Buyer Search and Price Dispersion:
A Laboratory Study

by Timothy Cason
Department of Economics, Purdue University
and
Daniel Friedman
Department of Economics, UCSC

January 31, 2000

Abstract:
Posted offer markets with costly buyer search are investigated in 18 laboratory sessions. Each period sellers simultaneously post prices. Then each buyer costlessly observes one or (with probability 1-q) two of the posted prices, and either accepts an observed price, drops out, or pays a cost to search again that period. The sessions vary q, the search cost, and the number and kind of buyers. Equilibrium theory predicts a unified very low (very high) price for q=0 (q=1) and predicts specific distributions of dispersed prices for q=1/3 and 2/3.

Actual transaction prices conform rather closely to the predictions, especially in treatments with many robot buyers. Individual buyer and seller behavior, however, differs systematically from the equilibrium predictions: buyers' reservation prices are biased away from the extremes and sellers' posted prices have positive autocorrelation and cross sectional correlation. Learning models can account for a portion of these deviations from equilibrium behavior.

Acknowledgements: We thank the NSF for funding the work under grants SBR-9617917 and SBR-9709874; Brian Eaton for programming assistance, and Garrett Milam, Alessandra Cassar, Sujoy Chakravarty and Sharad Barkataki for research assistance; and Jim Cox, Peter Diamond, Daniela DiCagno, David Easley, Andy Muller, Charles Plott, and audiences at LUISS, DePaul, UCSB, HKUST, the Economic Science Association and the Latin American Econometric Society Conferences for helpful suggestions. We retain all downside and residual upside responsibility.

Timothy N. Cason
Associate Professor
Department of Economics
Krannert School of Management
Purdue University
West Lafayette, IN 47907-1310

e-mail: cason@mgmt.purdue.edu
http://www.mgmt.purdue.edu/faculty/cason/
voice: (765) 494-1737
fax: (765) 494-9658

Daniel Friedman
Professor of Economics
212 Social Sciences I
University of California
Santa Cruz, CA 95064

(408) 459-4981 (fax 459-5900)
http://cash.ucsc.edu/Friedman/
dan@cash.ucsc.edu

Buyer Search and Price Dispersion: A Laboratory Study

1. Introduction

Transactions in modern economies occur predominantly in posted offer markets. Typically each seller posts a price and each buyer chooses a seller.[1] Buyer search is the main competitive force in such markets: as long as the posted prices differ, buyers tend to choose sellers with lower prices. In many circumstances buyer search enforces a unified, competitive price. But sometimes different sellers post substantially different prices for essentially the same good, and we have price dispersion. Our goal in this paper is to better understand when buyer search enforces the law of one price and when it allows price dispersion.

The goal is important for three complementary reasons. First, dispersed prices are a persistent fact of life, even when buyer search is relatively cheap (Brynjolfsson and Smith, 1999; Baye and Morgan, 1999). Second, despite a large theoretical literature, price dispersion is not yet well understood. The next several paragraphs will highlight some of the outstanding theoretical issues. Third, a deeper understanding of forces behind price dispersion may help macroeconomists construct better cost shock propagation models and business cycle models, and surely will help microeconomists construct better models of imperfect competition.

Theoretical debates on price dispersion go back at least to Bertrand (1883) and Edgeworth (1925). Bertrand argued that buyer search in a posted offer market will enforce a unified competitive price even when there are only two sellers. The intuition is simply that *undercutting* is profitable at any higher price; at least one seller will be able to substantially increase volume and profit by slightly decreasing his price. Edgeworth noted that the outcome is more complicated when the sellers have binding capacity constraints at the competitive price. He

[1] Similar considerations apply to variants such as labor markets, in which buyers (firms) post price (wage) and sellers (workers) search. For expositional simplicity we focus below on simple seller posted offer markets and

assumed myopic price adjustment and predicted a price cycle in which sellers reduce price in small increments when there is excess capacity but jump to much higher prices when the capacity constraints bind. Modern textbook treatments such as Tirole (1988) downplay price cycles and instead focus on dispersed prices, in the sense of mixed strategy Nash equilibrium. This literature relies on capacity constraints despite their short-run nature, and does not explicitly model buyer choice despite its crucial role.

Stigler (1961) started a large literature on buyer choice when search is costly but the distribution of posted prices is exogenous and constant. Gastwirth (1976) pointed out that buyer behavior and payoffs are quite vulnerable to misperceptions of the price distribution.

Diamond (1971) apparently was the first to model simultaneous buyer search and seller price setting, and he reached a surprising conclusion. If all buyers have positive search costs (uniformly bounded below by an arbitrarily small positive number), then the unique Nash equilibrium is a unified price, but at the monopoly rather than the competitive level. The intuition is that when other sellers charge a price below the monopoly price, it is more profitable to choose a slightly higher price, given buyers' search costs; we refer to this as cream *skimming*. Several later authors find equilibrium dispersed prices in models with heterogeneous buyers. For example, Salop and Stiglitz (1977) derive equilibrium price dispersion in a model where some buyers are costlessly and fully informed while the other buyers have prohibitive search costs and are totally ignorant of posted prices. Stahl (1989) shows that if some buyers have zero search costs while others have identical positive search costs then there is a unique symmetric NE in mixed strategies. The NE price distribution changes continuously from the monopoly price (Diamond) to the competitive price (Bertrand) as the fraction of zero search cost buyers varies from 0 to 1.

