FACULTY WORKING PAPER 91-0162
Political Economy Series #52

Equilibrium Price Stickiness

Aldo Rustichini
Department of Economics
Northwestern University

Anne P. Villamil
Department of Economics
University of Illinois

Bureau of Economic and Business Research
College of Commerce and Business Administration
University of Illinois at Urbana-Champaign

Papers in the Political Economy of Institutions Series

No. 3 George E. Monahan and Vijay K. Vemuri. "Monotonicity of Second-Best Optimal Contracts" Working Paper #1417

No. 4 Charles D. Kolstad, Gary V. Johnson, and Thomas S. Ulen. "*Ex Post* Liability for Harm vs. *Ex Ante* Safety Regulation: Substitutes or Complements?" Working Paper #1419

No. 5 Lanny Arvan and Hadi S. Esfahani. "A Model of Efficiency Wages as a Signal of Firm Value" Working Paper #1424

No. 6 Kalyan Chatterjee and Larry Samuelson. "Perfect Equilibria in Simultaneous-Offers Bargaining" Working Paper #1425

No. 7 Jan K. Brueckner and Kangoh Lee. "Economies of Scope and Multiproduct Clubs" Working Paper #1428

No. 8 Pablo T. Spiller. "Politicians, Interest Groups, and Regulators: A Multiple-Principals Agency Theory of Regulation (or "Let Them Be Bribed" Working Paper #1436

No. 9 Bhaskar Chakravorti. "Asymmetric Information, 'Interim' Equilibrium and Mechanism Design" Working Paper #1437

No. 10 Bhaskar Chakravorti. "Mechanisms with No Regret: Welfare Economics and Information Reconsidered" Working Paper #1438

No. 11 Bhaskar Chakravorti. "Communication Requirements and Strategic Mechanisms for Market Organization" Working Paper #1439

No. 12 Susan I. Cohen and Martin Loeb. "On the Optimality of Incentive Contracts in the Presence of Joint Costs" Working Paper #1445

No. 13 Susan I. Cohen and Martin Loeb. "The Demand for Cost Allocations: The Case of Incentive Contracts Versus Fixed-Price Contracts" Working Paper #1455

No. 14 Jan K. Brueckner and Kevin M. O'Brien. "Modeling Government Behavior in Collective Bargaining: A Test for Self-Interested Bureaucrats" Working Paper #1481

No. 15 Jan K. Brueckner. "Estimating a Bargaining Contract Curve: Prior Restrictions and Methodology" Working Paper #1490

No. 16 Peter C. Reiss and Pablo T. Spiller. "Competiton and Entry in Small Airline Markets" Working Paper #1497

No. 17 Pablo T. Spiller. "A Note on Pricing of Hub-and-Spoke Networks" Working Paper #1498

No. 18 Larry DeBrock. "Joint Marketing Efforts and Pricing Behavior" Working Paper #1500

No. 19 Frank A. Wolak and Charles D. Kolstad. "A Model of Homogenous Input Demand Under Price Uncertainty" Working Paper #1502

No. 20 Susan I. Cohen. "Reputation, Intertemporal Incentives and Contracting" Working Paper #1511

No. 21 Lanny Arvan and Antonio Leite. "A Sequential Equilibrium Model of Cost Overruns in Long Term Projects" Working paper #1514

No. 22 Jan K. Brueckner and Pablo T. Spiller. "Competiton and Mergers in Airline Networks" Working Paper #1523

No. 23 Hadi S. Esfahani. "Reputation, Product Quality, and Production Technology in LDC Markets" Working Paper #89-1525

No. 24 Hadi S. Esfahani. "Moral Hazard, Limited Entry Costs, and 'Introductory Offers'" Working Paper #89-1526

No. 25 Bhaskar Chakravorti. "Mechanisms with No Regret: Welfare Economics and Information Reconsidered" Working Paper #89-1527

No. 26 Susan I. Cohen. "Implicit Cost Allocation and Bidding for Contracts" Working Paper #89-1558

No. 27 Rafael Gely and Pablo T. Spiller. "A Rational Choice Theory of the Supreme Court" Working Paper #89-1559

No. 28 Rafael Gely and Pablo T. Spiller. "An Economic Theory of Supreme Court Statutory Decisions The *State Farm* and *Grove City* Cases" Working Paper #89-1560

No. 29 Rafael Gely and Pablo T. Spiller. "The Political Economy of Supreme Court Constitutional Decisions: The Case of Roosevelt's Court Packing Plan" Working Paper #89-1561

No. 30 Hadi S. Esfahani. "Reputation and Product Quality Revisited." Working Paper #89-1584

No. 31 Jan K. Brueckner. "Growth Controls and Land Values in an Open City." Working Paper #89-1594

No. 32 Jan K. Brueckner. "Tastes, Skills, and Local Public Goods." Working Paper #89-1610

No. 33 Luis Cabral and Shane Greenstein. "Switching Costs and Bidding Parity in Government Procurement of Computer Systems." Working Paper #90-1628

No. 34 Charles D. Kolstad. "Hotelling Rents in Hotelling Space: Exhaustible Resource Rents with Product Differentiation." Working Paper #90-1629

No. 35 Santiago Urbiztondo. "Investment Without Regulatory Commitment: The Case of Elastic Demand." Working Paper #90-1634

No. 36 Bhaskar Chakravorti. "Sequential Rationality, Implementation and Communication in Games." Working Paper #90-1636

No. 37 Pablo T. Spiller, Rafael Gely. "Congressional Control of Judicial Independence: The Determinants of US Supreme Court Labor Relations Decisions, 1949/1987." Working Paper #90-1637

No. 38 Hans Brems. "Dynamic Macroeconomics: Fiscal and Monetary Policy." Working Paper #90-1640

No. 39 Lanny Arvan. "Flexibility Versus Commitment in Strategic Trade Policy Under Uncertainty: A Model of Endogenous Policy Leadership." Working Paper #90-1651

No. 40 David T. Scheffman, Pablo T. Spiller. "Buyers' Strategies, Entry Barriers, and Competiton." Working Paper #90-1674

No. 41 Richard Arnould and Larry DeBrock. "Utilization Control in HMOs." Working Paper #90-1698

No. 42 Shane Greenstein. "Did Installed Base Give an Incumbent Any (Measurable) Advantages in Federal Computer Procurement?" Working Paper #90-1718

No. 43 Bhaskar Chakravorti and Charles M. Kahn. "Universal Coalition-Proof Equilibrium" Working Paper #91-0100

No. 44 Pablo T. Spiller. "A Rational Choice Theory of *Certiorari*: Hierarchy, Strategy and Decision Costs at the Courts" Working Paper #91-0110

No. 45 Pablo T. Spiller. "Agency Discretion Under Judicial Review" Working Paper #91-0111

No. 46 Charles D. Kolstad and Michelle H.L. Turnovsky. "Production with Quality Differentiated Inputs" Working Paper #91-0118

No. 47 Pablo T. Spiller and Santiago Urbiztondo. "Political Appointees vs. Career Civil Servants: A Multiple Principals Theory of Political Bureaucracies" Working Paper #91-0129

No. 48 Shane M. Greenstein. "Lock-In and the Costs of Switching Mainframe Computer Vendors: What Do Buyers See?" Working Paper #91-0133

No. 49 Pablo T. Spiller and Matthew L. Spitzer. "Judicial Choice of Legal Doctrines" Working Paper #91-0134

No. 50 Bruce D. Smith and Anne P. Villamil. "Government Borrowing Using Bonds with Randomly Determined Returns: Welfare Improving Randomization in the Context of Deficit Finance" Working Paper #91-0148

No. 51 Pablo T. Spiller and Santiago Urbiztondo. "Interest Groups and the Control of the Bureaucracy: An Agency Perspective on the Administrative Procedure Act" Working Paper #91-0152

BEBR

FACULTY WORKING PAPER NO. 91-0162

Papers in the Political Economy of Institutions Series No. 52

College of Commerce and Business Administration

University of Illinois at Urbana-Champaign

August 1991

Equilibrium Price Stickiness

Aldo Rustichini*

and

Anne P. Villamil**

*Department of Economics, Northwestern University, Evanston, IL
**Department of Economics, University of Illinois, Champaign,IL

ACKNOWLEDGEMENT:

Equilibrium Price Stickiness

Aldo Rustichini* Anne P. Villamil†

August 1991

Abstract

This paper develops theory and laboratory experiments to analyze "price stickiness." We study the sequential equilibria of a game that takes place over a sequence of "market periods." There are two risk neutral traders, a buyer and a seller, with differential information. The seller's cost of producing a unit of a fictitious good is known and constant in all market periods. The buyer's value for the good is a random variable governed by a Markov Process that is common knowledge. At the beginning of each period the unit's value is determined by "nature" and privately revealed only to the buyer. The market is organized as a Posted Offer Institution: the seller posts one price offer each period, and the buyer either accepts or rejects it. The market termination rule is a binary random variable. We show that for some parametric specifications, the model generates sticky prices. However, because of differential information price stickiness is an equilibrium phenomenon. We report results from fourteen laboratory experiments designed to compare the decisions of human subjects with those predicted by the theory.

---

*Department of Economics, Northwestern University, Evanston, IL 60208
†Department of Economics, University of Illinois, Champaign, IL 61820
We thank John Dickhaut, Christopher Phelan, John Rust, Neil Wallace and seminar participants at the University of Minnesota for helpful comments. We also gratefully acknowledge financial support from NSF grant SES 89-09242.

1 Introduction

The phenomenon known as "price stickiness" is an important and often discussed issue in macroeconomics. In most economic sub-fields, researchers presume that the price of a good or service adjusts to equilibrate supply and demand. However, some economists argue that this assumption is inappropriate for macroeconomic analyses because some wages and prices appear to adjust slowly. Further, there is an implicit presumption that price stickiness, though the phenomenon itself is often never defined precisely, is bad. Standard examples of price stickiness given in this literature are that some labor contracts are set for up to three years while some firms leave their product prices unchanged for long periods. Although there are numerous theories that purport to explain price stickiness, Blinder (1991, p. 89) reports:

> Most economists would, I think, agree that we know next to nothing about which of several dozen theories of wage-price stickiness are valid and which are not. We might have expected statistical tests to have weeded out the weaker theories by now, especially since many have been around for a long time. ... I think the main reason is that most of the theories are *empty* in the following specific sense: Either they involve unobservable variables in an essential way, or they carry no real implications other than that prices are "sluggish" in some unmeasurable sense, or both. This makes econometric modeling a blunt—perhaps even useless—investigative tool.

The purpose of this paper is to address two issues raised by Blinder's comments: First, virtually any pricing behavior (including prices that are inflexible downward and those that are "flat" for long periods of time) can be justified as "optimal" for some specification of agents' beliefs. However, in naturally occurring markets agents' beliefs about future events are often difficult to ascertain with any reasonable degree of confidence because the underlying probability distributions that the beliefs are based on are not well specified (i.e., are unobservable). One goal of this paper is to specify an environment where both agents' beliefs and the pricing behavior implied by theory can be specified and tested. Second, we wish to define price stickiness precisely in the context of our controlled environment, and evaluate the traditional presumption that stickiness is "bad."

The paper proceeds as follows. We study price stickiness in a market game designed to resemble a retail institution. A single buyer of a fictitious good experiences privately observed shocks to his/her valuation of the good. In contrast, the seller knows only the stochastic process that governs the good's value to the buyer. A sequential trading environment is formally specified in Section 2, and is designed to focus on the following questions regarding price stickiness:[1] How long do price adjustments by sellers lag behind "learnable" shocks to demand? And more fundamentally: How do agents form beliefs and pricing/buying strategies in well specified stochastic environments? In Section 3 we formulate the buyer's and the seller's problems as stochastic discounted dynamic programming problems, and derive stationary Markov strategies and equilibria of the market game. Section 4 contains a discussion of rational Bayesian equilibrium belief formation and a notion of equilibrium price stickiness. Section 5 describes the experimental implementation of the theory and Section 6 reports the results from fourteen laboratory experiments. Finally, Section 7 contains concluding remarks.

2 The Model

Consider a market with two risk neutral traders: a buyer and a seller. Trade between the agents takes place over a sequence of periods, and the market is organized as a Posted Offer Institution.[2] In each period, the seller may produce one unit of an indivisible good. The seller's production cost is fixed

[1]Our model departs from two main concerns of much of the existing price stickiness literature: the frequency of price changes by sellers and selection criteria for games with multiple Pareto rankable equilibria. However, we maintain the imperfectly competitive market structure common in these models. See Gordon (1991) for a survey of this literature.

[2]See Ketcham, Smith and Williams [1984] for a complete description of this institution.

and common knowledge. In contrast, the buyer has a reservation value for the good which follows a Markov Process that is common knowledge. The unit's value to the buyer, v, can take on one of two possible values, h or l, denoting high and low respectively. The Markov Process describes the period-to-period dependence (serial correlation) in the unit's value over the sequence of market periods, and has the following features: 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. That is, we assume that $0 < \alpha < 1/2$. Of course, three other cases are possible:

(i) When $\alpha = 0$, the initial draw determines the unit's value for all periods.

(ii) When $\alpha = 1/2$, the unit's values are completely uncorrelated.

(iii) When $1/2 < \alpha \leq 1$, the unit's values exhibit negative serial correlation.

We ignore (i) and (ii) because they eliminate price cycles (the phenomena that motivate our study), and (iii) because it does not correspond to most naturally occurring Markov Processes of economic interest (e.g., weather patterns, oil shocks, and monetary shocks all exhibit positive serial correlation).

