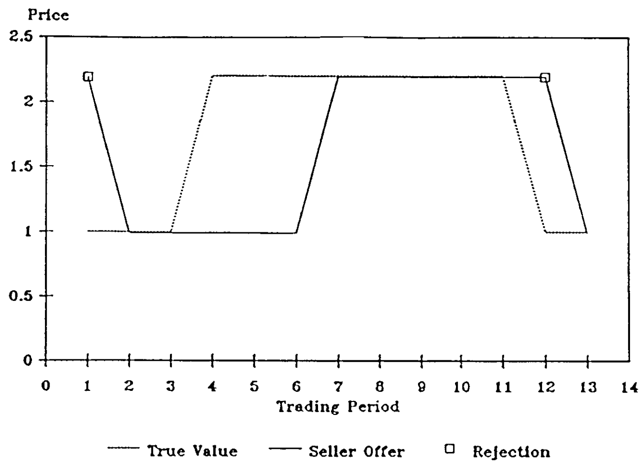

Figure 3a

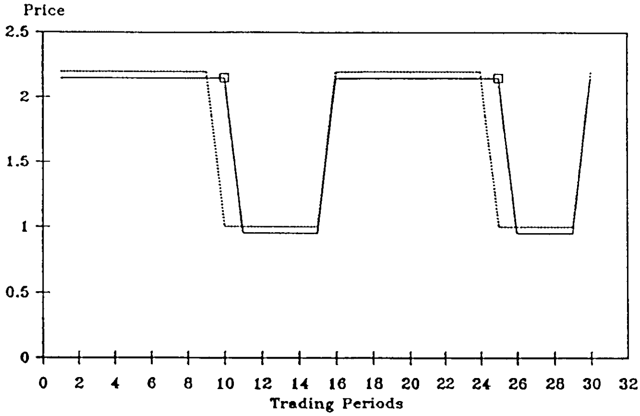

Figure 4a

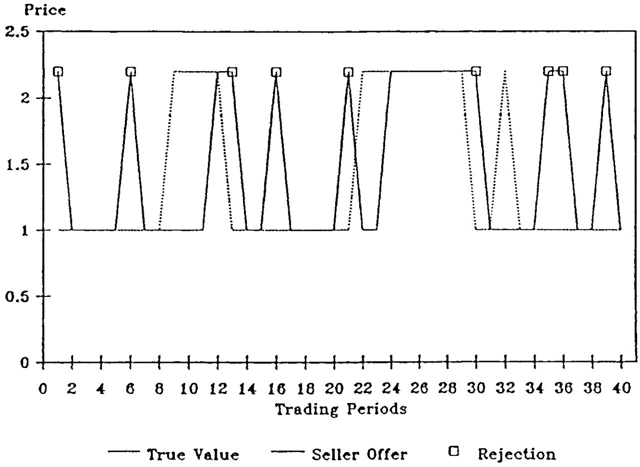

Figure 5a

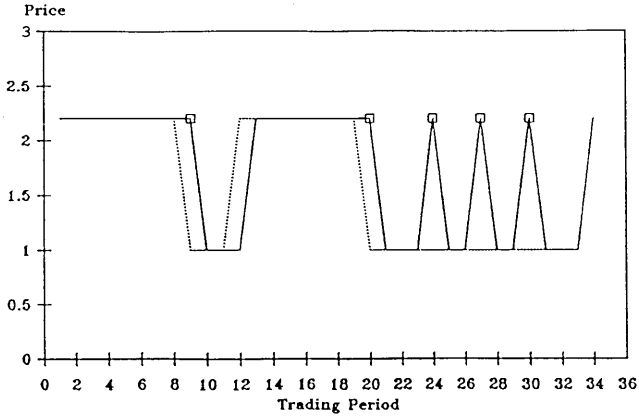

Figure 6a

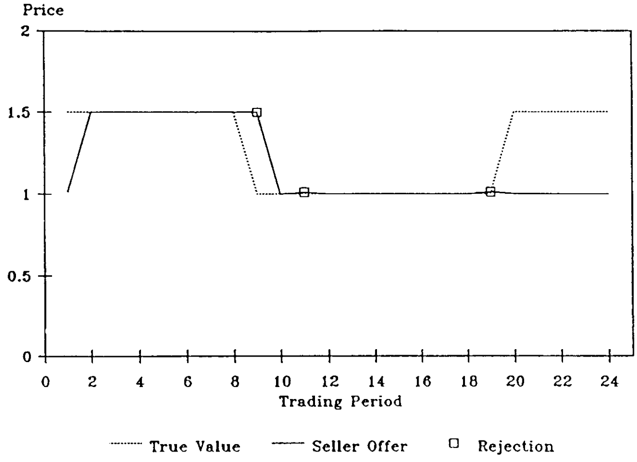

Figure 7a

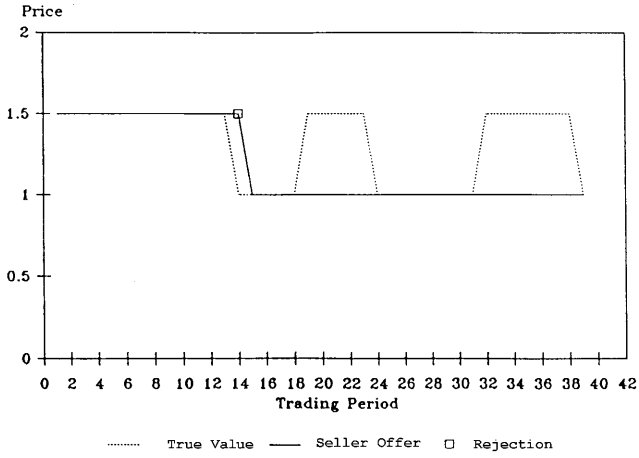

Figure 8a

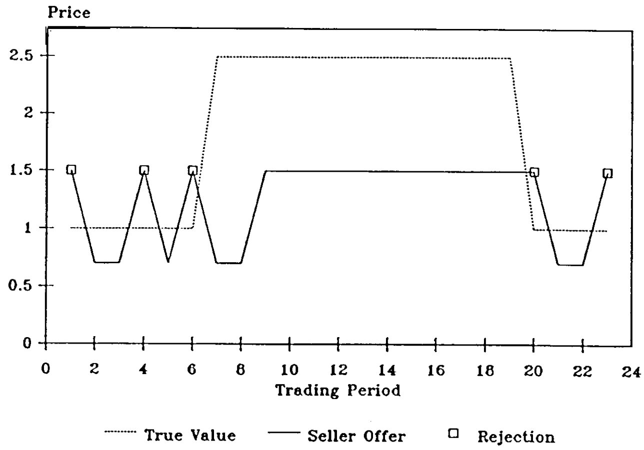

Figure 1

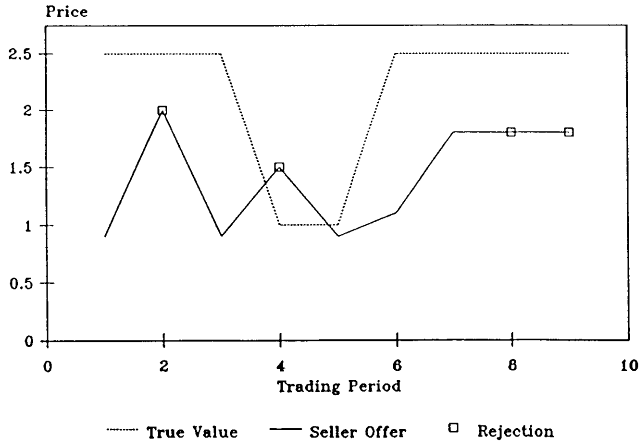

Figure 2

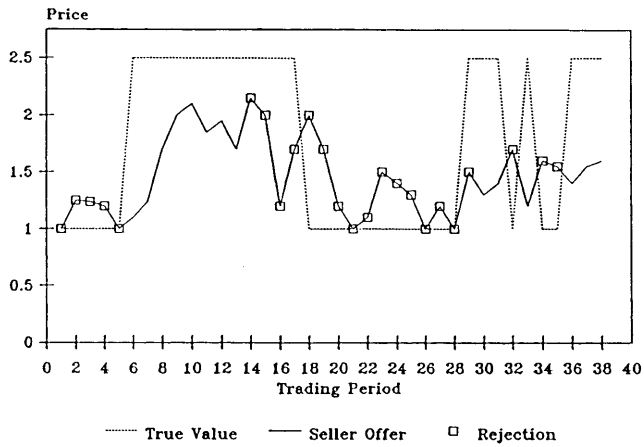

Figure 3

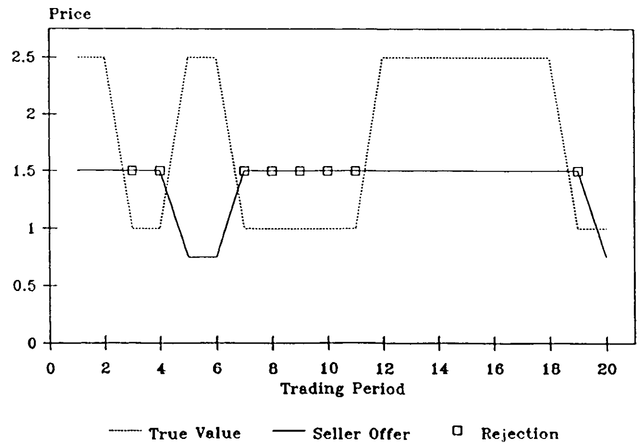

Figure 4

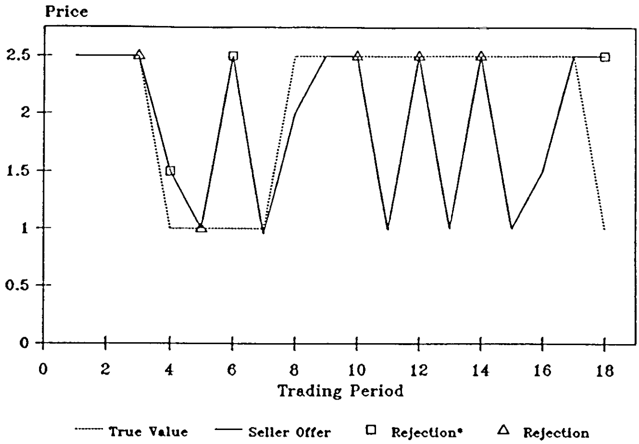

Figure 5

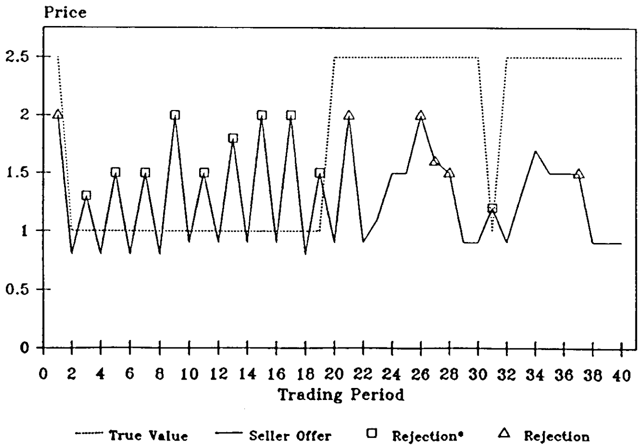

Figure 6

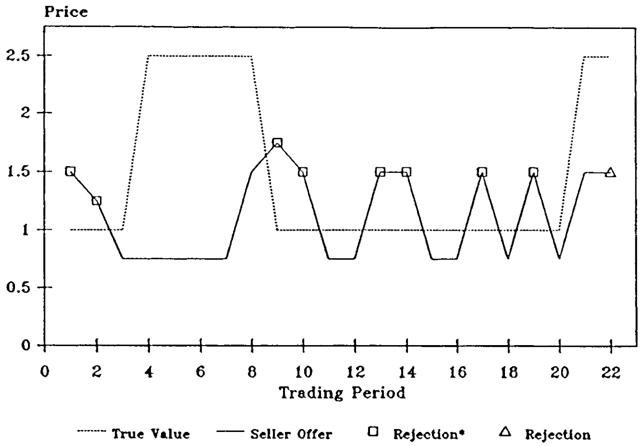

Figure 7

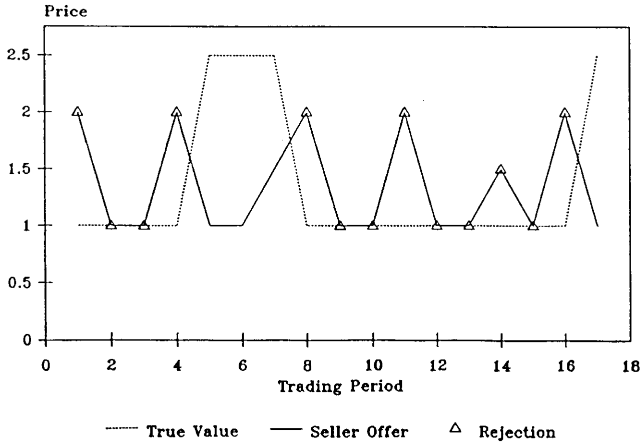

Figure 8

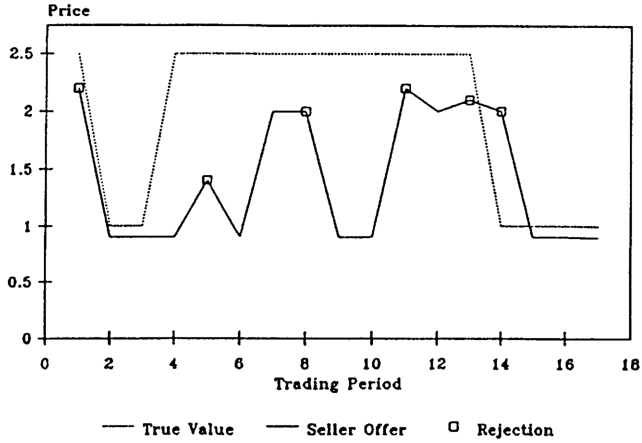

Figure 9

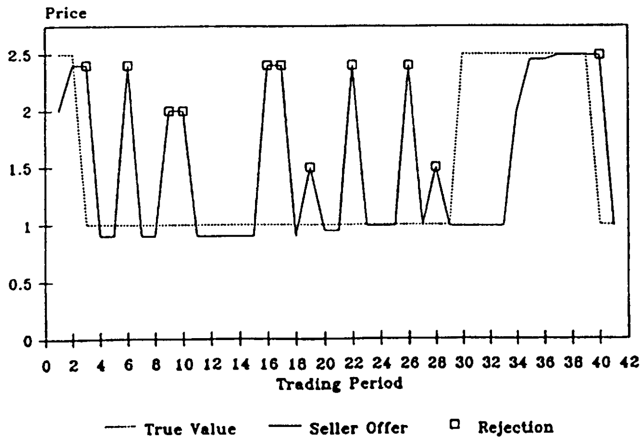

Figure 10

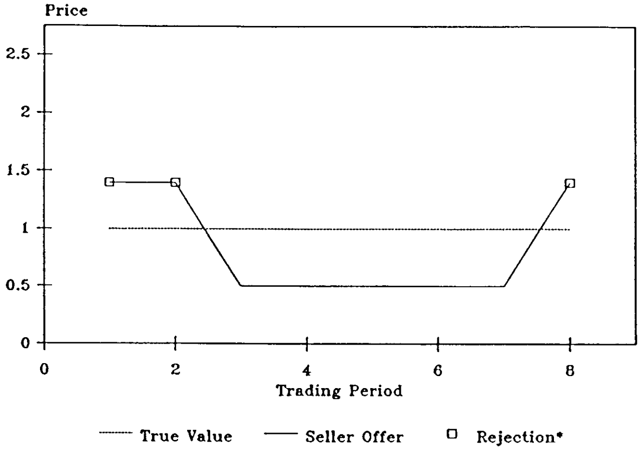

Figure 11

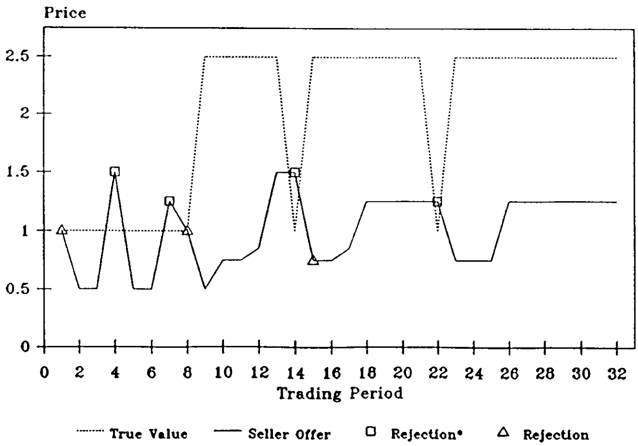

Figure 12