---

ignore the variants.

Burdett and Judd (1983) model price dispersion in starkest form, without resorting to heterogeneous buyers or small numbers of sellers or capacity constraints. In the sequential search version of their model the buyers decide the number of prices to sample at given search cost per price, while in the noisy search version the sample size is random. Both versions have dispersed price equilibria that shift systematically with the parameters of the search technology. The intuition is that the incentives to undercut and to skim balance over a specific range of prices that varies with the parameters.

The empirical relevance of these models is open to question. Gastwirth's results suggest that buyers might not behave properly until they know the equilibrium price distribution precisely. The results of Hopkins and Seymour (1999) suggest that convergence to dispersed price equilibrium is unlikely even with large numbers of well-behaved buyers and sellers. They show that the unified price Diamond is dynamically stable but all dispersed price equilibria are unstable under a wide class of learning dynamics.[2]

In this paper we present a laboratory study based mainly on the noisy search version of Burdett and Judd (1983) model. We simplify by eliminating sample sizes of size three and larger (they play a negligible role anyway) and by normalizing sellers' production cost to zero and buyers' willingness to pay to 2.00. The simplifications involve no true loss of generality but allow us to solve the model explicitly, and to compare the theoretical predictions directly to the lab data. Our experiment varies the search technology parameters and also the number and type (human or automated) of buyers. Thus we have a clear environment for testing the effect of buyer search on price dispersion.

Two earlier laboratory studies suggest that the theory may not predict very well. Abrams

[2] Indeed, in private conversations, two eminent theorists (neither apparently aware of the Hopkins and Seymour paper) conjectured to us that stable dispersions would be unlikely to emerge in laboratory experiments.

et al. (2000) study a posted offer laboratory market with search, and examine only cases predicted to have unified prices at the Bertrand and Diamond equilibria. They find that most prices were closer to the halfway point than to either the Bertrand or the Diamond extreme, and the sample size parameter had much smaller impact than predicted. Davis and Holt (1996) obtained similar results. Both studies gave sellers a public identification number and neither told traders the distribution of prices chosen in previous trading periods. On the other hand, in three of four relevant trial sequences, Grether, Schwartz and Wilde (1988) obtained results closer to the Diamond extreme in an experiment with anonymous sellers and with public information on the previous period's price distribution. Since these features are consistent with relevant models, we shall employ them.

Brown Kruse et al. (1994) study the role of capacity constraints in oligopoly. Their laboratory environment features four human sellers facing a constant elasticity demand curve and a simple rationing rule. They find that Edgeworth's price cycle theory explains the data better than three alternatives (competitive equilibrium, dispersed prices as given by the symmetric Nash equilibrium in mixed strategies, and tacit collusion), although none does especially well.

We begin in the next section by writing out the noisy buyer search model and stating three formal results. The first result demonstrates the optimality of buyer reservation price strategies taken for granted by Burdett and Judd (1983), and the second result repeats one of their main findings. The third result and its corollary give apparently new explicit formulas for the distribution of dispersed prices as a function of search costs and other parameters. Section 3 describes the experiment and lays out the research hypotheses.

Section 4 presents the results. Overall the data conform to theory much more closely than in previous posted offer market experiments with buyer search. The data track the comparative static predictions on search cost and sample size parameters remarkably well, especially in the treatment with large numbers of automated buyers. Still, there are systematic departures from the theory. Especially with relatively few human buyers, we see a tendency of sellers to move prices in a parallel fashion, with occasional sharp deviations to higher prices...

A concluding section summarizes the results, offers some interpretations and implications, and lists possible avenues for future research. Formal proofs are collected in Appendix A. Instructions to subjects are attached as Appendix B. Details of the experimental procedures can be found in a companion paper, Cason and Friedman (1999), that uses a subset of the same data to study the exercise of market power.

2. An Equilibrium Model

Our experiment is based on streamlined version the Burdett and Judd (1983) noisy search model. The model assumes a continuum of sellers with zero cost for producing a homogeneous good, and a continuum of buyers with identical search cost $c \geq 0$ and with identical willingness to pay (say \$2.00) for a single indivisible unit of the good. Each seller posts a single price at which he is prepared to sell as many units as buyers order. Sellers maximize profit py (= revenue, since production costs are zero), where the sales volume y depends on the seller's posted price p, on other sellers' prices, and on buyer characteristics.

Each buyer initially has an independent sample of sellers' posted prices; the sample contains one price with probability $q \geq 0$ or two prices with probability $1- q \geq 0$. The buyer can quit, or can purchase at the (lower) observed price p, or can search by paying c to obtain a fresh sample. New samples again are independent and either of size 1 (with probability q) or size 2 (with probability $1- q$). The values of q and c are common knowledge. After $m \geq 0$ searches buyer's payoff is – *mc* if she quits, and is 2.00 – *p* – *mc* if she purchases at posted price *p* from the last sample; there is no recall from earlier samples. Buyers maximize expected payoff.