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 and the value of each subsequent unit is determined by the stationary Markov Process described above). We assume that the complete description of the Process is common knowledge of both agents, but that the unit's current value is privately revealed only to the buyer. The seller then posts a single price offer. The buyer can either accept or reject the seller's price offer. If the buyer accepts, the unit is traded at the price posted by the seller. The buyer makes a profit on the unit equal to the difference between the unit's (random) value for the period and the posted price. The seller makes a profit on the unit equal to the difference between the posted price and the (known) cost. If the buyer rejects the offer the unit is not produced and is not traded, and both traders make zero profit. This concludes the current market period.[3]

Whether or not trade occurs in the next period is determined by the following probabilistic termination rule: The probability that the market ends each period is given by δ, a number between 0 and 1. For risk neutral traders, this is equivalent to assuming a discount factor of δ. Strictly speaking, the termination rule for the market is the binary random variable $\{Continue, Stop\}$. In the next Section we will state the seller's and the buyer's problems as discounted dynamic programming problems, and $0 < \delta < 1$ is required for these problems to be well defined. Fortunately, δ is a parameter that can be specified by the experimenter via the publicly announced termination rule. The trading and termination procedures are repeated for each subsequent period until the market ends.

3 Strategies and Equilibria

We focus on equilibria in *stationary Markov* strategies for the model described in Section 2. Given the distribution of the buyer's value for the initial unit, and a history of past price offers by the seller and answers by the buyer, the seller forms a belief (for any equilibrium of the game) about the buyer's value for the current unit of the good. In 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. Agents do not consider the history of game *(i.e.,* the sequence of previous price offers and answers) except for the previous price and answer information reflected

[3]There is no "price discounting" in this market—the seller posts a single price and does not revise it. See Davis and Holt [1990] for an analysis of the Posted Offer Institution with discounting but without value or cost shocks.

in the seller's current belief. To derive equilibrium strategies, we assume that the seller's initial belief is common knowledge.[4]

A stationary strategy for the seller is a map from the seller's belief to the space of price offers. We express the seller's belief by a number $w \in [0, 1]$, where w 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 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. The Appendix shows that equilibria in stationary strategies of this game have the following simple state dependent form: An equilibrium is characterized by a triple (w^*, p_l, p_h), which indicates the *critical belief of the seller*, the *low price offer* and the *high price offer*, respectively. Equilibrium strategies for each agent are a pair $(p(w), A(w, p, v))$, where

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

and

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

Thus, equilibrium strategies are a state dependent price offer by the seller denoted by (P_s), and a state dependent answer by the buyer denoted by (P_b), where Y denotes "yes" and N denotes "no." The seller's and the buyer's dynamic programming problems that these state dependent stationary strategies are solutions to are specified in Sections 3.1 and 3.2, respectively.

The equilibria of this game, given by the triple (w^*, p_l, p_h), are sequential equilibria (see Kreps and Wilson [1982]). We have specified $p(w)$ above. To complete our analysis of the equilibria we must provide a rule that the seller follows in order to form beliefs, both on and off the equilibrium path. We assume that on the equilibrium path, the seller follows Bayes rule. Recall

[4] In the experiment, this assumption is guaranteed by the fact that the first unit's value is drawn randomly from a known distribution.

that w denotes the seller's belief that the buyer's value for the current unit is high. Let $w'(\cdot)$ denote the seller's belief that the buyer's value for *next period's* unit is high. From Bayes rule, it follows that

$$w'(\cdot) = w(1 - \alpha) + (1 - w)\alpha.$$

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 it was low but changed states (the second term).

The stationary strategy of each agent is binary, thus there are four possible belief situations:

(i) $w'(p_h, Y, w) = 1 - \alpha$: If the seller posts the high price and the buyer accepts, then the seller believes that the unit's value is the same as in the current period (i.e., is high).

(ii) $w'(p_h, N, w) = \alpha$: If the seller posts the high price and the buyer rejects, then the seller believes that the unit's value has changed (i.e., is low).

(iii) $w'(p_l, Y, w) = w(1-2\alpha) + \alpha \equiv \hat{w}$: If the seller posts the low price and the buyer accepts, then the seller revises his/her belief regarding the unit's value according to the formula given by $\hat{w}$.

(iv) $w'(p_l, N, w) = \alpha$: We assume that if the seller posts the low price and the buyer rejects, then the seller believes that the unit's value has changed.

Beliefs (i), (ii), and (iii) are equilibrium path strategies that follow directly from $w' = w(1 - \alpha) + (1 - w)\alpha$. In particular, 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) is an "off the equilibrium path" strategy, so we must attribute to the agent some belief in order to complete the belief specification rule. We assume that when the buyer rejects the seller's low price offer the seller believes that $v = l$ so $w = 0$.

This is a plausible assumption because when $v = l$ the buyer loses nothing by rejecting the seller's equilibrium low price offer. However, if $v = h$ the buyer foregoes substantial profit by rejecting the seller's low price offer.

3.1 The Seller's Problem

A stationary strategy for the seller is a map $p: [0,1] \rightarrow I\!R_+$. That is, a stationary strategy is a map from the set of beliefs to the space of price offers. We assume that the seller knows his/her own beliefs, w, and his/her own value function defined over $[0,1]$ and denoted by $V_s(\cdot)$. Denote by $\mathbf{1}_{\{A=Y\}}$ the indicator function of the set $\{A = Y\}$, which is a random event from the seller's perspective. Let c denote the seller's (known) cost of producing each unit that is sold.[5] The seller's functional equation can be written:

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

3.2 The Buyer's Problem

Recall that a stationary strategy for the buyer is a function from beliefs, the price offer, and the unit's current value into the answer set $A \in \{Y, N\}$. We assume that the buyer knows the realization of the unit's current value, $v = h, l$, takes the seller's price and beliefs as given, and knows his/her own value function defined over $[0,1] \times I\!R_+ \times \{h, l\}$ and denoted by $V_b(\cdot)$. The buyer's functional equation can be written:

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

where the expectation is taken with respect to the probability induced by the stochastic process for the unit's value to the buyer, conditional on v being the current value. "Primes" denote next period's magnitude for the variable.

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

4 Discussion of the Solution

In the sequential market game that we consider, the seller moves first by posting a price, but has imperfect information (both ex ante and ex post) about the value of a sequence of units to the buyer. The seller knows his/her own cost, the probability structure that generates the buyer's sequence of values, and forms a belief about the unit's value in any given period. More specifically, the seller knows: the unit's value is either high or low, the Markov Process that governs the evolution of the unit's value, and the previous answers of the buyer. However, the seller does not observe directly the actual realization of the unit in any market period. The buyer moves second and responds to the seller's price offer—after the state of nature has been revealed. The buyer knows the seller's price offer and cost (and consequently the seller's profit), as well as the unit's current value.

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 *critical belief*, denoted by w^*, about the unit's value to the buyer. Recall that w denotes the probability that $v = h$. The seller's state contingent pricing strategy indicates that the seller should post 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 value is high; but should accept only the seller's low price offer if the value is low. Thus the buyer rejects high price offers only when the value is low, and otherwise accepts any price that does not exceed l. In the Appendix we prove that strategies (P_s) and (P_b) are indeed solutions to the seller's problem and the buyer's problem, respectively, in Lemma A.1 and Lemma A.2.

It is interesting to note that the buyer's *equilibrium* strategy sometimes reveals information to the seller, and when it does the information is revealed truthfully. At first glance this may seem odd. One may wonder why it is not optimal for the buyer to reject the seller's high price offer when the value is high in an attempt to mislead the seller and drive down the price. This strategy is not optimal in the market that we consider because lying is costly so the buyer's equilibrium strategy (P_b) involves an essential tradeoff. If the buyer accepts price offer p_h this action reveals information to the seller, *but* the buyer gets an immediate reward for telling the truth (i.e., the profit from the trade). On the other hand if the buyer lies (by rejecting p_h when in fact $v = h$) this action distorts the seller's belief, but the cost of lying is the profit foregone on the rejected current trade. Thus, the buyer faces a tradeoff between current profit and manipulating the seller's beliefs (in the hope of obtaining higher future profit)—in a sequential game with a random termination rule and an oscillating value sequence. The seller's equilibrium strategy (P_s) takes this tradeoff into account as w^* (the seller's critical belief) depends on δ (the termination rule) and on l and h (the oscillating values). In the remainder of this Section, we discuss the nature of equilibrium price paths, provide an estimate of w^* in terms of δ, h and l, and discuss "price stickiness" in our model and in naturally occurring markets.

4.1 Equilibrium Price Paths

An equilibrium in stationary strategies is identified with the triple (w^*, p_l, p_h). If such a triple describes all equilibria, then $p_l = l$ by the following argument: $p_l \leq l$ because if $v = l$ then any higher price offer by the seller would be rejected by the buyer, and the buyer would have no incentive to deviate from this answer strategy. Clearly, $p_l = l$ is a dominant strategy in equilibrium. The high price offer is bounded by $l < p_h \leq h$.

A key determinant of the *equilibrium price path* is the relationship between the critical belief w^* and 1/2, where 1/2 is the invariant measure of the Markov Process (which describes the evolution of each unit's value over the course of the game). Note that 1/2 is also the limit belief of the seller when no new information about the unit's value to the buyer is acquired. In fact, $\lim_{i\to\infty} \hat{w}^i = 1/2$ for every w.[6] Thus there are two possible cases:

(1) $w^* \geq 1/2$: In this case the belief region $[0, w^*]$ prevails and is invariant. The equilibrium price offer by the seller is p_l, which gives the seller no new information about the unit's value. Hence if $w \in [0, w^*]$ at some points, then all subsequent beliefs remain in this interval and converge to the limit belief 1/2. The unit's value to the buyer is equal to l in finite time with probability one, and in the following period the seller's belief is in $[0, w^*]$ for any previous belief. Thus, in this case prices attain the constant level p_l in finite time with probability one.

(2) $w^* < 1/2$: In this case, equilibrium prices follow a cyclical pattern for any realization of the buyer's value process. Consider beliefs $w < w^*$. The seller's equilibrium price offer is p_l, and 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 is finite)—then the seller's price offer becomes p_h. If the unit's value is low, then the buyer refuses the seller's offer, the seller sets his/her new belief to α, and the process described above begins again. If the unit's value is high, then the buyer accepts the seller's offer, and 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. Finally, when the unit's value does in fact become low, the buyer rejects the offer, and a new period of low price offers begins. The average length of the periods in which the seller makes low price offers is constant and given by:[7]

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

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

[7] The function $(\cdot)^i$ is defined in the Appendix.

4.2 The Seller's Critical Belief

We now provide an estimate of the magnitude of the seller's critical belief, w^*. An exact computation is difficult because w^* depends on the computation of the equilibria. However, in the Lemma below we provide an estimate of w^* in terms of the known parameters h, l, and δ. This estimate is essential for experimental analysis of the theory. Recall that w is a probability, thus we wish to obtain non-trival left and right bounds on w^* between 0 and 1. We begin by introducing the following notation:

$$Z_r \equiv \frac{p_l}{p_h};\ z_r \equiv \frac{l}{h(1-\delta)};\ Z_l \equiv \frac{p_l(1-\delta)}{p_h - p_l + p_l(1-\delta)};\ \text{and}\ z_l \equiv \frac{l(1-\delta)}{h-l+l(1-\delta)}.$$

Clearly we have: $z_l \leq Z_l < z_r(1-\delta) \leq Z_r$, since $p_l = l$ and $p_h \leq h$ in equilibrium.[8] Also at $w = 1$ we must have, in equilibrium, $p_h \geq h(1-\delta)$, or the seller would deviate and post any price below h that he/she expects the buyer to accept with probability one. Therefore, $Z_r \leq z_r$.

The following Lemma provides an estimate of the seller's critical belief.

Lemma. $z_l \leq w^* \leq z_r$.

Remark. To simplify the analysis, normalize $p_l = l = 1$. Observe that if $h > \frac{2}{1-\delta}$, then $w^* < 1/2$ so the equilibrium outcomes have persistent cycles of the nature described previously. Alternatively, if $h < 2-\delta$, then it follows from the inequality in the Lemma that $w^* > 1/2$ and the equilibrium outcomes converge to the low price offer with probability one. These results are essential for the experimental analysis of our theory, and are consistent with the following intuition: When the seller believes that the unit's value is high, $v = h$, and consequently believes that the buyer has a large consumer surplus (that the seller wishes to exploit), the seller posts a high price to increase revenue *and* to acquire information about the unit's value. In the experimental analysis of the model in Section 6 we choose h to be sufficiently

[8]Unlike the low price, $p_l = l$, the equilibrium high price, p_h, need not be unique (i.e., the distribution of exchange surplus is indeterminate).

high (i.e., \$2.50 or \$2.20) in some experiments, given $l = \$1$ and $\delta = 0.05$, to ensure that the seller has an incentive to acquire information. In equilibrium the high price is fully revealing, so when h is high enough theory predicts the seller pursues the cyclical pricing strategy even at the cost of some occasional lost trades. In other experiments we choose h to be low enough (i.e., \$1.50) to ensure that it is not optimal for the seller to attempt to extract information from the buyer, so theory predicts prices attain the constant level p_l.

Proof. When $\delta = 0$, the seller is indifferent between posting the low or high price if $w = p_l/p_h$. However, when $\delta > 0$ price offer p_h gives the seller additional information about the unit's value. Thus, for $\delta > 0$, $w^* \leq Z_r$.

If w satisfies[9]

$$p_l + \delta V_s(\hat{w}) - \delta V_s(\alpha) - w[p_h + \delta(V_s(1-\alpha) - V(\alpha))] \equiv M > 0,$$

then $w < w^*$. But $M \geq \frac{1}{1-\delta}[p_l(1-\delta) - w(p_h - \delta p_l)]$; so if $\frac{p_l(1-\delta)}{p_h - p_l + p_l(1-\delta)} > w$, then $w \leq w^*$. We have shown that $w^* \geq \frac{p_l(1-\delta)}{p_h - p_l + p_l(1-\delta)} \equiv Z_l$, and this concludes the proof.