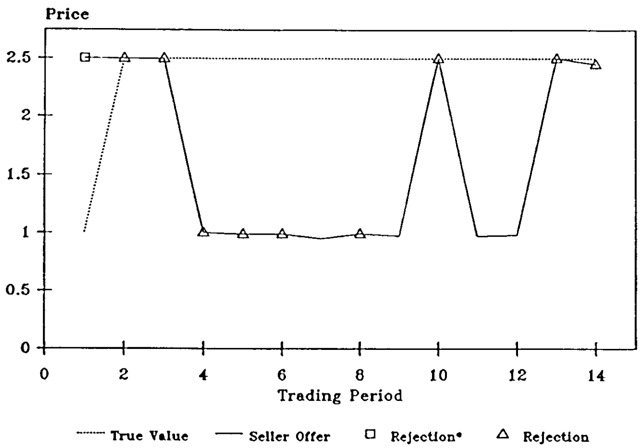

Figure 13

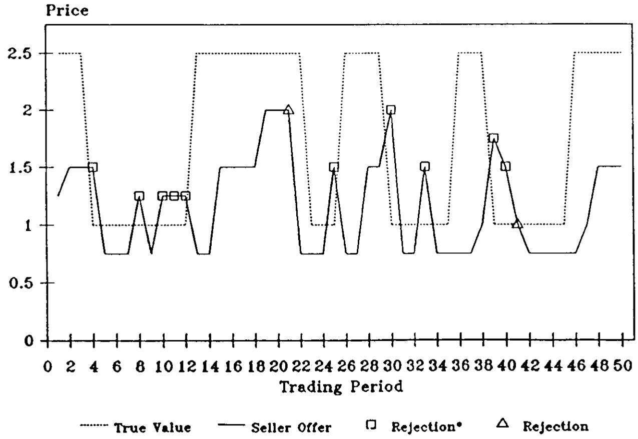

Figure 14

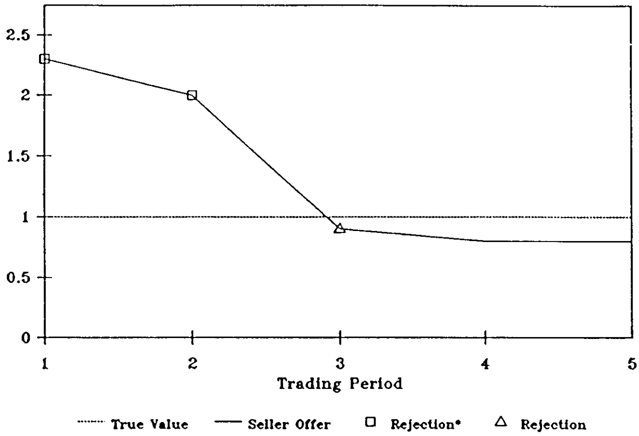

Figure 15

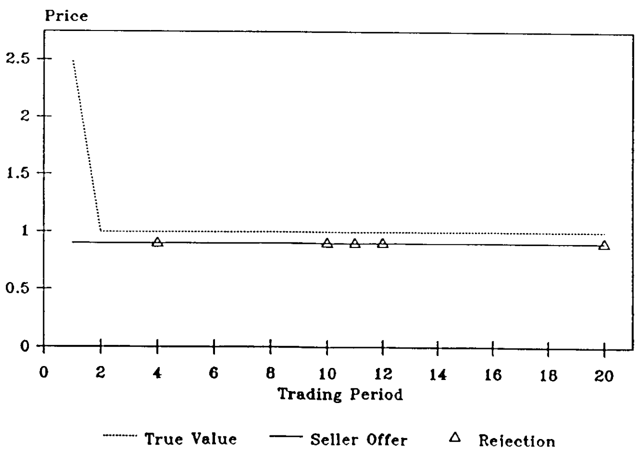

Figure 16

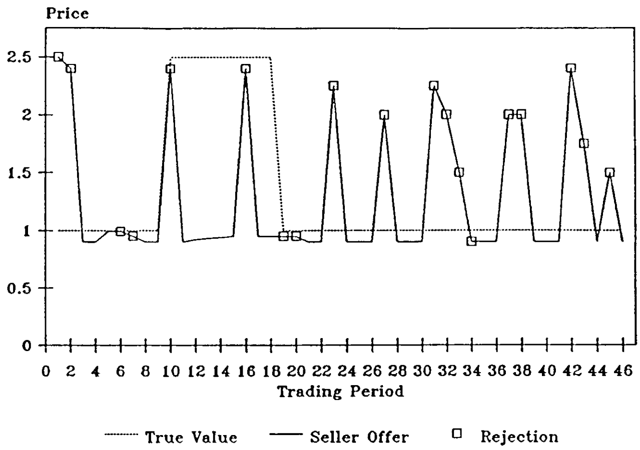

Figure 17

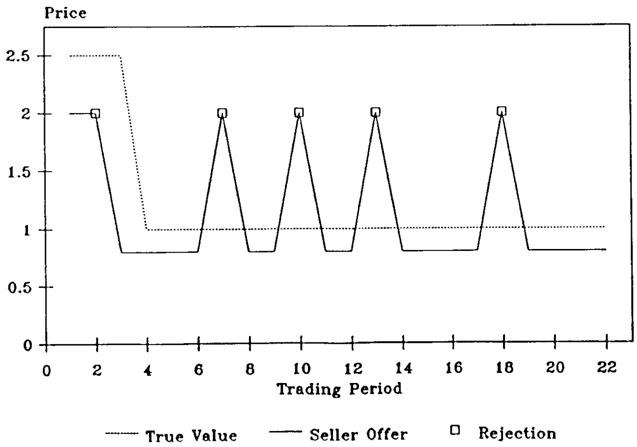

Figure 18

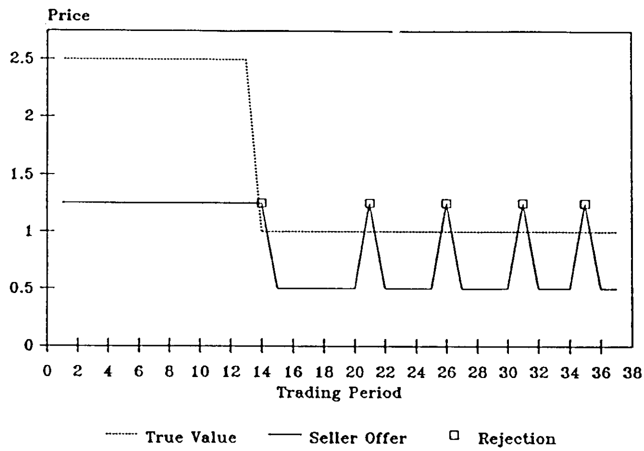

Figure 19

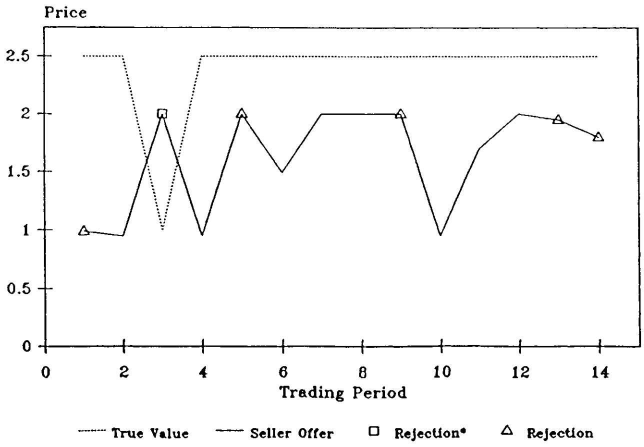

Figure 20

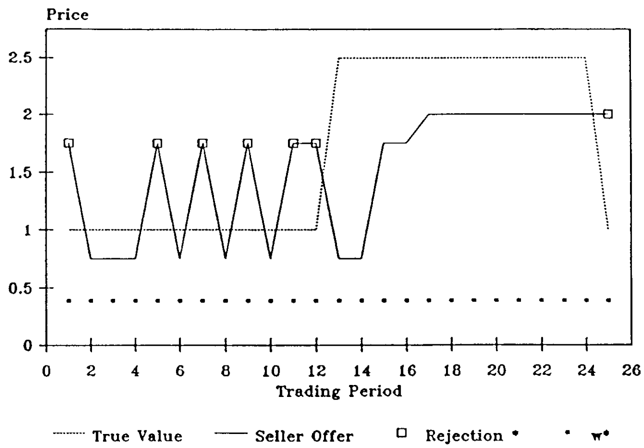

Figure 21

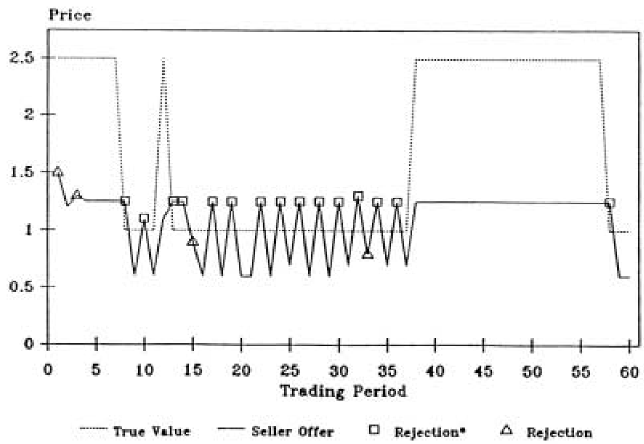

Figure 22

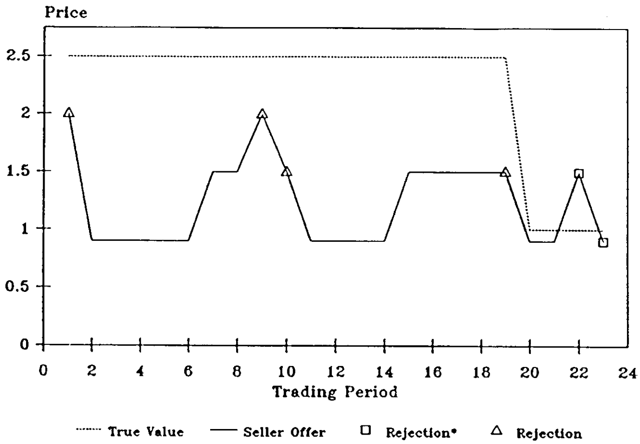

Figure 23

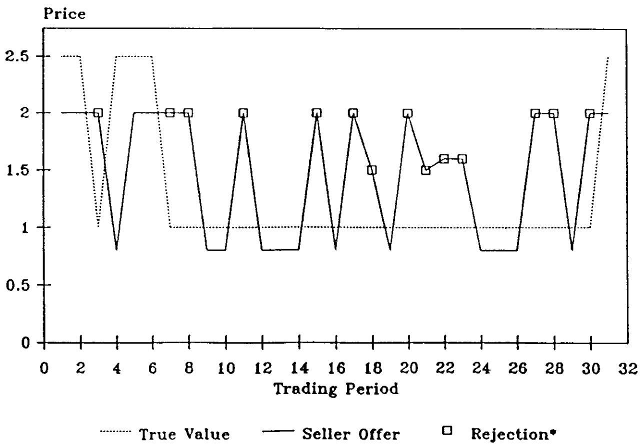

Figure 24

and 37 is the total number of outcomes; and $a = 150/251$, where 150 is the number of possible high prices $p_h \in [1.01, 2.50]$ and 251 is the total number of prices in $[0, 2.50]$. For (P_s) the theory's hit rate is high but the accuracy is low. When h = \$2.50 (Experiment 1a) the maximum possible m for (P_s) is 