Buyers are endowed with a fixed belief *F* about the distribution of sellers' posted prices; *F* is not altered by observed samples. For now *F* is an arbitrary cumulative distribution function with support contained in $(0, \infty)$, i.e., *F* is right-continuous and increasing, and $F(p) = 0$ (resp. $=1$) at all points *p* below (resp. above) its support. Recall that the lower price in a sample of size 2 has distribution $(1 - (1-F(p))^2)$, so the overall distribution of the lowest price in a sample not yet drawn is $G(p|\, q) = qF(p) + (1 - q)(1 - (1-F(p))^2)$.

A buyer is said to follow a *reservation price strategy* if there is some $p^* \leq 2.00$ such that she purchases her unit at the lowest price *p* in the current sample if *p* is at or below p^*, and she searches again if *p* is above the reservation value p^*. A buyer's strategy is said to be *optimal* with respect to *c*, *q* and *F* if no other buy/search/quit plan yields higher expected payoff.

2.1 Analysis

Consider the equation

$$(1) \qquad c = \int_0^z (z-p)\,G(dp|\, q)\,.$$

The notation $G(dp|\, q)$ indicates a Stieltjes integral with dummy variable *p*, so the right hand side (RHS) of (1) is the incremental benefit of search, the expected price reduction, when the best current price is *z*. The left hand side is, of course, the incremental cost of another search.

Proposition 1. *Given search cost* $c>0$*, sample size parameter* $q \in (0,1)$ *and perceived posted price distribution F, there is a unique solution* $z=z^*(c,q,F) \geq 0$ to *equation (1). If the solution* z^* *is no greater than the willingness to pay 2.00, then it is optimal for each buyer to follow a reservation price strategy with* $p^*=z^*$.

All proofs are collected in Appendix A.

How should sellers choose price when all buyers follow identical reservation price strategies? The model assumes Nash equilibrium for a continuum of sellers, i.e., a distribution *F* of sellers' posted prices such that, given *q* and *p**, no seller can increase profit by unilaterally changing price. This comes down to an equal profit condition for all prices in the support of the distribution, which quickly leads to the following result.

Proposition 2. *Given sample size parameter q Î (0,1) and identical reservation price strategies with reservation value p**≤ 2.00 *for all buyers, there is a unique Nash equilibrium distribution F for sellers' posted prices. On its support interval [qp*/(2-q), p*], the distribution takes the value*

(2) $$F(p) = 1 + \left(1 - \frac{p^*}{p}\right)\frac{q}{2-2q}.$$

In equilibrium, the reservation price *z*(c,q,F)=p** is consistent with the distribution *F* and conversely. More formally, for given exogenous parameters *c* and *q*, say that (*F*, *p**) is a Burdett-Judd *Noisy Search Equilibrium* (NSE) if (a) *F* is a Nash equilibrium distribution with respect to *q* and *p**, and (b) *p** defines an optimal reservation price strategy with respect to *c*, *q* and *F*. Our main result is

Proposition 3. *For each sample size parameter q Î (0,1) and positive search cost c<2.00(1-q), there is a unique NSE* (*F, p**). *The NSE reservation price is p* = c/(1-q) and hence the NSE price distribution is* $F(p) = 0.5(2-q)/(1-q) - 0.5[cq/(1-q)^2]\, p^{-1}$ *on the support interval [cq/((2-q)(1-q), c/(1-q)].*

Corollary. *In NSE, buyers never search. For q Î (0,1) and c Î (0, 2.00(1-q)),*

1. *the NSE posted and transacted price densities, upper and lower endpoints, and mean and median are as listed in Table 1;*

2. *all sellers earn the same profit $mqc/(1-q)$, where m is the number of buyers per seller; and*

3. *the NSE demand for a seller posting price p is $y = m\,s(p)$, where the relative share is $s(p) = q + 2(1-q)(1-F(p))$ and F is the NSE price distribution given in Proposition 3.*

Note that in Table 1 the transaction prices are lower than the corresponding posted prices in the sense of first-order stochastic dominance. The intuition simply is that buyers who see two prices transact at the lower price.

For $0<q<1$ and moderate $c>0$, the NSE features dispersed prices, as illustrated in Figure 1. It is not hard to show (see Burdett and Judd) that prices are unified in the limiting cases. Specifically, if $q=0$ or if $c=0$, the NSE price distribution is degenerate at $p=0$ and we have the competitive or Bertrand equilibrium. If $q=1$ and $c>0$, the NSE distribution is degenerate at the buyers' true willingness to pay, $p=2.00$, and we have the monopoly or Diamond equilibrium.

2.2 Empirical Issues

The model says that the incentive to undercut dominates when $q=0$ (so every buyer sees more than one seller's price), and the incentive to skim dominates when the search cost c is positive (however small) and $q=1$ (so every buyer sees just one seller's price). For intermediate values of q and positive c, the incentives balance over a specific range of prices and we have price dispersion. As the parameters change, the balance shifts in the particular manner given in Table 1 and Figure 1.