4.3 Price Stickiness

The market game described in this Section has several interesting features that may provide insight into "price stickiness" phenomena in actual markets. First, observe that three different pricing patterns can be optimal, depending on the seller's beliefs:

(i) $p(w) = p_l$: If the seller posts only low price offers, it is always optimal for the buyer to accept such offers (regardless of the value realization). This "always post a low price" strategy is optimal for the seller if the seller believes the probability that the buyer's value is high (w) is weakly below the critical value (w^*).

[9] This equation describes the situation where the value to the seller of posting the low price is strictly greater than the expected value of posting the high price.

(ii) $p(w) = p_h$: If the seller posts only high price offers, it is optimal for the buyer to accept such offers if the value is high. If the value is low, the buyer rejects the offer. This "always post a high price" strategy is optimal for the seller if the seller believes the probability that the buyer's value is high (w) is always above the critical value (w^*).

(iii) $p(w) = (p_l, p_h)$: Price cycles (i.e., oscillations between the high price and the low price) are optimal when the seller's beliefs oscillate about w^* and the length of the cycle is given by L in Section 4.1.

The precise nature of the optimal pricing pattern depends on the seller's beliefs, which in turn depend on the Markov Process and the buyer's decisions. The crucial feature of the solution is that in general, virtually any seller pricing strategy is optimal *given some seller belief specification*. Agents' beliefs about probability structures are often difficult to elicit in naturally occurring markets because actual probability structures are often unobservable. Fortunately, probability structures such as the Markov Process in this paper can be easily controlled in laboratory markets. Our experimental design (specified in the next Section) makes essential use of this feature of the market game.

The second interesting feature of this market game is that "price stickiness" can be defined precisely: We define price stickiness as a situation where the seller's price is not responsive to a change in the buyer's demand (i.e., a value transition). The model predicts periods of "unchanged" prices, and the length of these periods of "flat prices" (on average) is endogenously determined by L. However, in contrast with the traditional presumption that price stickiness is "bad," the non-instantaneous adjustment of prices in our model is an *equilibrium* phenomenon; it results from the seller's attempt to learn the Markov Process. Consequently, the full information prices, $p_l = l$ and $p_h = h$ (where the seller captures all of the exchange surplus), provide a benchmark for measuring the welfare loss associated with price stickiness (or more precisely—differential information).[10]

In conclusion, the differential information price stickiness model that we propose has the following features. First, differential information reduces market efficiency relative to full information when $h > \frac{2}{1-\delta}$ because agents systematically forego trade on units for which the seller posts p_h (to acquire information about the unit's current value) but the value is low. Second, the buyer prefers the differential information solution ($p_l = l$ and $l < p_h \leq h$) to the full information solution because in general it allows the buyer to obtain some exchange surplus. Finally, when information is revealed in the differential information equilibrium it is revealed truthfully—but the equilibrium is *not* fully revealing.[11] The fact that the equilibrium strategy is incentive compatible but not fully revealing provides insight into the following criticism often levied against models with differential information.[12] If differential information causes an economy to achieve Pareto inferior allocations (relative to the full information case), why don't agents willingly reveal their private information to the market and thereby attain the Pareto superior full information allocations? In our environment the answer is clear. Although a social planner is indifferent to the distribution of exchange surplus, individual agents are not. By not announcing the true value, the buyer captures some exchange surplus when the seller posts the low price and his/her value is high (i.e., case (iii) of the Bayesian belief formation rule in Section 3). In contrast, in the full information case the seller fully extracts all available exchange surplus. We discuss the robustness of this result in the final section, and turn now to an experimental analysis of the model.

[10]This is the full information monopoly solution. Multiple bilateral monopoly solutions are inapplicable because of the sequential nature of the game: the seller moves first and posts a "take-it" or "leave-it" price. Further, "two-part pricing" and other intertemporal surplus extraction schemes are inapplicable due to the absence of commitment mechanisms.

[11]Recall case (iii) and the definition of w in Section 3.

[12]For example, this criticism is sometimes made of the Lucas [1972] "Islands Model."

5 Experimental Design

The trading rules of the Posted Offer Institution are reported in Ketcham, Smith, and Williams [1984]. This experimental market is of interest because it resembles naturally occurring retail institutions.[13] We consider a market with a single seller and a single buyer that is conducted over a sequence of "trading periods." There is a single unit of a fictitious good each period. The Posted Offer 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 single unit of 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.[14] The actual value in any period is determined randomly:

(i) The initial value is drawn from a known equal distribution, so the probability that the first unit's value to the buyer is high (i.e., h) is 1/2 and the probability that it is low (i.e., 1) is 1/2.

(ii) All subsequent values are determined by a (first order) Markov Process with the following characteristics: $P(v = 1|v = 1) = P(v = h|v = h) = 1 - \alpha = 0.9$, and $P(v = h|v = 1) = P(v = 1|v = h) = \alpha = 0.1$, where $P(\cdot|\cdot)$ denotes a conditional probability.

The explanation of (i) is obvious. From (ii) it follows that the probability that the current value is the same as last period's value is 90 percent and the probability that it has changed is 10 percent. In general, when α is small there is persistence in the process so there is a big payoff over time to the seller from trying to extract information about the unit's true value. The invariant measure of the distribution is 1/2.

[13] See Plott [1982] and [1989] for surveys of experimental economics.

[14] Based on the computations in the Remark in Section 4.2, price cycles will occur whenever $h > \$2.10$, given that $l = \$1$ and $\delta = 0.05$.

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 value of the *first unit* was high; otherwise it was low.
- In periods $2, \ldots, end$, if the outcome was a 1 through 18, the value of the *current unit* was the same as last period's value; otherwise it changed.

Several examples of stochastic processes were shown to subjects in the Instructions. Further, subjects were told that the invariant measure of the process was 1/2. Specifically, they were told that any *long sequence* of values generated by the value determination rule would have the following characteristics:

(1) On average about half of the values would be high and half would be low, if the experiment lasted for many market periods.

(2) There would be period-to-period dependence in the unit's value over the course of the experiment.

Both the seller and the buyer knew the seller's cost (i.e., $c = 0$), and both knew the value determination rule. However, the current value of the unit 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 ex post was public knowledge.

Finally, the termination rule used in all experiments was stochastic. Subjects were undergraduate and graduate students at the University of Illinois. They were told that the experiment would last between twenty minutes and three hours, and that the final market period was uncertain and would be determined by the following stopping rule: The experimenter would roll a 20-sided die at the end of every market period. If the outcome was a 1 the experiment would end; otherwise it would continue for the next market period. This termination procedure corresponds to a discount factor of $\delta = 0.05$.

6 Experimental Results

We report the results of the following series of experiments that correspond to the theory in Section 4 and the design in Section 5:

(i) six experiments with inexperienced sellers and simulated buyers where $h > \frac{2}{1-\delta}$: equilibrium prices predicted by theory have persistent cycles:

(ii) three experiments with experienced sellers and real buyers where $h > \frac{2}{1-\delta}$: equilibrium prices predicted by theory again have persistent cycles; and

(iii) two experiments with inexperienced sellers and simulated buyers where $h < 2 - \delta$: equilibrium prices predicted by theory converge to p_l.

The first series of experiments was designed to investigate the seller's actual pricing strategy when the buyer always followed the stationary Markov equilibrium strategy (P_b). The second series of experiments was designed to introduce strategic uncertainty into the agents' problems. That is, in the first series of experiments the seller was told that the buyer would make decisions according to decision-rule (P_b): accept the seller's price offer (regardless of whether it is high or low) if the value is high, but accept only the seller's low price offer if the value is low. In the second series of experiments both the buyer and the seller had no information about eachother's decision-rule other than the information contained in the instructions and their observations about eachother's decisions as the game progressed. Thus agents faced strategic uncertainty in addition to differential information about the stochastic process. The third series of experiments was designed as a "check" on the consistency of the model. Finally, in all series (i) through (iii) experiments we elicited the seller's belief about the unit's value in the next period. They were not paid for reporting their beliefs thus the data we report for w' are only suggestive.

We also report the results of a series of pilot experiments:

(iv) three experiments with inexperienced sellers and real buyers where $h > \frac{2}{1-\delta}$: equilibrium prices predicted by theory have persistent cycles.

The pilot experiments correspond to the theory in Section 4 but deviate from the experimental design in Section 5 in the following important way: Subjects were told that there was a 5 percent chance that the experiment would terminate each period to induce $\delta = 0.05$. However, they were also told that when one market experiment ended they would rotate roles and a new market experiment might begin, depending on time constraints. Unfortunately this role rotation procedure undermines the effect of the random termination rule, thus this series of experiments does not constitute an accurate test of the theory. We report the results for completeness.

6.1 Inexperienced Sellers-Simulated Buyers: Cycles

We report the results from the first series of experiments in Figures 1 through 6. Figures 1a through 6a report the results of contract prices and rejections when $h > \frac{2}{1-\delta}$, where $h = \$2.50$ in Figure 1a and $h = \$2.20$ in the remainder. The theory in Section 4 predicts persistent cycling and all of our results are consistent with this prediction. In Figures 1a, 2a, 3a, and 5a the true value of the first unit was low. The seller offered a high price for unit 1, the buyer rejected the offer, and the seller immediately offered the low price for the next trade.[15] In Figures 4a and 6a the true value of the first unit was high. The seller offered a high price for unit 1, the buyer accepted the offer, and the seller maintained the high price offer until it was rejected. There was variance among the sellers in L, the length of periods in which the seller made low price offers. The theoretical prediction is 6 periods when $h = \$2.50$ and 9 periods when $h = \$2.20$.[16] The observed period length of the first cycle of low price offers reported in Figure 1a, where $h = \$2.50$,

[15]All sellers offered a low price that was "behaviorally consistent" with the theoretical equilibrium prediction of $p_l = 1$ because in experimental markets subjects often require a commisssion of \$0.05 to \$0.10 to induce them to trade zero profit units.

[16]Use the definition of L from Section 4.1 to compute these predictions where $\hat{w}^t$ follows from the Bayesian updating rule in Section 3 and w^* follows from z_l in Section 4.1.

is 2. The length of the first cycle of low price offers in Figures 2a through 6a, where $h = \$2.20$, is 3, 5, 5, 4 and 3 respectively. In general, the data reported in Figures 1a through 6a correspond to the theoretical equilibrium price level predictions, the length of the low price cycles are shorter than the theoretical L prediction, and the length of the high cycles is consistent with the theoretical prediction. Sellers' earnings were \$59.23, \$36.40, \$16.89, \$49.40, \$39.40, and \$43.00 respectively.

Figures 1b through 6b report sellers' (unpaid) beliefs that the unit's value in the next period will be high: $w' = w(1-\alpha) + (1-w)\alpha$. The equilibrium predictions in the Figures are derived from the four belief situations delineated in Section 3. The seller's critical belief is $w^* = .388$ in Figure 1b and $w^* = .442$ in Figures 2b through 6b, where w^* is marked by an asterisk on each of the figures.[17] Figures 2b, 3b, 4b, and 6b are somewhat consistent with the Bayesian updating predicted to occur when $w'(p_t, Y, w) \equiv \hat{w}$ (i.e., case (iii)), while Figures 1b and 5b are not consistent with the theoretical predictions for case (iii) beliefs. The reports are generally consistent with the theoretical beliefs in cases (i), (ii), and (iv).

6.2 Experienced Sellers-Real Buyers: Cycles

We report the results of the second series of experiments in Figures 7 through 9 when $h > \frac{2}{1-\delta}$, where $h = \$2.20$ in Experiments 7 and 8 and $h = \$2.50$ in Experiment 9. Theory again predicts persistent cycling, and all subjects in this series of experiments had participated in a series (i) experiment. Their experience profiles are as follows: In Experiment 7 the seller had participated in Experiment 1 and the buyer was in Experiment 5. In Experiment 8 the seller was in Experiment 2 and the buyer was in Experiment 4. In Experiment 9 the seller was in Experiment 3 and the buyer was in Experiment 6.

Experiment 7 lasted for 13 periods and the price data are reported in

[17] This follows from the computation of z_t in Section 4.2.

Figure 7a. There was one value transition: the first 7 units were low and the remainder were high. The seller posted the low price for the first unit and the buyer accepted, an equilibrium answer. The seller posted high prices for two of the next three units, a deviation from equilibrium strategy (P_s), and was rejected, an equilibrium answer. The seller posted the equilibrium low price thereafter (for 8 periods), consistent with the equilibrium L prediction of 9. The seller's profit was $10 and the buyer's profit was $7.20. Interestingly, the buyer always followed equilibrium strategy (P_b), even when he/she received zero profit on a trade. Figure 7b indicates that the seller's (unpaid) belief reports were not consistent with the theoretical w' predictions.

Experiment 8 lasted for 63 periods and the price data are reported in Figure 8a. The buyer frequently deviated from equilibrium strategy (P_b), and this appears to have severely impaired the seller's ability to learn the Markov Process. The buyer's behavior may be explained by two factors. First, subjects in market experiments often require a commission of $0.05 to $0.10 to induce them to trade zero profit units.[18] Thus, the "behaviorally acceptable" equilibrium low price may have been $0.90. Inspection of the data reveals that the buyer accepted only 8 trades and rejected 55. The lowest price the seller offered was $0.90 on two occasions, and the buyer accepted this price both times. The buyer accepted $0.99 on five occasions when the unit's value was high and once (in the last period) when it was low. Second, the buyer's opportunity cost of deviating from equilibrium strategy (P_b) was small. The seller's profit was $7.74 and the buyer's profit was $6.26. Had the buyer followed (P_b) the seller would have earned $65.06 in additional profit but the buyer would have earned only $0.95 in additional profit.[19] Figure 8b indicates that the seller's (unpaid) belief reports were not

[18]No commissions were paid in any of these experiments because one prediction of the theory is that equilibrium strategy (P_b) is incentive compatible. The absence of commissions provide a strong test of this prediction.