0.30. When h = \$2.20 (Experiments 2a, 3a, 4a, 5a, 6a) the maximum possible m for (P_s) is 0.46. When h = \$1.50 (Experiments 7a, 8a) the maximum possible m is 0.67. Most m in Table 1 are close to the maximum possible value. Since buyer behavior was simulated, m for (P_b) is not applicable. This is indicated by *NA* in column (P_b) of Table 1.

Table 2 reports summary statistics for the twenty four series (i), (ii), and (iii) experiments. For example in Experiment 1, $m = r - a = 0.27$ for (P_s), which is computed as follows: $r = 20/23 = 0.87$, where 20 is the number of

Table 1

Experiment	P_s	P_b	max m_s
1a	0.26	NA	0.30
2a	0.31	NA	0.46
3a	0.37	NA	0.46
4a	0.39	NA	0.46
5a	0.23	NA	0.46
6a	0.31	NA	0.46
7a	0.59	NA	0.67
8a	0.67	NA	0.67

Table 2

Series i			Series ii			Series iii		
Experiment	P_s	P_b	Experiment	P_s	P_b	Experiment	P_s	P_b
1	0.22	0.50	9	0.16	0.21	17	0.18	0.35
2	0.07	0.17	10	0.20	0.50	18	0.22	0.45
3	−0.13	0.24*	11	0.27	0.50	19	0.32	0.50