The model is crisp and intuitive, but is it useful empirically? NSE is a static, one-shot equilibrium of a continuum economy. One might ask whether finite numbers of buyers and sellers interacting dynamically in a real market (field or lab) will converge to NSE. There are at least six empirical issues:

1. Real markets have only finite numbers of buyers, so sellers posting the same price often will

be sampled by different numbers of buyers and earn different profits. Sellers' differing experience is likely to encourage differing behavior and that may delay convergence to equilibrium. With more buyers per seller, sampling variability is reduced and convergence may be faster.

2. Human buyers are unlikely to follow strictly a reservation price strategy, much less identical reservation price strategies as called for in NSE. Again this source of variability may delay or bias convergence.
3. Human sellers are unlikely to draw prices randomly and independently each period from a common distribution. The same issue arises in zero-sum matrix games. O'Neill (1987) argued that laboratory evidence confirmed the Nash equilibrium theory because the overall choice frequencies approximated the Nash frequencies. Brown and Rosenthal (1990), on the other hand, showed that individual subjects' choices were correlated and biased, contrary to the best-response behavior presumed in Nash equilibrium. We will examine the correlations of sellers' prices across and within periods, but for us the primary question is whether NSE correctly predicts the overall price distribution and its response to parameter shifts.
4. As noted in the introduction, Hopkins and Seymour (1999) show that the unified price Diamond equilibrium is dynamically stable but all dispersed price equilibria are unstable under a wide class of learning dynamics, even granting the large numbers assumption.[3] Hence the dispersed price NSE may be unreachable empirically.
5. Finite numbers opens the possibility of strategic manipulation by sellers or buyers. By contrast, the NSE model assumes price-taking behavior, and therefore may give inaccurate

[3] Their argument applies directly to Burdett and Judd's nonsequential search model. Hopkins and Seymour appear to believe that their argument applies equally to the noisy search model, but we are not quite convinced. An important difference is that the nonsequential search model can have multiple equilibria (including a dispersed equilibrium as well as the Diamond equilibrium), while the NSE is unique.

predictions.

6. The NSE prediction regarding the law of one price is especially interesting and subtle. The law is predicted to hold only in the extreme cases ($c = 0$, or $q = 0$ or 1) and the unified price in these cases is also extreme at $p = 0$ (the competitive price) or $p = 2.00$ (the monopoly price). Such extreme outcomes are hard to obtain in laboratory (or field) markets. On the other hand, the intermediate cases ($c > 0$ and $0<q<1$) could well produce a unified price at some intermediate level each period rather than dispersed prices.

3. Laboratory Procedures

The experiment uses the posted offer market institution described in Cason and Friedman (1999, p. 73-77). We summarize the key features here for completeness. As shown in Figure 2, each period each seller uses a scroll bar to post a single price. Each buyer sees one or two of the sellers' current posted prices, and clicks the appropriate button to indicate her choice: quit for that period ("Reject"), accept a posted price she sees, or search. If the "Search" button is clicked, the search cost is immediately deducted from the buyer's profit, a fresh sample of one or two posted prices is displayed and the buyer again has the same three choices. After all buyers have transacted or quit, the buyers and sellers review their own profit or loss and see all posted and transaction prices for the past period. Prices are displayed in random order to obscure seller identities. The next period then begins with new random matching of buyers to sellers.

Subjects were recruited from undergraduate classes in Economics and Biology at UCSC and Purdue. They received the written instructions attached as Appendix B, were assigned randomly to buyer and seller roles, and at the end of the session received total profits, on average about \$20 for sessions that lasted about 90-100 minutes. Sellers had zero production cost and no capacity constraint. Buyers had an induced $2.00 value for a single indivisible unit. Seller costs, buyer values, and the treatments described below were posted on the blackboard, displayed on traders' screens and announced in the instructions.

3.1 Treatments

Our experiment features three treatment variables: search cost c, sample size q and buyer population. The search cost is controlled at two levels: 20 cents and 60 cents. The probability q that any sample of prices has only one seller's price (rather than two different sellers' prices) is controlled at four levels: q = 0, 1/3, 2/3 and 1. When a buyer searches, the new sample is drawn with replacement of earlier samples. The values are chosen to provide a good separation of predicted (NSE) prices, as shown in Table 2.

The buyer treatments vary the number of buyers per seller and the buyer search strategy in order to treatments address the first two empirical issues raised earlier. The baseline buyer population is six human buyers. The alternative treatments replace the human buyers by computer algorithms (or "robots") that follow NSE reservation price strategies. The robot implementations are straightforward, with the number of robots controlled at 6, 12 and many. For example, with 12 robot buyers, c=0.60 and q=1/3, the equilibrium reservation price is $0.90 and (with six sellers) m =2. We tell the sellers that there are no human buyers and there are 12 automated buyers, each of which immediately accepts the lowest price it sees if it is below $0.91 and otherwise keeps searching.