[19]The asymmetry in foregone profits stems from the seller's aggressive pricing strategy (i.e., $0.99 and $2.10) and the buyer's unusually large number of low values (49/63).

consistent with the theoretical w' predictions.

Experiment 9 lasted for 41 periods and the price data are reported in Figure 9a. The seller's low price eventually settled at \$0.75, mid-way between the theoretical equilibrium price of \$1 and the "equal profit split" low price of \$0.50. The seller's high price appeared to settle at \$1.25, the "equal profit split" price when the unit's value is high. The length of the seller's low price offer ranged between 3 and 11 periods, but was 6 periods on average (excluding the final sequence of low prices), consistent with the theoretical L prediction when $h = \$2.50$. The buyer deviated from equilibrium strategy (P_b) in periods 3, 8, 9, 10, and 15, but his/her behavior was otherwise consistent with the theory. The seller's profit was \$24.75, the buyer's profit was \$24.25, and the values were equally split between high and low outcomes. The seller's (unpaid) belief reports were generally not consistent with the theoretical predictions for w'.

6.3 Inexperienced Sellers-Simulated Buyers: $p_l = 1$

We report the results of the third series of experiments in Figures 10 and 11, where $h = \$1.50$ in both experiments. Theory predicts that prices will converge to p_l with probability one, and $p_l = 1$. Experiment 10 lasted for 24 periods and Experiment 11 lasted for 39 periods. The data reported in Figures 10a and 11a are generally consistent with the theoretical prediction. Figure 10a indicates that the seller posted a high price ($p_h = \$1.01$) in periods 11 and 19, a deviation from (P_s), but all other decisions correspond to it. The price decisions in Figure 11a correspond exactly to the pricing behavior predicted by theory. The sellers' profits were \$24.51 in Experiment 10 and \$45.50 in Experiment 11. Figures 10b and 11b report the seller's actual and equilibrium beliefs, where $w^* = .655$.[20] Figure 10b indicates that the seller's (unpaid) belief reports are consistent with the theoretical predictions in the

[20] This follows from the computation of z_l in Section 4.2.

first ten periods, but do not display the gradual convergence toward 1/2 in the second half of the experiment predicted by theory. However, belief region $[0, w^*]$ prevails and is invariant after the seller's high price offer is rejected in period 9 as predicted by theory. Figure 11b indicates a similar pattern.

6.4 Inexperiened Sellers-Real Buyers: Pilots

We report price data for three pilot experiments in Figures 12, 13, and 14, where $h = \$2.50$. We did not elicit beliefs from sellers in these experiments, and the termination procedure used in this series of experiments effectively undermined the random termination rule (δ) necessary for the theoretical p, w^*, and L predictions. Subjects were told that when one experiment ended they would rotate roles and a new market experiment might begin, but they were given no explicit information about the number of rotations. Although δ was not well specified, we report benchmark predictions for $\delta = .05$.

The data reported in Figure 12 display the persistent cycles and the state dependent pricing predicted by theory. The experiment lasted for 150 market periods with 4 role rotations. The first "phase" lasted for 92 periods, and the true value of the first unit was low. The seller posted declining high prices for the first six units and was rejected on all trades. The seller then posted a low price for two periods that was accepted, but again posted the high price for the next four periods and was rejected. The seller posted a low price for the next twelve periods and was rejected in five of the twelve periods. The subjects often deviated from equilibrium strategies (P_s) and (P_b), and the data suggest that they may have been engaged in an implicit bargaining process. Eventually they tacitly agreed on the "equal profit split" low price of \$0.50 and the "equal profit split" high price of \$1.25. Thus, subjects conformed to the two-state cyclical pricing rule, but the level of p_l was below the theoretical equilibrium prediction, and the actual length of low prices was somewhat lower than the $L = 6$ predicted by theory. The number of non-equilibrium buyer rejections was dramatically lower in the second half of the experiment.

The data reported in Figure 13 are also generally consistent with the predictions of the theory by the second half of the experiment. The experiment lasted for 75 periods with 3 role rotations. Subjects again tacitly agreed on a low price, in this case of \$0.90, which is behaviorally consistent with the theory, and a high price in the range \$1.25-\$1.75. The data suggest that the subjects again may have engaged in a bargaining process.

Finally, the data reported in Figure 14 are strikingly inconsistent with the predictions of the theory when l = \$1, h = \$2.50, *and* $\delta = 0.05$. The experiment lasted for 94 periods with 4 role rotations. The seller and the buyer tacitly agreed to post and always accept a *single* price offer of \$0.95. However, the theoretical analysis in Section 4 indicates that the cyclical nature of the optimal contract price path depends on both the Markov Process *and* the fact that the experiment may terminate (with probability 0.05) after each market period. Because we indicated to subjects that even if the experiment ended they might switch roles and participate in a new experiment, it appears that they (quite rationally) ignored the random termination rule. If the market is certain to last for many periods, it is optimal for subjects to agree on a single price because the invariant measure of the Markov Process is 1/2 (over many market periods about half of the values will be high and half will be low, so the expected exchange surplus each period is \$1.75). A fixed price of \$0.95 gives the seller a certain profit of \$0.95 each period and the buyer an average profit of \$0.80 each period.

7 Concluding Remarks

This paper develops theory and laboratory experiments designed to study "price stickiness." Our theory indicates that observed prices that appear to be sticky may in fact be generated by rational equilibrium learning about an underlying stochastic process. Although contract prices are "set" for long periods of time, price stickiness is distinct from the persistence in the underlying process. Equilibrium price stickiness occurs when the seller occasionally and intentionally foregoes trade in order to learn the Markov Process that governs the unit's value. The data from laboratory experiments are generally consistent with the predictions of the theory. Unlike previous studies, the theory and experiments provide a framework in which potential sluggishness in price adjustment can be measured precisely. We find little evidence that contract prices are inflexible downward, a result that is consistent with findings by Carlton (1986) for various naturally occurring markets.[21]

Our model can be easily extended to the case of a single seller and multiple buyers, where buyers are randomly chosen to "shop" each period. We have also developed theory and experiments to study the stationary equilibria of the Posted Bid Institution where a single seller with a two-state cost that follows a Markov Process and multiple (five) buyers with known values for a single unit of a fictitious good trade. The equilibrium strategies for this game predict similar episodes of equilibrium price stickiness, but the Posted Bid Institution does not have a clear analog in naturally occurring markets. The introduction of multiple sellers in the Posted Offer Institution is an interesting but difficult problem in this setting that remains for future research. Finally, the unpaid belief data elicited in these experiments suggest that this may be a fruitful environment for studying expectation formation. See Keane and Runkle (1990) for an excellent discussion of the difficulty of testing rationality in naturally occurring markets.

[21]In series (i) and (iii) experiments where the seller knew the simulated buyer was following strategy (P_b), prices were never inflexible downward or sluggish: The seller took rejection of p_h as a perfect signal of a value transition. In series (ii) experiments, "sluggish adjustment" was observed once in Experiment 7 (in period 3); seven times in Experiment 8 (in periods 54 through 60) after persistent deviation by the buyer from (P_b); and eight times in Experiment 9 in periods 2, 4, 9, 10, 20, 33, 36, and 37.

8 Appendix

The function $\hat{(\cdot)}: [0,1] \to [0,1]$ is defined by $\hat{w} \equiv (1-2\alpha)w + \alpha$. Let $\hat{(\cdot)}^i$, for $i = 1, 2, \ldots$, denote its ith iterate. Let the function $\check{(\cdot)}: [\alpha, 1-\alpha] \to [0,1]$ denote the inverse of $\hat{(\cdot)}$, i.e., $\check{w} \equiv (w-\alpha)(1-2\alpha)^{-1}$, and let $\check{(\cdot)}^i$, for $i = 1, 2, \ldots$, denote its ith iterate.

We consider first the problem of the seller. The seller faces a buyer with a policy (P_b) given by

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

The optimal policy of the seller (P_s) is determined by the solution to the Seller's Dynamic Programming Problem, specified in Section 3.1. Recall that P_s is given by the function:

$$p(w) = \begin{cases} p_h & \text{if } w \ge w^*; \\ p_l & \text{if } w < w^*. \end{cases}$$

Lemma A.1. *For a given answer policy of the buyer, the value function of the seller exists, is continuous, and gives a policy of the form P_s.*

Proof. We can write the functional equation which defines the value of the seller as:

$$V_s(w) = \max\{p_l + \delta V_s(\hat{w}); (1-w)\delta V_s(\alpha) + w[p_h + \delta V_s(1-\alpha)]\}$$

The right hand side of this equation defines a map from continuous functions, uniformly bounded above by $p_h(1-\delta)^{-1}$ and below by zero, into the same space. Therefore a unique fixed point exists. We now prove that the policy has the stated form.

For simplicity, denote two functions f and g by:

$$f(w) \equiv p_l + \delta V_s(\hat{w}); \text{ and } g(w) \equiv (1-w)\delta V_s(\alpha) + w p_h + w\delta V_s(1-\alpha).$$

Note that $V_s(w) = \max\{f(w), g(w)\}$, and that f, g are continuous. Since $f(0) > g(0)$, $f(1) < g(1)$, there exists a non-empty set of w's such that $f(w) = g(w)$. We now prove that this set is a singleton. By a standard argument, V_s is concave and so is f, while g is linear. Hence, the set of intersection points is a closed interval, in the interior of (0, 1) because of the inequalities at the boundary. Let w^* denote the right extreme of this interval, i.e., $w^* = \max\{w: f(w) \geq g(w)\}$. Let $A \equiv \delta V_s(\alpha)$, $B \equiv p_h + \delta[V_s(1-\alpha) - V_s(\alpha)]$; now for any $w \in (\max(0, \check{w}^*), w^*)$, one can easily find $f'(w) = \delta(1-2\alpha)B < B = g'(z)$. Hence the interval is the singleton $\{w^*\}$.

Lemma A.2. *For a given price policy of the seller, the value function of the buyer exists, is continuous from the right, and gives a policy of the form P_b.*

Proof. For the (fixed) value w^* in the buyer's policy, let $w_i \equiv \check{w}^{i*}$, for $i = 0, 1, 2, \ldots$. Note that there is only a finite number of non-negative w_i values. Then define the function space S_w^*:

$S_w^* \equiv \{V_b: [0,1] \times \{p_h, p_l\} \times \{h, l\}: V_b(w, \cdot, \cdot)$ is continuous for every $w \in [0,1]$; $V(\cdot; p; v)$ is right continuous and continuous on $[w_{i-1}, w_i)$ for every i, and every p, $v\}$.

With the topology induced by the supremum norm, this is a complete metric space. Consider now the functional equation defining the buyer's value function (see Section 3.2). The right hand side defines all operations from S_w^* into S_w^*. For instance, when $p = p_l$, $v = h$, the corresponding equation is:

$$V_b(w, p_h, h) = \max\begin{cases} h - p_l + \delta[(1-\alpha)V_b(\hat{w}, p(\hat{w}), h) + \alpha V_b(\hat{w}, p(\hat{w}), l)]; \\ \delta[(1-\alpha)V_b(1-\alpha, p(1-\alpha), h) + \alpha V_b(1-\alpha, p(1-\alpha), l)]\} \end{cases}.$$

The other three cases define three similar equations. A standard application of the Contraction Mapping Theorem to space S_w^* gives a unique fixed point.

References:

A. Blinder (1991), "Why are Prices Sticky?," *American Economic Review* **81**, 89-100.

D. Carlton (1986), "The Rigidity of Prices," *American Economic Review* **76**, 637-658.

D. Davis and C. A. Holt (1990), "List Prices and Discounts," Discussion Paper.

D. Gordon (1990), "What is New Keynesian Economics," *Journal of Economic Literature* **28**, 1115-1171.

M. Keane and D. Runkle (1990), "Testing the Rationality of Price Forecasts: New Evidence from Panel Data," *American Economic Review* **80**, 714-735.

J. Ketcham, V. L. Smith, and A. Williams (1984), "A Comparison of Posted-Offer and Double Auction Pricing Institutions," *Review of Economic Studies* **51**, 595-614.

D. Kreps and R. Wilson (1982), "Sequential Equilibria," *Econometrica* **50**, 863-894.

R.E. Lucas (1972), "Expectations and the Neutrality of Money," *Journal of Economic Theory* **4**, 103-124.

C. Plott (1982), "Industrial Organization Theory and Experimental Economics," *Journal of Economic Literature* **20**, 1484-1527.

C. Plott (1989), "An Updated Review of Industrial Organization: Applications of Experimental Methods," R. Schmalensee and R.D. Willig (eds.), Handbook of Industrial Organization, Volume II, Elsevier Publishers, North Holland.