4	0.05	0.50	12	0.24	0.41	20	0.04	0.14
5	0.07	0.22*	13	0.11	−0.14*	21	0.16	0.50
6	0.02	0.35	14	0.22	0.46	22	0.13	0.43*
7	0.08	0.45	15	0.20	0.30	23	0.27	0.28
8	0.10	0.09*	16	0.40	0.25	24	−0.05	0.05
max m	0.40	0.50	max m	0.40	0.50	max m	0.40	0.50
$\mu_s^i = 0.06$	$\mu_b^i = 0.315$		$\mu_s^{ii} = 0.23$	$\mu_b^{ii} = 0.311$		$\mu_s^{iii} = 0.17$	$\mu_b^{iii} = 0.450$	

*Higher if zero-π trades are excluded.

successful predictions and 23 is the total number of outcomes (see Figure 1); and $a = 150/251 = 0.60$, where 150 is the number of $p_h \in [1.01, 2.50]$ and 251 is the total number of prices in $[0, 2.51]$. Since $h = \$2.50$ in all experiments, the maximum hit rate of (P_s) is high ($r = 1$ if the seller follows (P_s) perfectly) but the accuracy is low ($a = .6$). Thus in all 24 experiments the maximum possible m for (P_s) is 0.40 in Table 1. The maximum possible m for (P_b) is 0.50 since $r = 1$ if the buyer follows (P_b) perfectly and $a = .5$ since there are two possible answers (yes or no).