There were two modifications to this general procedure. First, to avoid trivialities when q=0, we set the reservation price to ten cents instead of zero. Second, the many-robots treatment (described to subjects as 600 robots) uses the continuum formula given at the end of the Corollary. Thus a seller posting price p_i sells $y_i = m\, s(p_i)$ units, where m =100 and

(3) $s(p_i) = q + 2(1- q)(1-F^a(p_i)) = q + 2(1- q)(r_i-1)/(n-1)$.

Thus seller i sells to a fraction of the m 2(1- q) potential buyers who see two prices; the fraction can be expressed as in the Corollary using the actual (empirical) distribution F^a for that period, or can be expressed directly in terms of the rank r_i of his price p_i among the $n \leq 6$ sellers currently posting prices at or below the reservation price p^*.

3.2 Design

In pilot tests of the new user interface we varied both the number of buyers and sellers, but the experiment is designed to hold constant the number of sellers at six. Each run (usually 20-30 consecutive periods) holds constant all treatments, but the value of q switches across runs in a balanced fashion. The buyer treatment (with one human plus three robot conditions) and the search cost treatment are constant within each session but vary across sessions.

Table 3 lays out the design of our 18 sessions. It approximates a factorial design with 4 buyer conditions $\times$ 2 search cost conditions. Most of the eight cells include at least one session at each site (indicated by the UC- or PU- prefix for UCSC and Purdue respectively) and most cells include one or more sessions with subjects experienced in a previous session (indicated by an –x suffix). The numerical part of the session name indicates the calendar sequence, essentially random. Most sessions have four runs, with the first two using the extreme values of the q parameter. Experienced sessions required less instruction time, so we were able to use longer runs or more runs. There are a few minor irregularities, e.g., one or two extra or fewer periods in the human buyer sessions due to time constraints. Also, an early many-robot session used 45 robots instead of the later standard 600 robot algorithm, and the first experienced session had missing subjects and so was run with 5 human buyers and sellers rather than 6. A number of diagnostic tests on price distributions and other outcomes disclosed no impact of these irregularities on our conclusions.

4. Results

We begin in Section 4.1 with graphs of the raw data and summary statistics. Tests of the main comparative static predictions are presented in Section 4.2. Section 4.3 compares actual price distributions to the exact and "noisy" equilibrium predictions in the dispersed price treatments. Section 4.4 examines the behavior of individual buyers, and Section 4.5 looks at individual sellers.

4.1 Overview

Figures 3, 4 and 5 summarize some illustrative runs. The solid circles represent posted prices that result in transactions, and the open circles represent unaccepted price offers. Figure 3 shows the third run of session UC9x, featuring 6 experienced human sellers and a continuum of robot buyers that followed equilibrium reservation price strategies for a 20 cent search cost. The equilibrium in this third run (q=1/3) has prices in the interval [0.06, 0.30], and all posted prices fall into the predicted interval except one that misses by a penny. The distribution, however, seems to oversample the higher prices and there appears to be cross sectional and serial correlation.

Figure 4 summarizes the q=2/3 run of session PU2, which featured 12 robot buyers and 60-cent search costs. Most prices are inside the predicted range (now [0.90, 1.80]); the main exception is a single seller pricing about 10 cents too low. The first half of the run looks like the beginning of an Edgeworth cycle, as mean prices decline gradually via undercutting until near the bottom of the price range when sellers begin to favor skimming. But the "cycle" does not repeat itself and prices look dispersed in the second half of the run. Similar patterns can be seen in the q=1/3 and q=2/3 runs with human buyers, as illustrated in Figure 5.

Graphs of the other 71 runs show considerable variability but suggest some general tendencies. The extreme unified price predictions hold up quite well for q=0 and 1 runs with many robot buyers, but prices are less extreme with fewer (and especially human) buyers. The dispersed price predictions also seem to fare well in most sessions with many robot buyers and in some (but not all) sessions with fewer robots and humans. Prices seem correlated in varying degrees across buyers and across time.

Figure 6 illustrates the overall price trends. Panel A shows mean transaction price by period and treatment in all sessions with human buyers. Mean prices in the first few periods are rather tightly clustered in a 20 cent interval but gradually spread out with the predicted ordering: highest in the q=2/3, c=60 treatment and lowest in the q=1/3, c=20 treatment, with the other two treatments in between and close together. The separation never quite reaches the predicted degree, however, and narrows somewhat towards the end, mainly from a decline the last 10 periods of the q=2/3, c=60 runs. Panel B shows the corresponding data for the sessions with many robot buyers. Here the magnitude of separation as well as the ordering is about right, and the final mean prices are only a few cents above the NSE predictions. Mean transaction prices for the 6 and 12 robot buyer treatments shown in panels C and D are quite similar to panel B. Mean posted prices are a bit higher and more variable but generally parallel to the mean transacted prices; the figures are omitted to conserve space.

It is standard in laboratory markets to measure the efficiency with which potential gains from trade (or surplus) are actually realized. Our experiment has a $2.00 potential surplus for each buyer each period, which is dissipated when buyers search or fail to transact. Sellers almost never price above robot buyers' reservation prices, so efficiency is virtually 100 percent in all robot buyer sessions. Table 4 shows that efficiency is usually high in the human buyer sessions as well, averaging in excess of 97 percent in the q<1 treatments. Gross trading efficiency (ignoring search costs) exceeds net efficiency only slightly because buyers rarely search. Efficiency is lowest, but still averages over 92 percent, in the extreme q=1 treatment, perhaps because human buyers resist the equilibrium that awards all surplus to sellers.