V. L. Smith (1976), "Experimental Economics: Induced Value Theory," *American Economic Review* **66**, 274-279.

Figure 1a

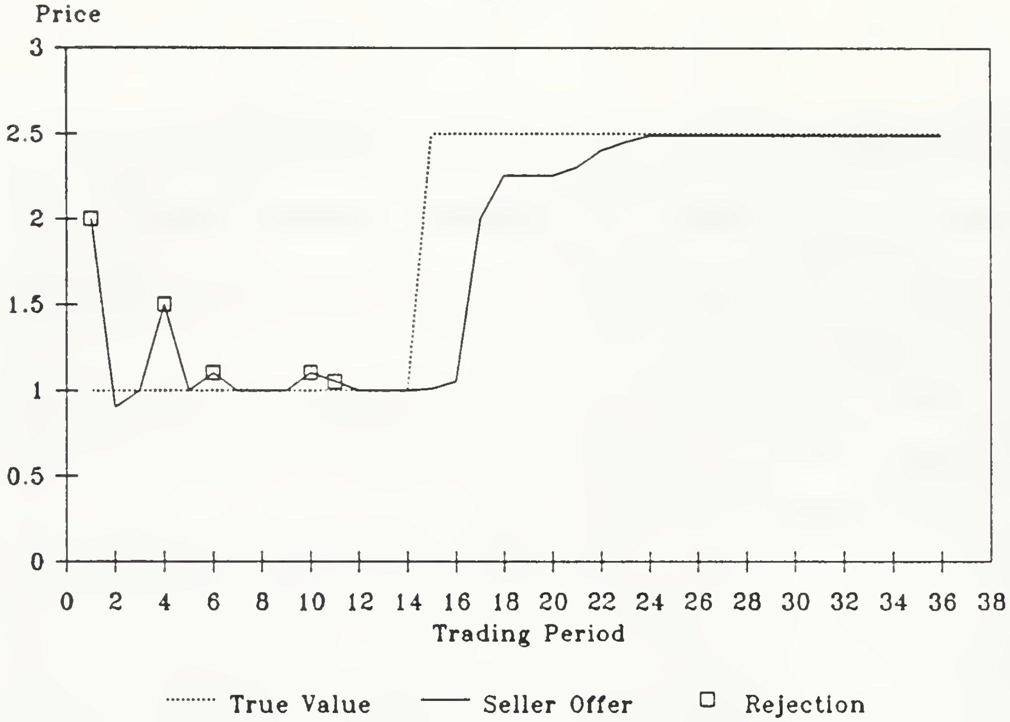

Figure 1b

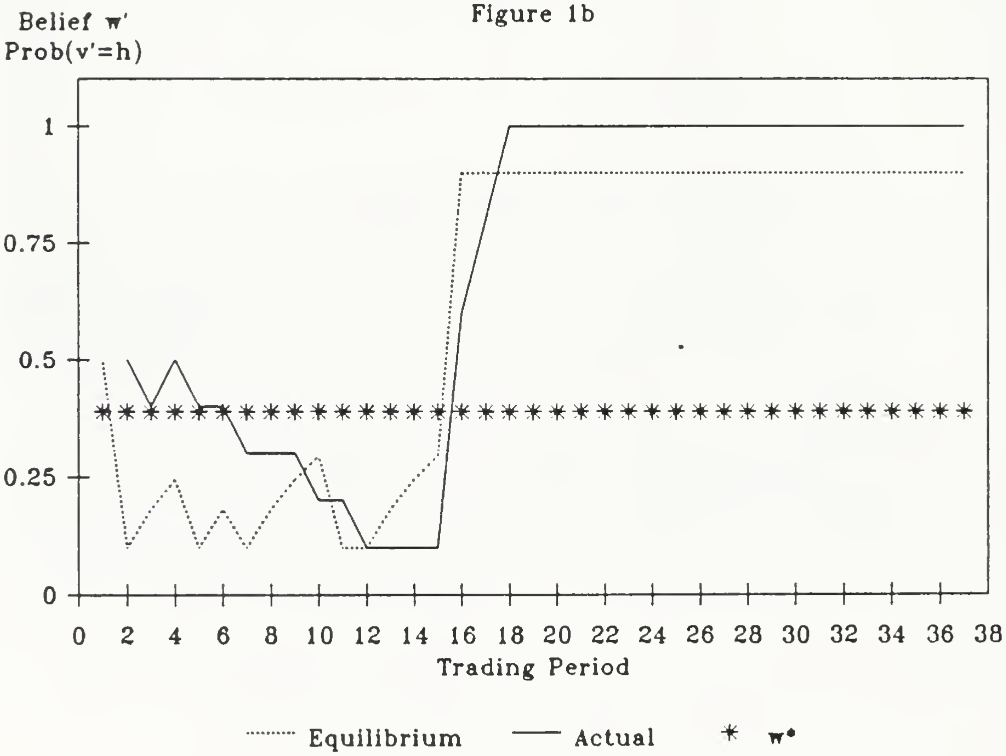

Figure 2a

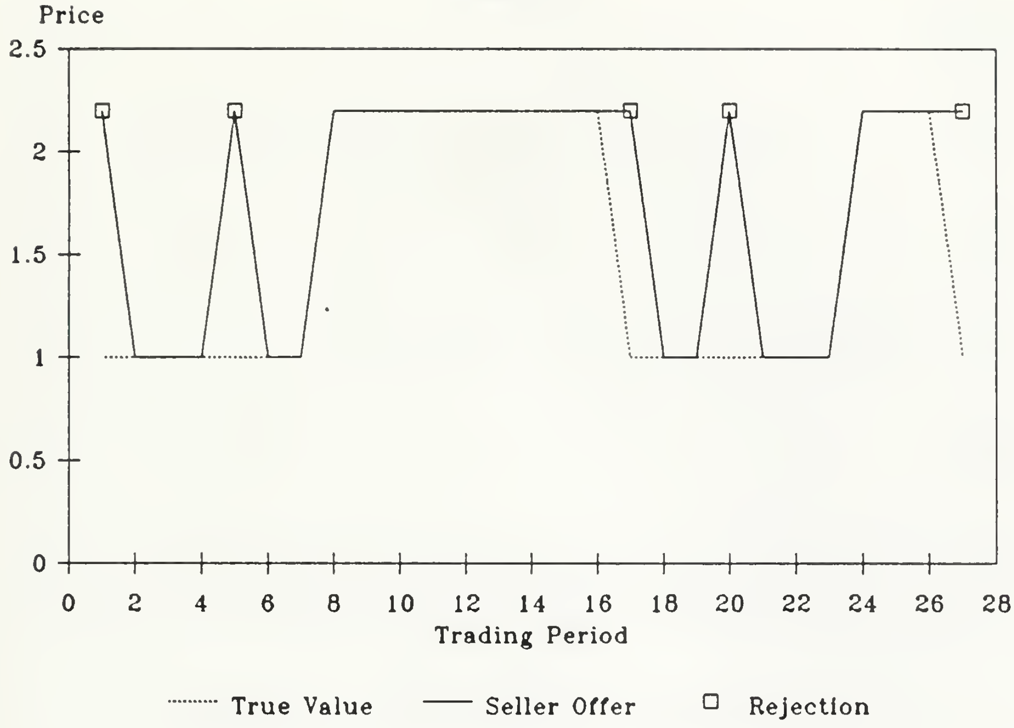

Figure 2b

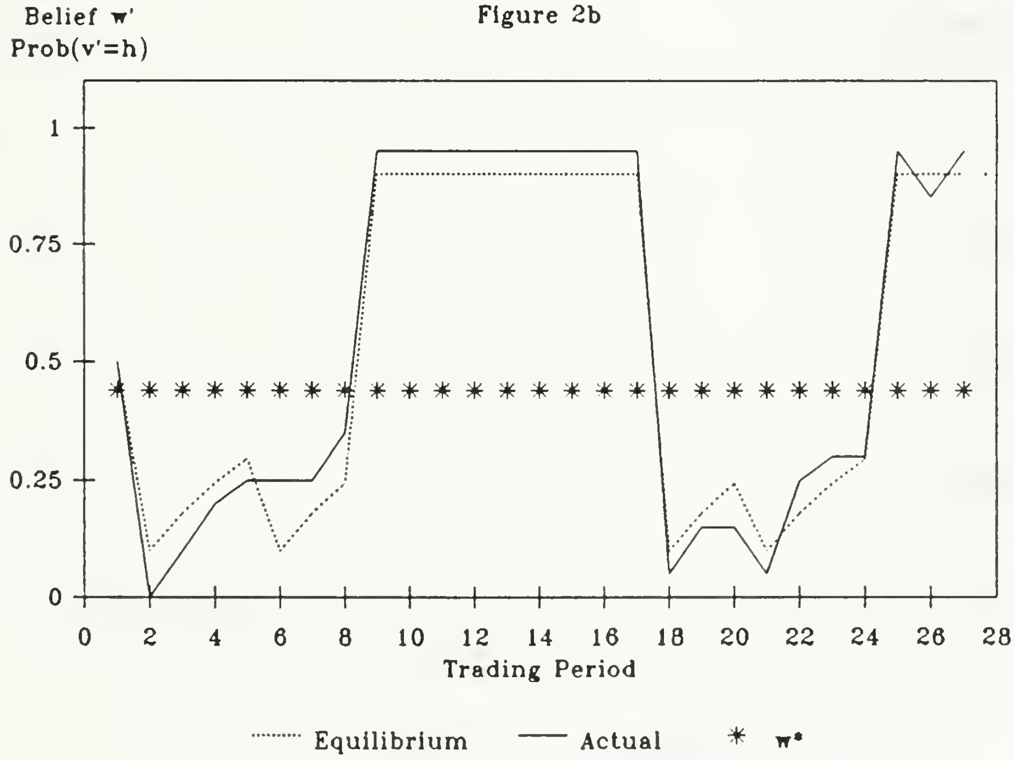

Figure 3a