In the series (i), (ii), and (iii) experiments subjects have different levels of experience in each series. Thus, Table 2 reports the mean measure of predictive success for the eight experiments in the series i, ii, iii for the seller and the buyer, respectively. These means are denoted by μ_r^e, where $e = i, ii, iii$ denotes the experiment series and $r = b, s$ denotes buyer or seller. In the series (i), (ii) and (iii) experiments, $\mu_s^e = 0.06$, 0.23, and 0.17, for the seller in $e = i, ii, iii$. These means are all clearly below the maximum possible m. Most interestingly, however, is the fact that the means increase between the series (i) and (ii) experiments, which we interpret as evidence that sellers are learning (P_s), *but* decrease between the series (ii) and (iii) experiments.

The results in Table 2 suggest that experience affects sellers and buyers differently. Our a priori conjecture was that it would make subjects more likely to behave in accordance with the equilibrium strategies. The data indicate that buyers use equilibrium strategy (P_b) more frequently when they are more experienced: the maximum possible m was 0.50 and the respective mean measures of predictive success for (P_b) were μ_b^e = 0.315, 0.311, and 0.45 for $e = i, ii, iii$. Observe that by the time buyers were twice experienced with the *institution* they were remarkably close to the maximum measure of predictive success for (P_b) (i.e., m = 0.50). However, seller's use of (P_s) declined in series (ii) to (iii) experiments: μ_s decreases from 0.23 to 0.17 and the variance of (P_s) increases from σ_s^{ii} = 0.086 in series ii to σ_s^{iii} = 0.127 in series iii. The variance of (P_s) in the series i experiments was σ_s^i = 0.097. The variance of (P_b) in the experiments was σ_b^{ii} = 0.214 and σ_b^{iii} = 0.132.