Table 5 shows that sellers' share of this high trade surplus is usually greater than predicted by the NSE model, except when more than half of the $2.00 surplus is predicted to accrue to sellers. In the q=1/3 and q=2/3 treatments (Panel A), sellers' share of the exchange surplus nearly always exceeds the NSE level. Nevertheless, the NSE model is able to describe the main differences in seller earnings across treatments. The data provide remarkably strong support for the highly asymmetric predicted distribution of exchange surplus with q=0 (Panel B). When q=0 sellers receive only about one or two percent of the exchange surplus, except for the human buyers treatment. Even for the human buyers treatment, the mean profits shown in Table 5 are misleadingly high because they are influenced by outliers; median seller profits in the human buyers treatment with q=0 are 5 cents with 20-cent search costs and are 34 cents with 60-cent search costs.

4.2 Comparative Statics

The main predictions from the equilibrium model are how transaction prices vary with search cost c and sample size 2-q. Table 6 summarizes the evidence, using simple and very conservative nonparametric Wilcoxon tests to compare mean transaction prices in pairs of runs. For tests on the within-session treatment variable q each session contributes one independent observation—the difference in mean prices between q runs. For tests on the across-session search cost treatment variable we pair sessions with identical buyer types and (where possible) experience conditions. Each pair of sessions contributes a single observation.

As predicted by the NSE model, there is an insignificant (4 ± 5 cents) difference in transaction prices between high and low search cost runs when q = 0 and also when q = 1 (now – 1 ± 12 cents). In other comparisons, NSE predicts significant differences and the predictions are remarkably accurate. For example, the predicted difference between high and low cost when q = 1/3 is 20 cents and the actual is 29 ± 3, and the predicted difference for q = 2/3 is 80 cents and the actual is 59 ± 12. In every case but one, the differences are significant and in the predicted direction. The only exception is the low search cost q = 1/3 vs. q = 0, where the actual difference of 8 ± 8 cents is insignificant but still very close to the 9 cent predicted difference. Actual differences are somewhat smaller than the very large predicted differences in the q = 2/3 vs. q = 1 comparisons mainly because the mean prices for q = 1, although nearly equal at $1.64 and far higher than the other mean prices, still fall short of the extreme prediction of $2.00. One should bear in mind that q = 1 runs always either begin the session or follow q = 0 runs, so the mean price calculation includes trials when the price is low but rapidly increasing.

4.3 Price Distributions

The histograms in Figures 7 and 8 confirm that the predicted and actual distributions shift in parallel as we vary the q and c parameters. But the histograms also suggest some systematic prediction errors. For the human buyer sessions (Figure 7), the observed range of prices is wider than predicted, usually extending more towards the middle and away from the 0.01 and 2.00 extremes. For the many robots sessions (Figure 8), the predicted price ranges seem quite accurate, but not the predicted skewness.

Table 7 more carefully examines the location and skewness of actual prices where the NSE predictions should be sharpest, in the final 10 periods of each run. As suggested by earlier tables on profits and transaction prices, the posted prices track the predictions rather well but are biased towards 100 cents. (The only exception is rather minor: the q=2/3, c=60 human buyer median and mean fall one or two cents on the wrong side of 100.) The bias in both mean and median is always largest in the six Human Buyers condition and often (but not always) smallest in the Many Robot Buyer condition. Actual skewness often falls below the NSE values, although 11 of the 16 skewness estimates have the predicted positive sign. A more robust directional measure of skewness is the sign of the mean minus the median. This difference has the predicted positive sign in 14 of 16 cases. Thus the skewness predictions do better than the histograms suggest.

Not surprisingly, tests such as Kolmogorov-Smirnov (K-S) strongly reject the null hypothesis that the data come precisely from the NSE distribution. But perhaps the discrepancy is mere noise. To test this conjecture, we first note that the noisy price $x = p + z$ has density $f \bullet g = h(x) = \int_{-\infty}^{\infty} f(y) g(x-y) dy$, where the noiseless density $f(p)$ is given in Table 1 and the noise density $g(z)$ is assumed to be mean zero Normal with unknown variance s^2. Thus the noisy NSE price density is

$$h(x) = k \int_a^z \frac{1}{y^2} \exp\left(-\frac{(y-x)^2}{2s^2} \right) dy, \text{ where } k = \frac{cq}{2(1-q)^2 \sqrt{2p}\, s}$$

and a and z are given in Table 1. We solve for $h(x)$ numerically for s ranging between 1 and 30 cents and truncate to the interval [0, 200], as illustrated in Figure 9 for case q=2/3 and search cost=20 cents. Increasing the noise level s increases the mass outside the NSE range [a, z] and decreases skewness. We search for the value of s that best fits the empirical price distribution function, measuring fit (as in K-S) by the maximum difference between the empirical and noisy cumulative distribution functions. Separate comparisons are made for each buyer type, q, and search cost treatment. The best fit involves a modest level of noise when q=2/3—often less than s =15. For the more theoretically skewed distributions with q=1/3, by contrast, the best fits typically involve s greater than 15. Human buyer sessions generally produced higher estimates of s than sessions with robot buyers. The noisy densities are never skewed left, so they also provide poor fits to empirical distributions that oversample higher prices.