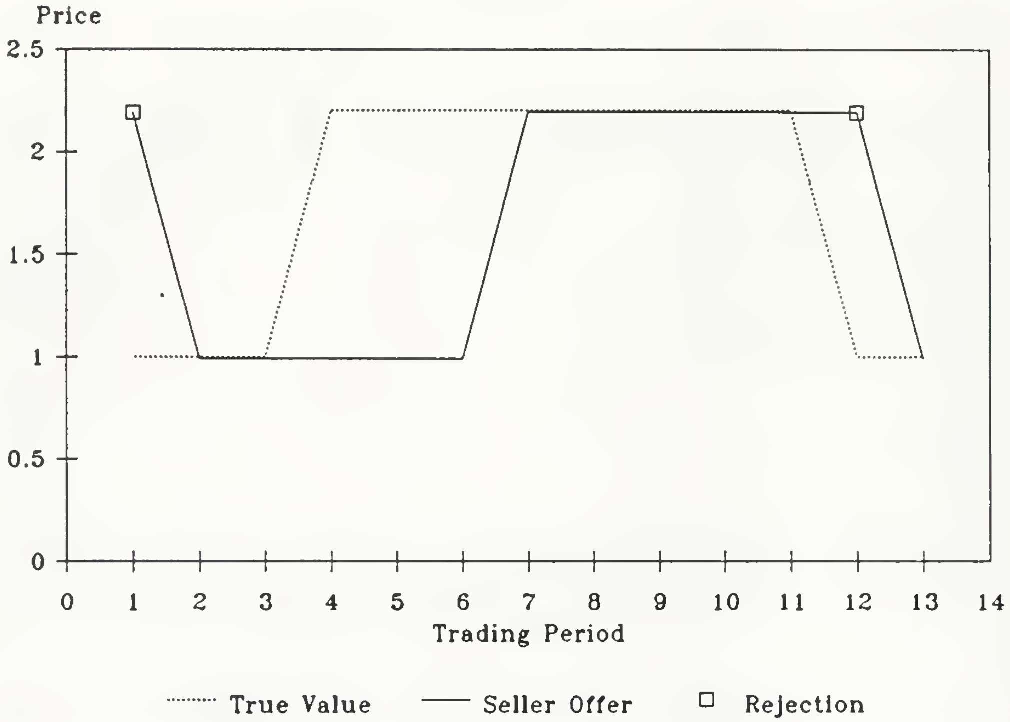

Figure 3b

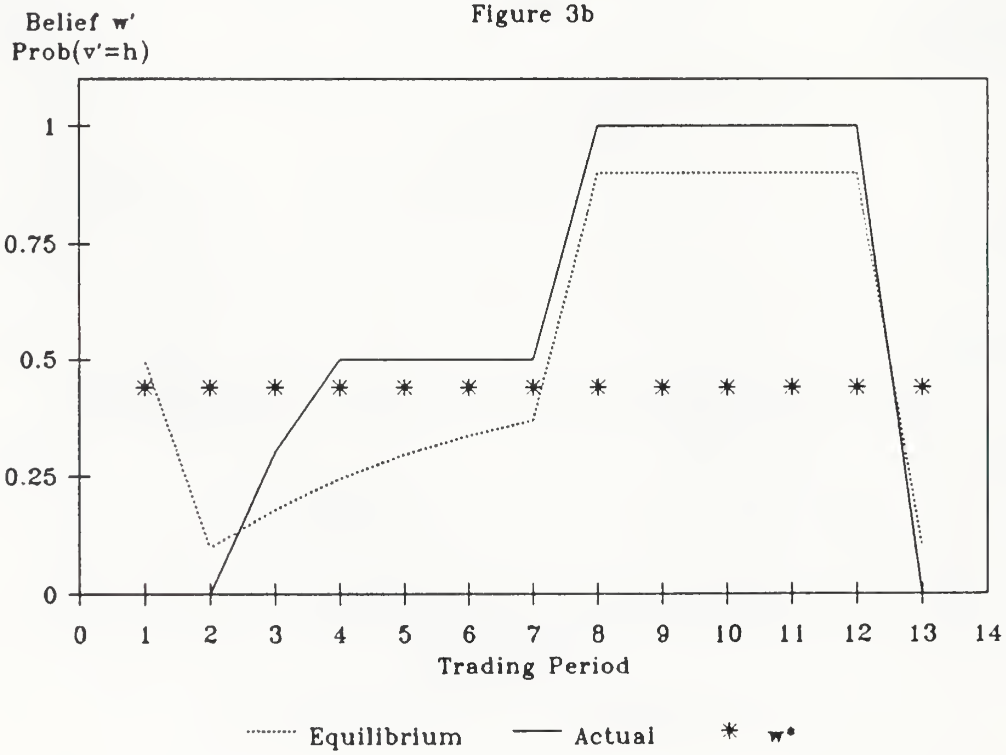

Figure 4a

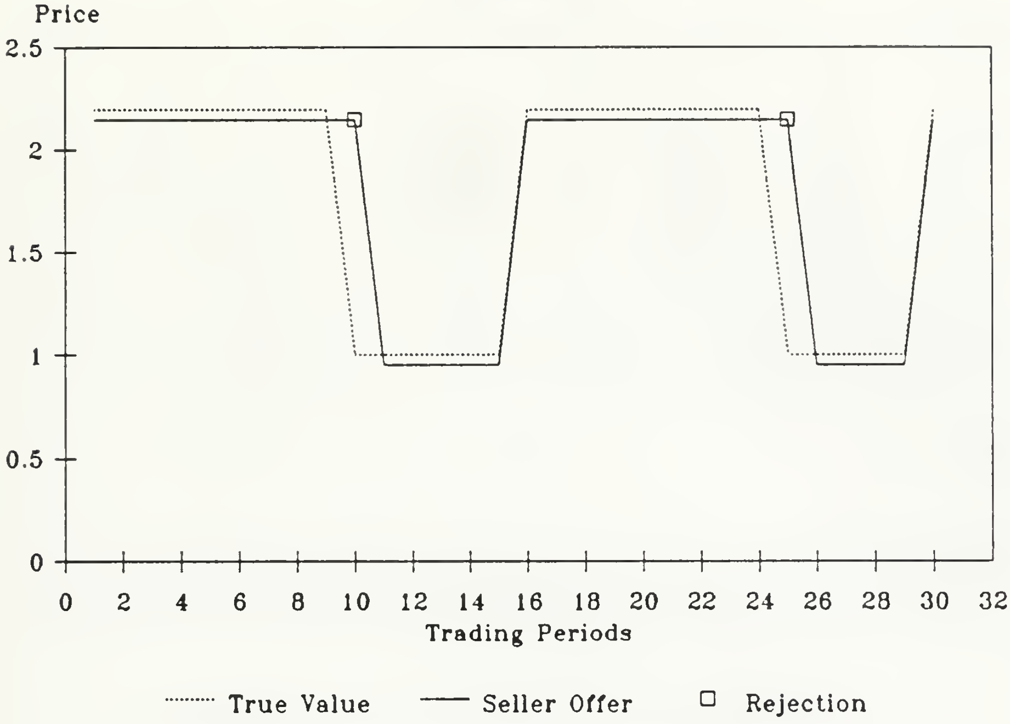

Figure 4b

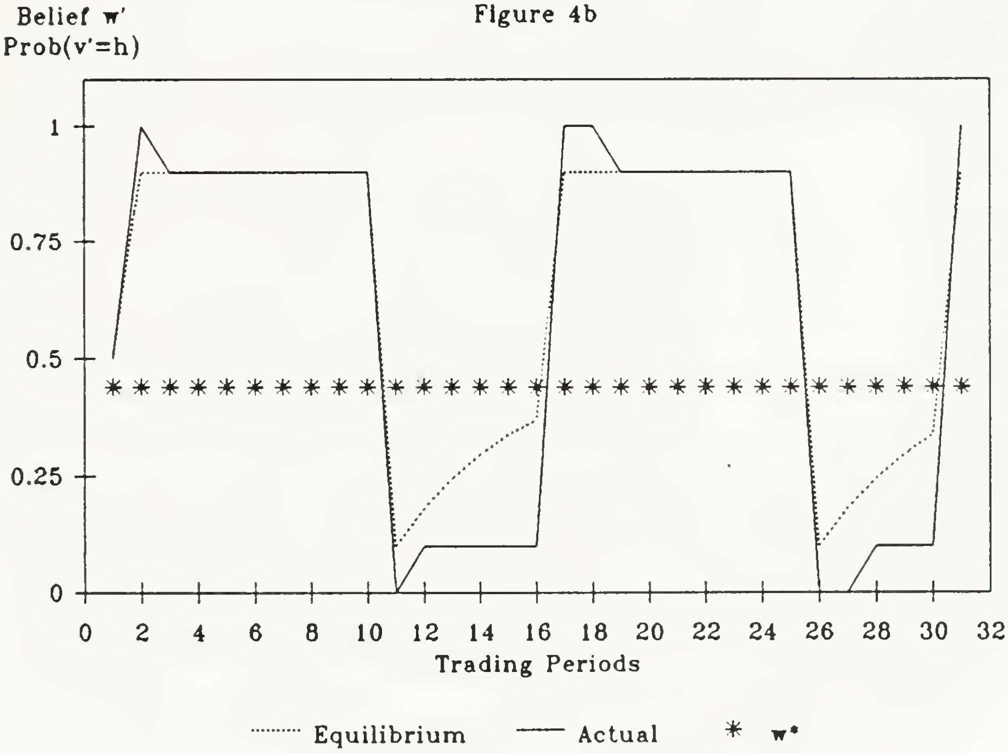

Figure 5a

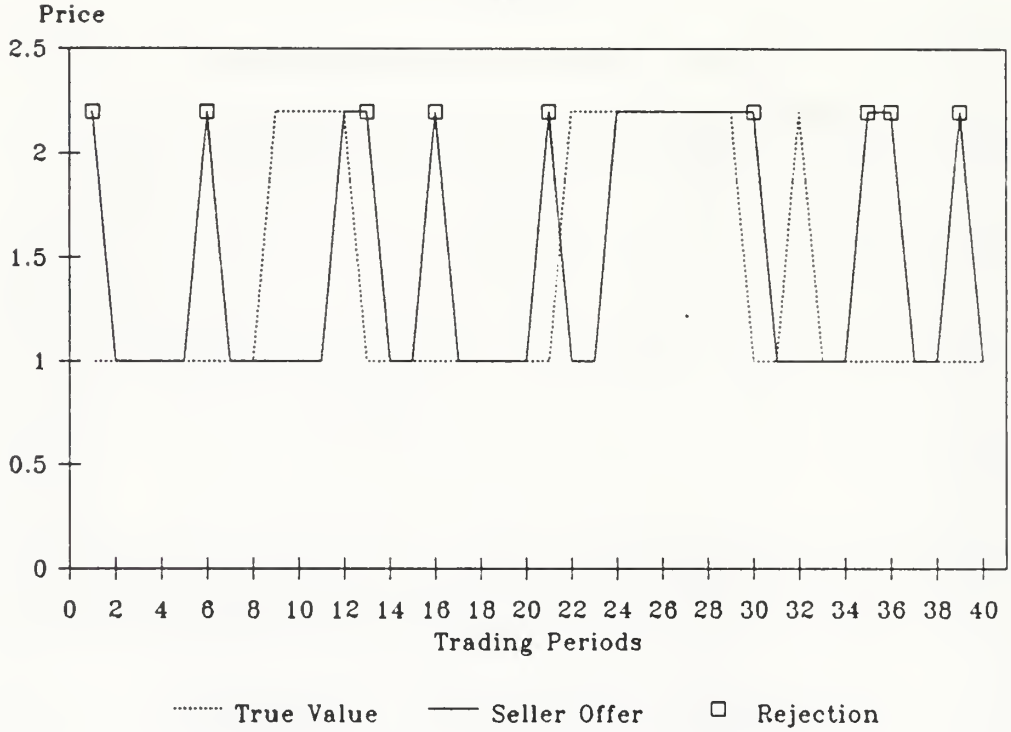

Figure 5b

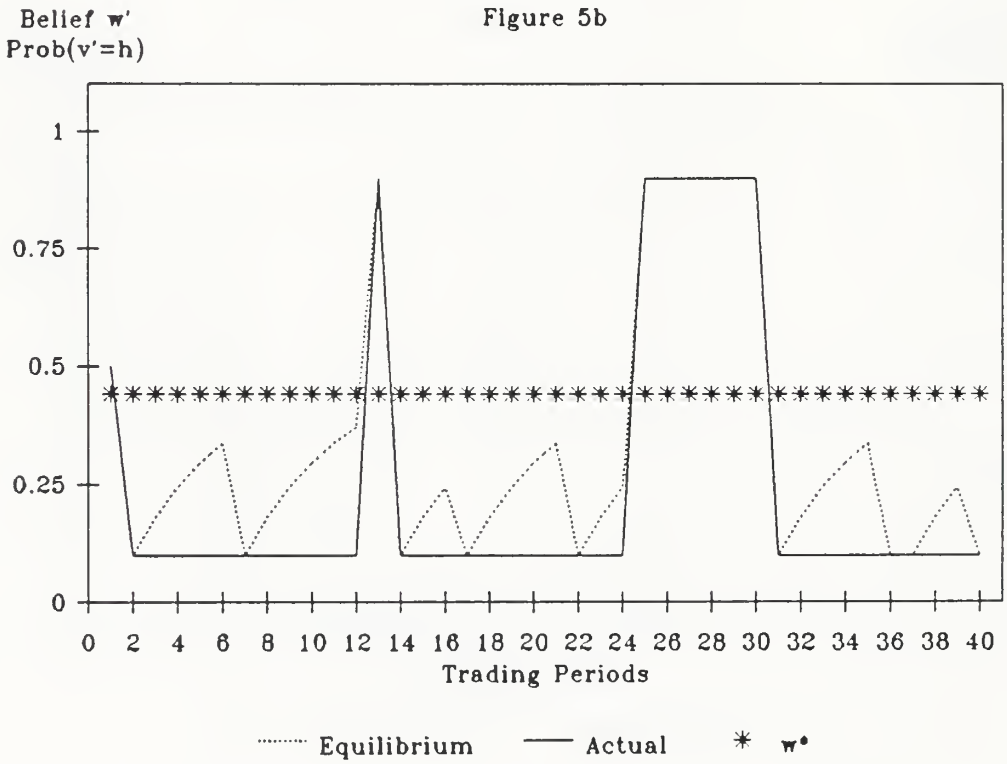

Figure 6a

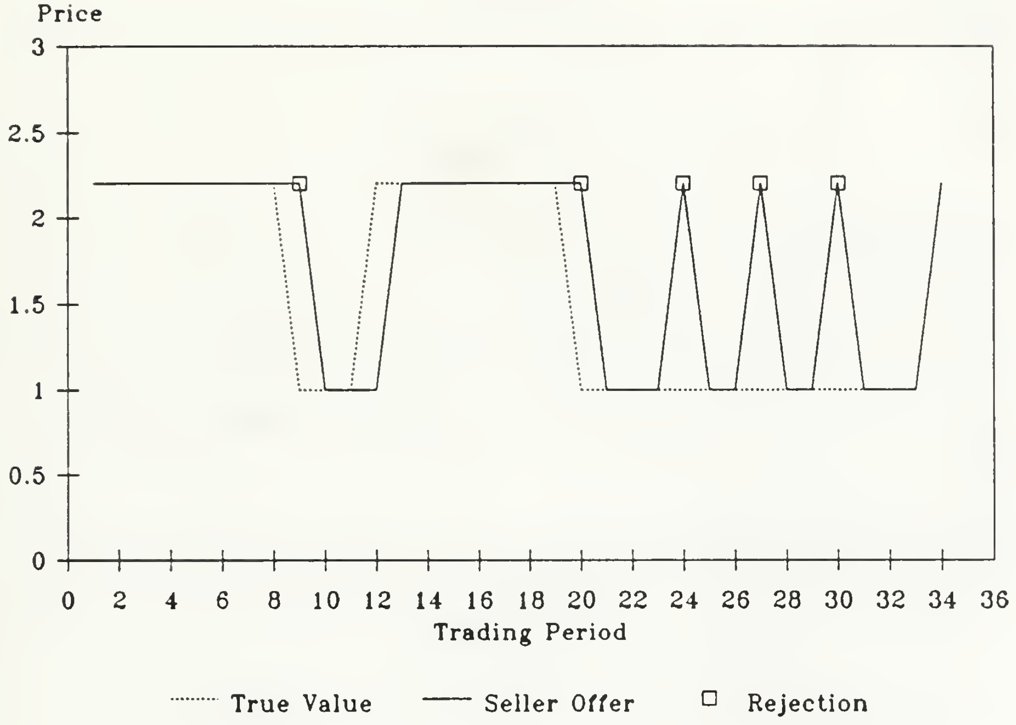