The increase in the variance of (P_s) suggests that there may have been something peculiar about some individual series (iii) experiments. Table 2 indicates that the measures of predictive success for the seller in each experiment in this series are m_s^{17} = 0.18 in Experiment 17, m_s^{18} = 0.22 in Experiment 18, m_s^{19} = 0.32 in Experiment 19, m_s^{20} = 0.04 in Experiment 20, m_s^{21} = 0.16 in Experiment 21, m_s^{22} = 0.13 in Experiment 22, m_s^{23} = 0.27 in Experiment 23, and m_s^{24} = −0.05 in Experiment 24. The maximum m_s = 0.40. Experiments 20 and 24 account for most of the increased variance, thus we now discuss the four series *iii* experiments with the worst seller performance in detail, with special emphasis on these two high variance cases.

In Experiment 20, buyer B20 (i.e., Subject 4 with experience profile S2, B12, B20: see Appendix A) frequently deviated from equilibrium strategy (P_b), resulting in m_b^{20} = 0.14. This buyer behavior is interesting in view of the fact that in Experiment 12 this individual was again a buyer (B12) with m_b^{ii} = 0.41, indicating that he consistently followed equilibrium strategy (P_b). Indeed in Experiment 12 this buyer had only 3 non-equilibrium rejections, and 2 of these 3 rejections were for zero-profit trades. Experiment 12 lasted for 32 periods. In contrast, in Experiment 20 this buyer had 5 non-equilibrium rejections in 14 periods, and only 1 of the 5 was a zero profit trade. The main difference between the buyer's behavior in Experiments 12 and 20 appears to be that in Experiment 12 the seller never offered a price higher than \$1.50 and by the end of the experiment p_h was consistently \$1.25. However, in Experiment 20 the seller consistently offered prices above \$1.50 (often \$2.00) and the buyer appeared unwilling to accept such a skewed profit split. The buyer may have been "punishing" the seller (and himself) in the hope of attaining a lower price for future trades. The seller was very slow in discerning this implicit lesson in the buyer's strategy, despite the fact that he had been a seller in Experiment 3. Although the measures of predictive success in Experiments 3 and 20 did show improvement (i.e., m_s^3 = −0.13 vs. m_s^{20} = 0.04), this subject had persistent trouble with the seller role but not with the buyer role (i.e., this subject was a buyer in Experiment 11 and m_b^{11} = 0.50).

In Experiment 21 the buyer followed equilibrium strategy (P_b) perfectly. This subject had two previous experiences as a seller (S5 and S13). The experiment lasted for 25 periods. The seller, however, had been a buyer twice previously (B8 and B14). This subject displayed no understanding of L, the optimal number of low price offers, and this is reflected in $m_s^{21} = 0.16$. This experiment suggests to us that S21's lack of prior experience with the seller role may have been more significant than B21's lack of prior experience with the buyer role, perhaps because buyers are much more passive in the Posted Offer institution than sellers.

In Experiment 22 the buyer followed equilibrium strategy (P_b) quite closely, and had previous experience as a buyer and seller (B7 and S14). The experiment lasted for 60 periods. The seller had been a seller once before (S6), and a buyer once before (B13). In Experiment 13 the buyer had an unusually long string of high values (13 out of 14 units). When this subject became a seller again in Experiment 22, he chose a relatively low p_h of \$1.25 (apparently to induce the buyer to tell the truth in the high state) but displayed no understanding of L. Perhaps his experience as a buyer in Experiment 13 (when he saw an unusually long string of high values) impeded his learning of L and resulted in $m_s^{22} = 0.13$.

In Experiment 24 the buyer followed equilibrium strategy (P_b) perfectly, and had two previous seller experiences (S7 and S16). The experiment lasted for 32 periods. The seller had been both a buyer and a seller previously (B5 and S15), however Experiment 15 lasted for only 5 periods so this subject's experience with the seller role was very limited. Interestingly in Experiment 15, $m_s^{15} = 0.20$ and in Experiment 24 $m_s^{24} = -0.05$ (for this same subject). Thus, this seller appears to have "unlearned" from his brief seller experience! In fact, closer inspection of Figure 24 reveals that a combination of an unusually large number of low realizations and imperfect comprehension of L (perhaps due to limited experience with the seller role) accounts for the extremely low $m_s^{24} = -0.05$.

The detailed search for anomalies in Experiments 20, 21, 22, and 24 indicates a common factor — experience with the buyer/seller role (not just with the institution) appears to be important when demand evolves according to an exogenous stochastic process with serial correlation. That is, whether a subject has been a buyer or a seller, and his/her particular experiences in the role are important when subjects are differentially informed and demand flutuates with some predictability. Subjects must not only attempt to learn the demand process, a seller must first form an opinion as to whether or not the buyer is revealing information about v truthfully.

Unusual prior experiences seem to affect sellers more than buyers in these experiments. Clearly, the buyer's role is easier to learn in the Posted Offer institution (i.e., the buyer simply accepts or rejects the price posted by the seller). In contrast, the seller must both choose a price *and* choose the optimal number of periods in which to make the low price offer (L). This dual task reflects the seller's more complicated problem: The seller is using price to both earn current profit on a trade and to learn the buyer's demand process (in the hope of making higher future profit). However, because sellers are informationally disadvantaged and realize that buyers may reject mutually profitable current trades in order to manipulate the seller's beliefs, subjects must learn about the institution, their own role, *and* their counterpart's role. This observation suggests that in future experiments it may be desirable to control for experience with only a particular role. That is, conduct a series of experiments where subjects are either buyers or sellers in all three experiments. These results can be compared with the results reported in Appendix A where role assignment was determined randomly over the course of the three experiments.