So far we have looked at empirical distributions that pool observed prices across periods. Such pooling cannot distinguish between true price dispersion and the alternative prediction, noted at the end of the theory section, that there is a unified price each period that changes across periods. To make the empirical distinction, we measure dispersion separately each period as the standard deviation of posted prices across sellers. Table 8 reports a robust summary statistic, the median dispersion measurement over the final 10 periods in each run. The results confirm NSE theory quite nicely. Dispersion is very low where NSE calls for unified prices—in the many robot buyer treatments with q=0 and q=1—and is about the right magnitude where NSE predicts dispersion. We also ran a regression (not shown) of the dispersion measure on various explanatory variables including the treatments, their interactions and time. The estimates indicate that dispersion tends to decline in early periods of a run.

4.4 Buyer Behavior

The results so far broadly support the main NSE predictions but include several interesting discrepancies. To deepen our understanding of the data, we now explore individual behavior, beginning with human buyers.

Table 9 offers estimates of all human buyers' reservation prices in the q = 1/3 and 2/3 runs. The maintained assumption is that each buyer has a constant reservation price but may occasionally err either (a) by purchasing at a higher price or (b) by not purchasing (quitting or searching) at a lower price. The estimates are the lowest reservation price that minimizes total errors subject to the constraint of equal numbers of type (a) and (b) errors.[4] The number of periods—and therefore the number of opportunities to violate the reservation price rule—varies across sessions, but is typically about 30 (see Table 3). For 48 of the 70 buyers we are able to identify reservation prices with no violations of either type.

The table indicates that the estimated reservation prices are less extreme than predicted. A large majority of the estimates are between the midrange (100) and the MSE prediction, and in 69 of 70 cases, the estimates are above the predictions for less than 100 and below the predictions for more than 100. (The one exception, inexperienced Buyer 5 in PU7 for q=1/3, is off by only 5 cents.) Equally important, the reservation price estimates are consistent with comparative static predictions in that they are higher for higher q (with only 6 exceptions) and for higher search costs. The estimated reservation prices differ somewhat by session, which is not predicted by NSE but is consistent with session differences in posted prices. For example, posted prices rarely exceed 100 in UC5 but do so frequently in PU5, consistent with the estimates reported near the top of Table 9.

We should also note that buyers occasionally quit—that is, refuse to buy in a particular period at any available price. Overall we observe 58 quits, compared to 277 searches and 3512 purchases. Posted prices almost never exceed the buyers' value of the good ($2.00), so most of the quits are demand withholding, analogous to rejected proposals in ultimatum games. Consistent with ultimatum game results and with Ruffle (2000), quits are most common at high prices that leave little surplus to buyers. Two-thirds of the quits (39 out of 58) occur in the q=1 treatment in which buyers receive no surplus in NSE, and also about two-thirds occur at posted

---

[4] For 6 of the 70 buyers we cannot identify reservation prices with an equal number of violations of each type, so we report reservation prices that minimize the difference between the violation counts of each type. Results are

prices that offer buyers less than one-third of the surplus.

4.5 Seller Behavior

Are the observed price dispersions symmetric across sellers, supported by independent, identically distributed draws from a common distribution like that in Table 1, or are they asymmetric with (for example) some sellers at a given moment favoring undercutting strategies and other sellers favoring skimming? To address this question, we first estimated the simple autocorrelation coefficient corr(p_{it}, p_{it-1}) for each seller *i* separately for the *q*=1/3 and 2/3 treatments. The mean autocorrelation across the 212 estimates is 0.455, and 193 of the estimates are greater than zero. We used a Monte Carlo simulation to determine the 95th percentile of the estimated correlation coefficients based on iid price draws from the equilibrium distributions. Table 10 shows that 164 of the 212 sellers' estimated correlation coefficients exceed the relevant 95th percentile. Therefore we reject the null hypothesis of iid draws from the equilibrium distribution. The rejection rate is highest for human buyers and for *q*=1/3.[5]

Prices are also correlated across sellers within a given period. We show this by estimating the correlation coefficient corr(p_{it}, $\bar{p}_{-it}$), where $\bar{p}_{-it}$ denotes the mean price posted by the sellers other than seller *i* in period *t*. We estimated this correlation separately for each seller for each *q*=1/3 and *q*=2/3 run. We again use a Monte Carlo simulation to estimate the distribution of correlation coefficients according to the NSE null hypothesis, again centered at zero. Table 11 shows that 159 of the 212 sellers' estimated correlation coefficients exceed the relevant 95th percentile and therefore reject the null hypothesis of iid draws from the equilibrium distribution. The patterns are less pronounced here than in the previous table, but it appears that the rejection

---

substantially unchanged for reservation price estimates that simply minimize the total number of violations.