Figure 6b

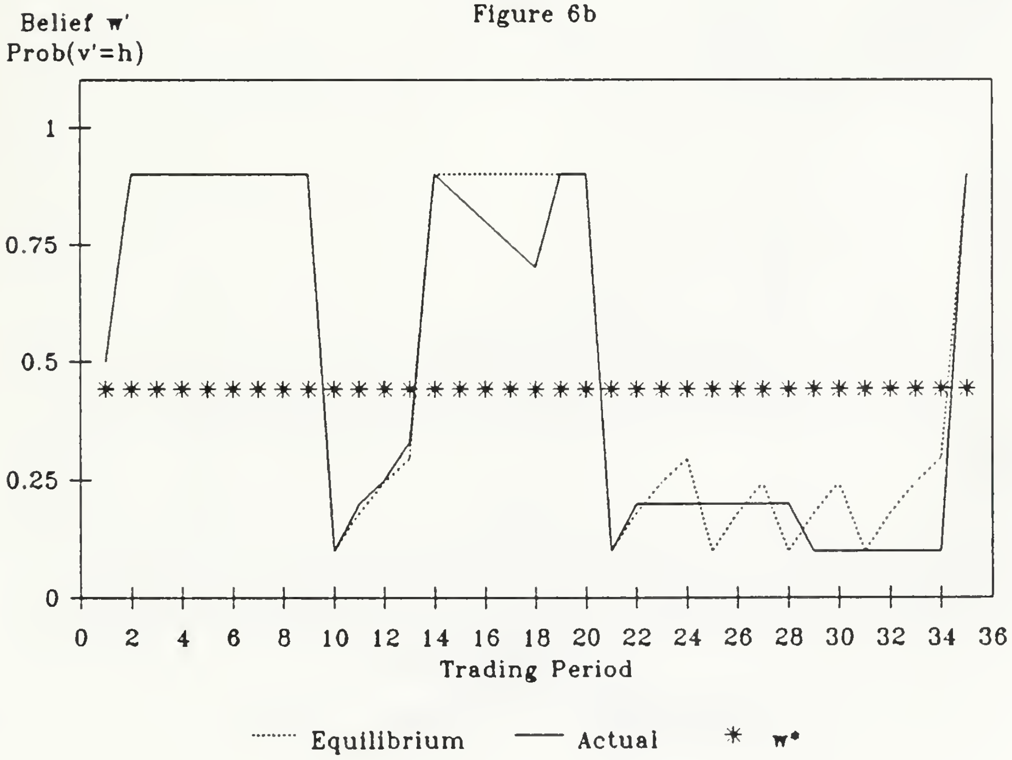

Figure 7a

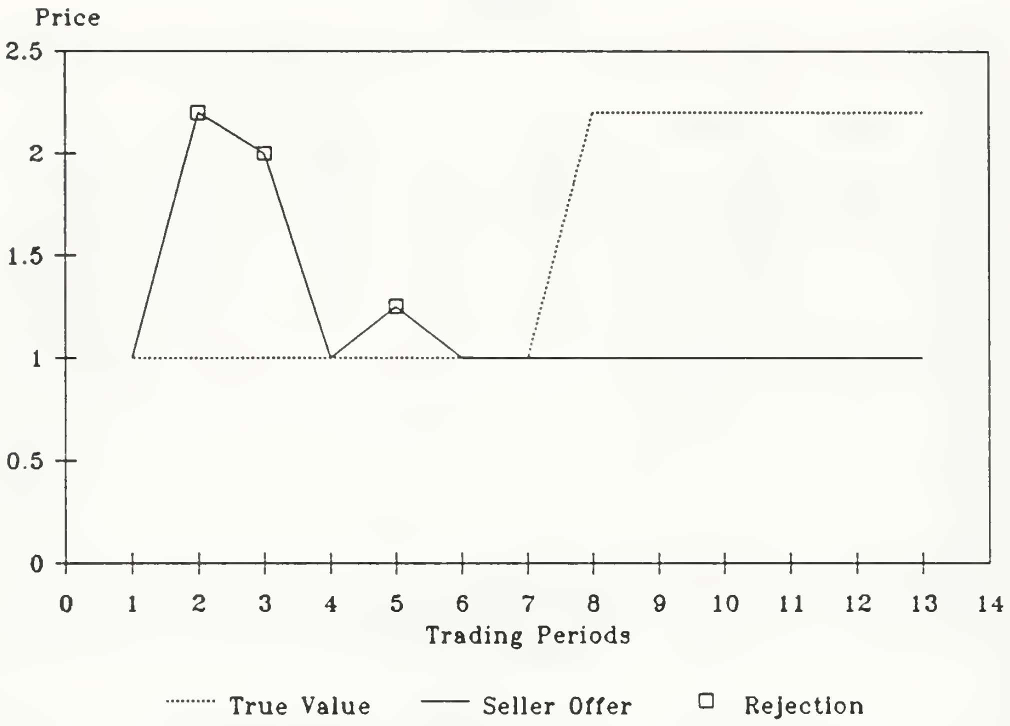

Figure 7b

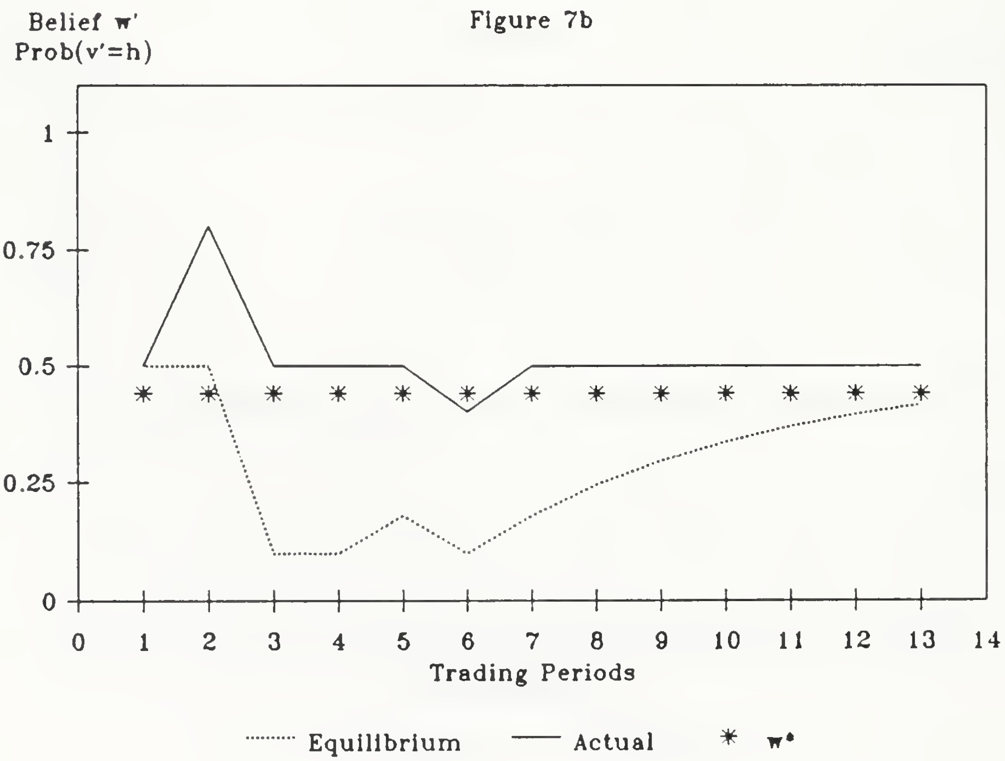

Figure 8a

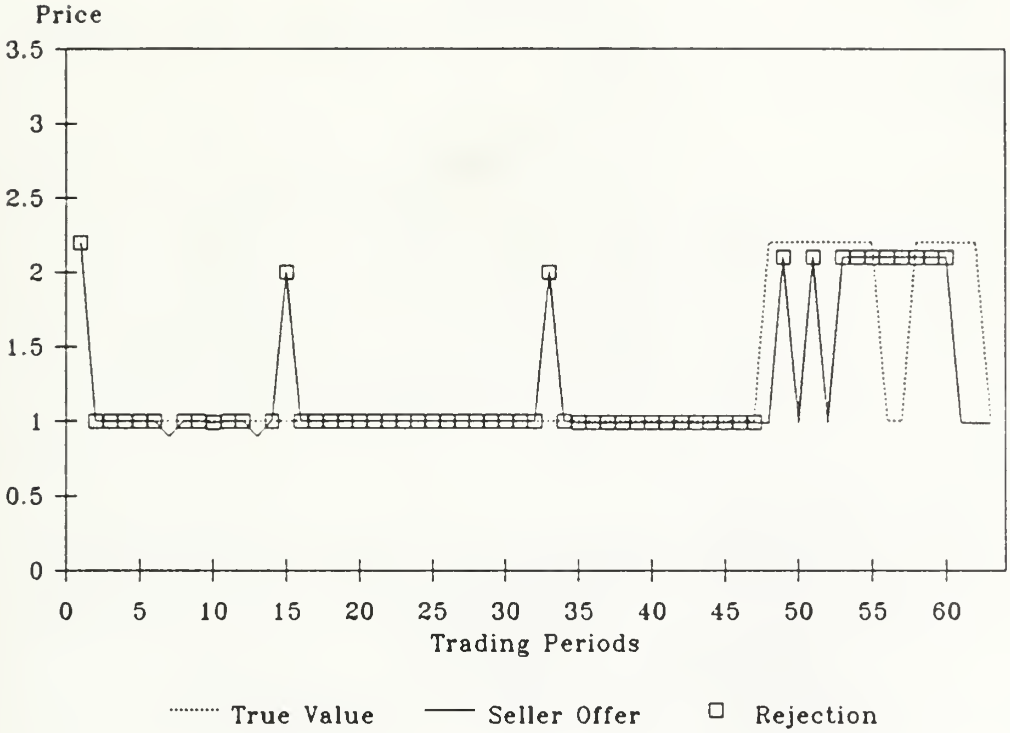

Figure 8b

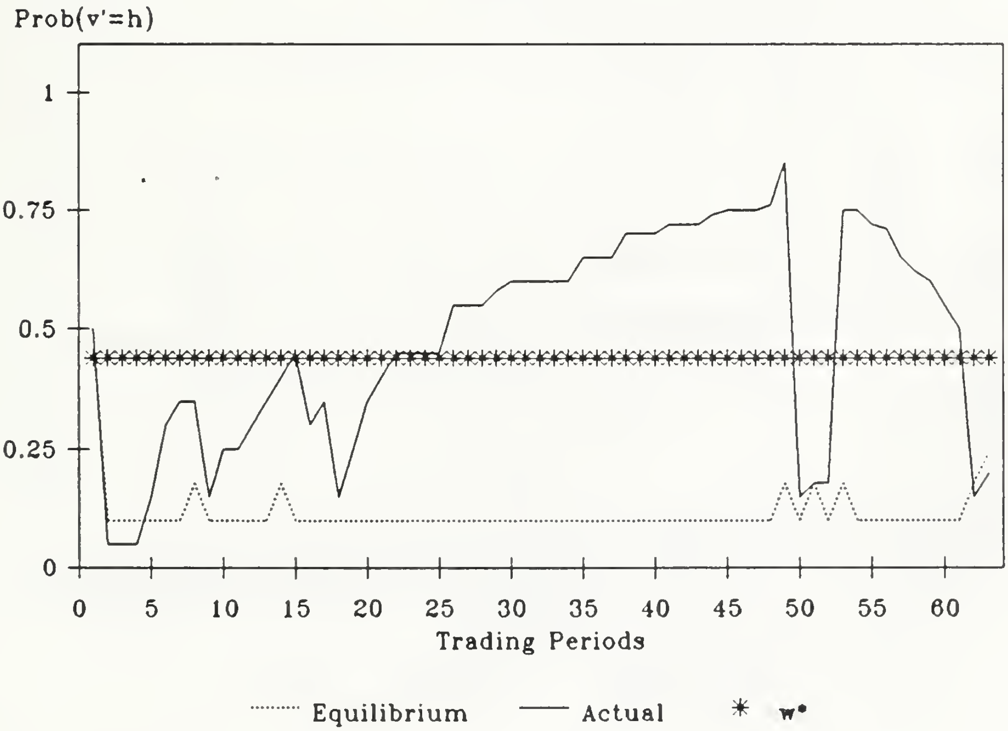

Figure 9b

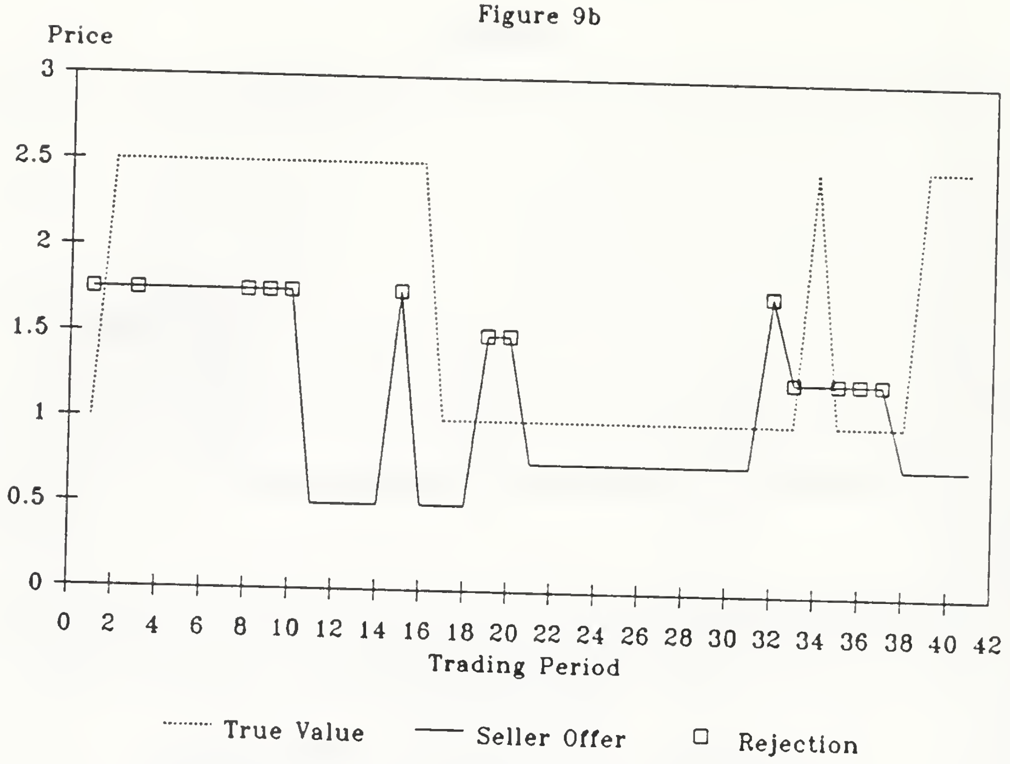

Figure 9b

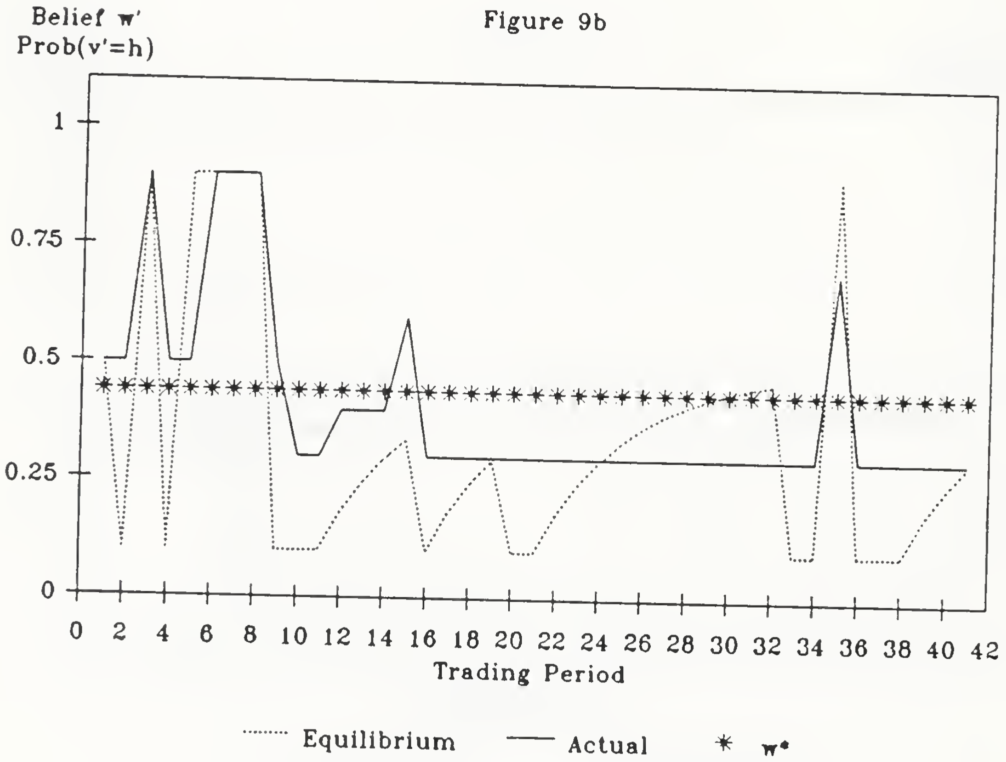