5 Concluding remarks

This paper studies intertemporal pricing and purchasing strategies in Posted Offer laboratory markets with demand values that follow a Markov Process. The following features of the model are important: First, differential information reduces market efficiency relative to full information when cyclical pricing is optimal because agents systematically forego trade on units for which the seller posts p_h to learn about the unit's current value when the value is low. Second, the buyer prefers the differential information (price cycle) solution to the full information solution because in general it allows the buyer to obtain some exchange surplus. Third, *in equilibrium information is revealed truthfully* but the equilibrium is *not* fully revealing (i.e., when the seller offers p_l the buyer's answer is uninformative).

The results from 32 laboratory experiments indicate that the Posted Offer institution performs reasonably well in revealing information in stochastic, intertemporal settings despite some pronounced inefficiencies (relative to full information) that are inherent in its structure. In particular, pricing patterns that appear to be "sticky," as well as periodic forgone trades, are part of an optimal intertemporal equilibrium pricing strategy. Since intertemporal pricing patterns are inextricably connected to agents' beliefs about the underlying nature of uncertainty that governs the system, laboratory experiments are useful in this setting because they allow researchers to observe and control the probability structures on which agents' actions are based. Thus, they may prove to be a useful tool for testing intertemporal, stochastic theories. However, the our results indicate that the nature of subject experience (i.e., with the institution versus with the role) matters much more in this setting than in previous static, deterministic Posted Offer experiments.

6 Appendix A: Experience profiles

All subjects in the series (a) experiments summarized in Table 1 were inexperienced. That is, they had never before participated in a Posted Offer experiment with stochastic demand. The experience profiles of the 16 subjects in the series (i), (ii) and (iii) experiments are given below. Subject 1 was the inexperienced buyer in experiment 1 (i.e., B1), the once experienced seller in experiment 11 (i.e. S11) and the twice experienced seller in experiment 19 (i.e., S19). After each role indicator, the associated measure of predictive success for that subject from Table 2 is reported. This provides a very rough profile of the subjects' learning over the course of the experiment. Following this notation, the experience and performance profiles for each subject are:

Subject experience and performance profile

Subject 01: B1 (m_1^b = .50), S11 (m_{11}^s = .27), S19 (m_{19}^s = .32).
Subject 02: S1 (m_1^s = .22), B10 (m_{10}^b = .50), S18 (m_{18}^s = .22).
Subject 03: B2 (m_2^b = .17), S09 (m_9^s = .16), B17 (m_{17}^b = .35).
Subject 04: S2 (m_2^s = .07), B12 (m_{12}^b = .41), B20 (m_{20}^b = .14).
Subject 05: B3 (m_3^b = .24), S12 (m_{12}^s = .24), B19 (m_{19}^b = .50).
Subject 06: S3 (m_3^s = −.13), B11 ($m_1 1^b$ = .50), S20 (m_{20}^s = .04).
Subject 07: B4 (m_4^b = .50), S10 (m_{10}^s = .20), S17 (m_{17}^s = .18).
Subject 08: S4 (m_4^s = .05), B09 (m_9^b = .21), B18 (m_{18}^b = .45).
Subject 09: B5 (m_5^b = .22), S15 (m_{15}^s = .22), S24 (m_{24}^s = −.05).
Subject 10: S5 (m_5^s = .07), S13 (m_{13}^s = .11), B21 (m_{21}^b = −.05).
Subject 11: B6 (m_6^b = .35), B15 (m_{15}^b = .30), B23 (m_{23}^b = .28).
Subject 12: S6 (m_6^s = .02), B13 (m_{13}^b = −.14), S22 (m_{22}^s = .13).
Subject 13: B7 (m_7^b = .45), S14 (m_{14}^s = .22), B22 (m_{22}^b = .43).
Subject 14: S7 (m_7^s = .08), S16 (m_{16}^s = .40), B24 (m_{24}^b = .50).
Subject 15: B8 (m_8^b = .09), B14 (m_{14}^b = .46), S21 (m_{21}^s = .16).
Subject 16: S8 (m_8^s = .10), B16 (m_{16}^b = .25), S23 (m_{23}^s = .27).

References

Davis, D., Harrison, G., Williamson, A.: Convergence to nonstationary competitive equilibria: an experimental analysis. Discussion Paper (1992)

Davis, D., Holt, C.A.: Price rigidities and institutional variations in markets with posted prices. Economic Theory **9**(1): 63–80 (1996)

Davis, D., Holt, C.A.: Conspiracies and secret price discounts. Economic Journal **108**: 1–21 (1998)

Hoffman, E., Plott, C.: The effect of intertemporal speculation on the outcomes in seller posted-offer auction markets. Quarterly Journal of Economics **96**: 223–241 (1981)

Ketcham, J., Smith, V.L., Williams, A.: A comparison of posted-offer and double auction pricing institutions. Review of Economic Studies **51**: 595–614 (1984)

Plott, C.: An updated review of industrial organization: applications of experimental methods. In: Schmalensee, R., Willig, R.D. (eds.) Handbook of industrial organization, Vol. II. Amsterdam: Elsevier Publishers, North Holland 1989

Rustichini, A., Villamil, A.P.: Intertemporal pricing in markets with differential information. Economic Theory **8**: 211–227 (1996)

Selten, R.: Properties of a measure of predictive success. Mathematical Social Sciences **21**: 153–167 (1991)

Smith, V.L.: Experimental economics: induced value theory. American Economic Review **66**: 274–279 (1976)

Smith, V.L.: Markets as economizers of information: experimental examination of the 'Hayak hypothesis'. Economic Inquiry **20**: 165–179 (1982)

Villamil, A.P.: Termination procedures in laboratory markets. University of Illinois Discussion Paper (1999)