Figure 10a

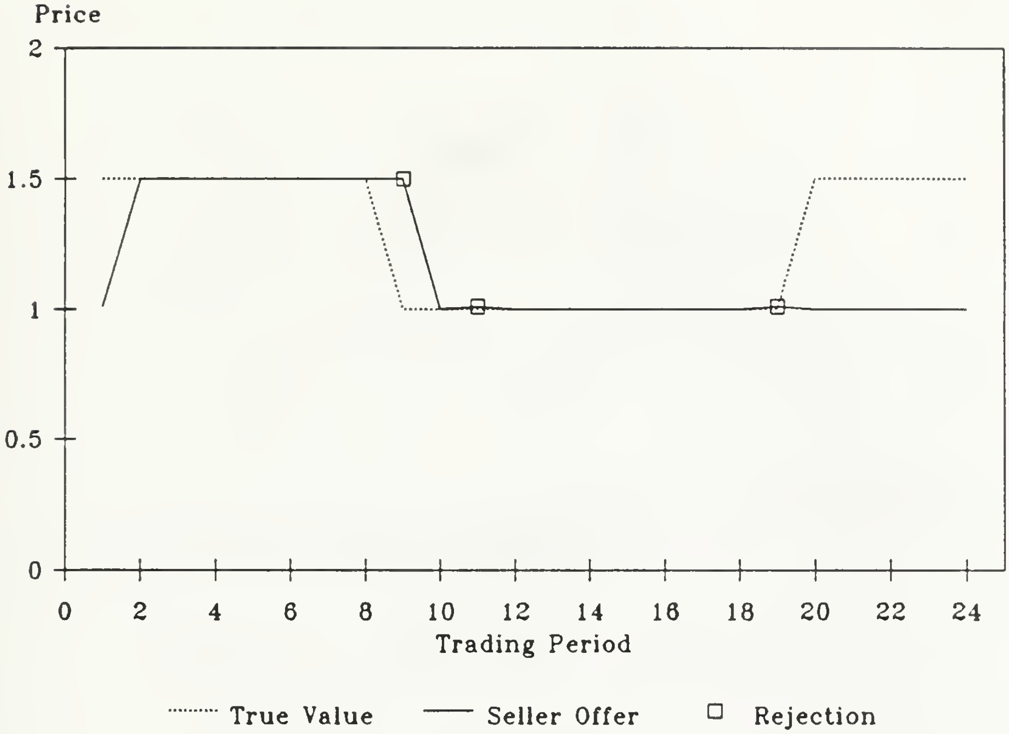

Figure 10b

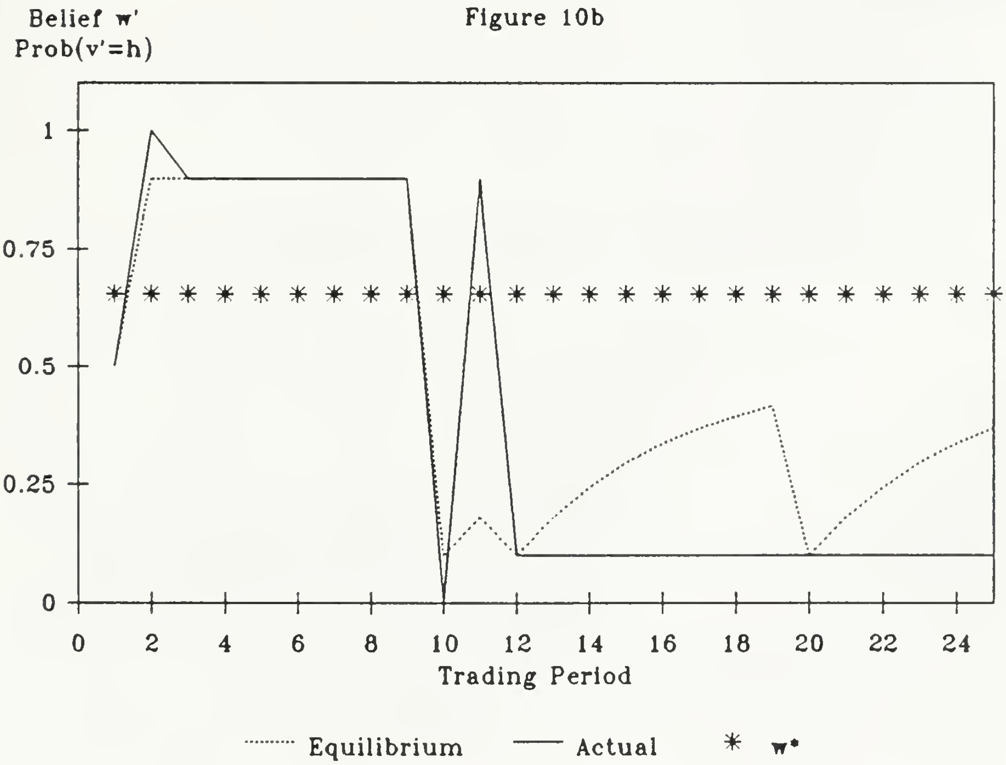

Figure 11a

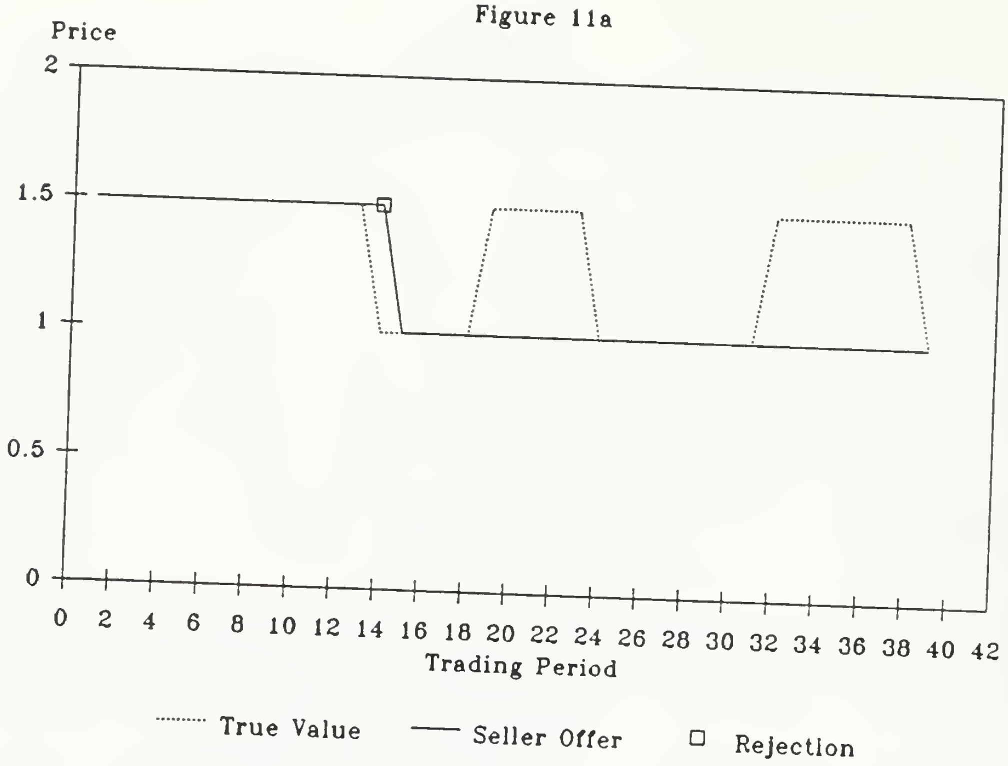

Figure 11b

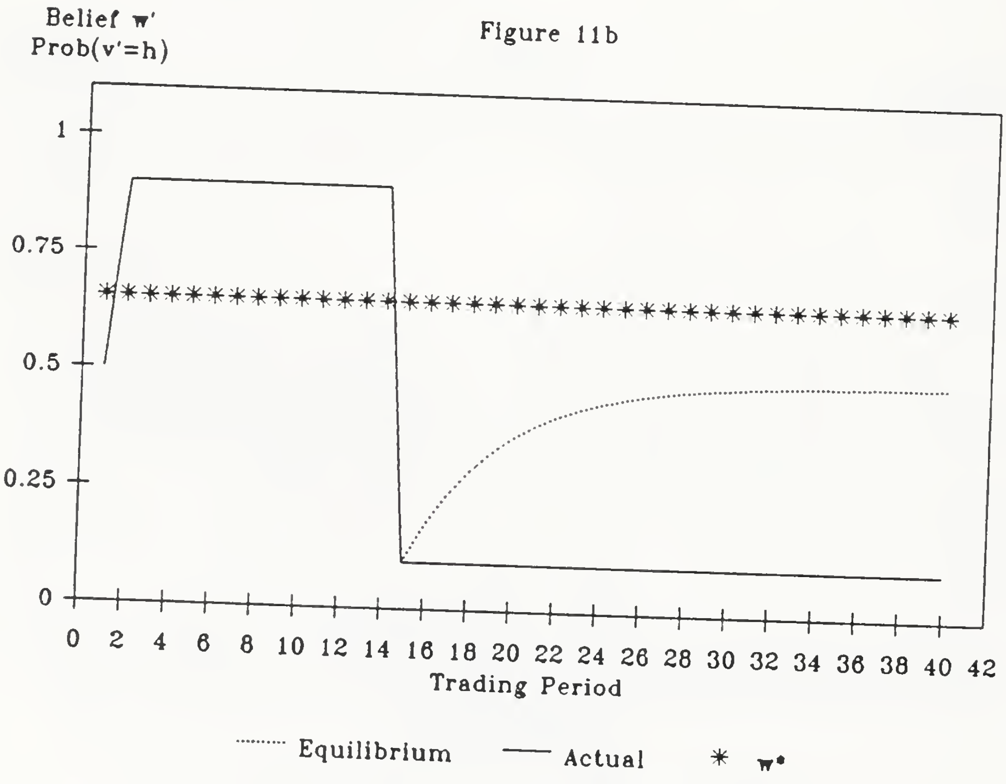

Figure 12

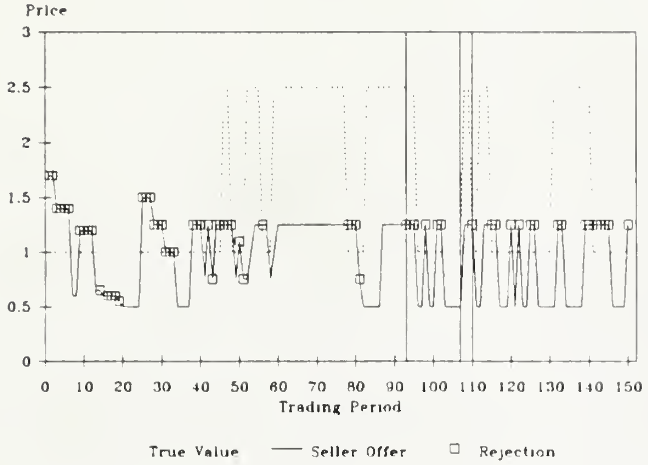

Figure 13

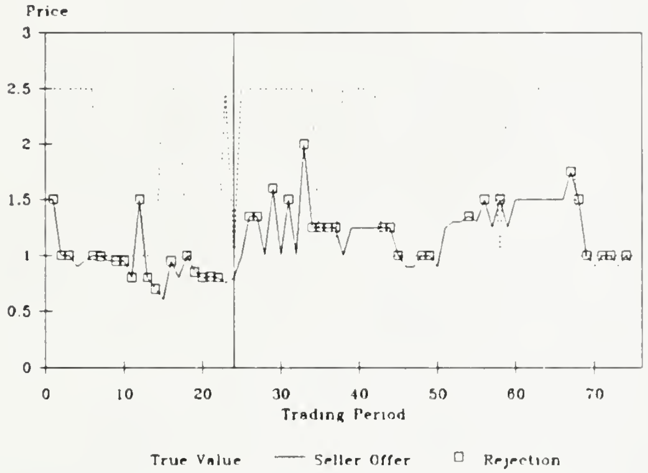

Figure 14