[5] We also estimated a simple cross-sectional OLS regression using the autocorrelation estimate as the dependent variable, with the various experimental treatments as explanatory variables. The estimates indicate that the autocorrelation is also higher with experienced subjects.

rate is lowest for the 6 robot buyers treatment and is lower when *q*=2/3 in the 20-cent search cost treatment.[6] Brown Kruse *et al*. (1994) also found that prices in their experiment were significantly correlated both across time and across sellers. Clearly there are non-trivial price dynamics here not captured in the static NSE model.

Next consider Figure 10, which shows that changes in posted price are not quite symmetric. Overall one sees more small price decreases balanced by less frequent large price increases, and the asymmetry is especially striking for the many robot sessions. This pattern is consistent with the impressions gleaned from Figures 3-5, and with the observation that undercutting suggests many small decreases but skimming suggests an occasional large increase. Table 12 shows that price decreases are more frequent for sellers who sold no units last period, especially for low values of *q*. In particular, the ratio of price decreases to increases when a seller fails to sell any units last period falls monotonically from four for *q*=0 to one for *q*=1.

An exhaustive evaluation of price dynamics and learning models is beyond the scope of this paper, but we will briefly explore two approaches, a myopic deterministic price adjustment model and a stochastic logit best response model with adaptive expectations. For the myopic price adjustment model, define the Edgeworth price P_{it}^{E} as seller *i*'s profit maximizing price assuming that rival sellers all maintain their prices from the previous period *t*-1. Following Brown-Kruse *et al.* (1994) we estimated a linear equation in which the dependent variable is the current price change, and the main explanatory variable is the current Edgeworth price minus the actual price chosen in the previous period. We also included the lagged value of this independent variable. When pooling all periods with *q*=1/3 and *q*=2/3, we estimate

[6] A simple cross-sectional OLS regression using the correlation coefficient as the dependent variable again indicates that across-seller price correlation increases with experience.

(4) $$P_{it} - P_{it-1} = -1.25 + 0.17\left(P_{it}^{E} - P_{it-1}\right) - 0.06\left(P_{it-1}^{E} - P_{it-2}\right);$$ Adj. R^2=0.084; N=6039.
(0.19) (0.01) (0.01)

Standard errors are shown in parentheses, and all three coefficients are highly significant in this pooled dataset. The Edgeworth coefficient estimate is very significantly positive but far below the complete adjustment level of one, indicating only a partial adjustment.[7] This result is similar to Brown Kruse *et al*. and shows some tendency of sellers to adjust price in the direction of last period's *ex post* optimal price.

The next model rectifies two limitations of this simple approach. It considers all available price information, not just prices in the last period or two, and it smoothes the deterministic but discontinuous decision rule. For example, if an undercutting strategy is expected to be almost as profitable as a skimming strategy, it would not be chosen in the simple approach but would be chosen almost as often in the next model. This smoothed decision (or noisy best response) approach is increasingly popular and underlies the Quantal Response Equilibrium by McKelvey and Palfrey (1995) and Chen, Friedman and Thisse (1997).

The key variable is $p_i^e(p)$, the seller *i*'s expected payoff from posting price *p*. The expectation obviously depends on the rivals' posted prices. Following Capra et al. (1999a, 1999b) we assume uniform initial beliefs over the set of relevant prices and update using "geometric fictitious play," adaptive expectations based on all observations of rivals' earlier choices. Sellers update last period's weight *w* of an unchosen price to rw and update the weight of an observed price to rw+1. (If two rival sellers post the same price, the weight is increased by two instead of one.) The discount parameter $r \leq 1$ specifies the degree to which old observations

[7] We also estimated this equation separately for the 4 buyer populations × 2 *q* values (1/3 and 2/3) × 2 search cost treatments. The estimated Edgeworth coefficient was significantly positive in all 16 cases, although it was always closer to zero than to one. Thirteen of the estimated coefficients on the lagged term were not significantly different from zero, with the other three estimates significantly negative.

are less important than new. Belief probabilities are calculated each period by dividing the weight by the sum of all weights. A seller *i* who posts price *p* and has such beliefs summarized in the cumulative distribution function *F* earns expected profit $p_i^e(p)$ = *qp* + 2*p*(1-*q*)[1-*F*(*P*)].

The expected profits determine choice probabilities using the convenient logit rule:

$$(5) \qquad P_i(j) = \frac{e^{l p_i^e(j)}}{\sum_{\text{all } k} e^{l p_i^e(k)}}.$$

The parameter l captures the sensitivity of choices to expected payoffs; the price with highest expected payoff is increasingly likely as l becomes larger. The indexes *j* and *k* denote prices in one or two cent intervals, depending on the range of feasible prices.

We set r=0.75 and estimate l using standard maximum likelihood techniques separately for each dataset, as shown in Table 13.[8] In all cases the l estimates are significantly positive, indicating that sellers indeed are more likely to choose prices that have (according to our calculations) higher expected profits. The magnitude of the estimates is rather small, however, indicating that the "best" prices are not chosen with an overwhelmingly high probability. The median l estimate over the 8 datasets with 20-cent search costs is 0.174, and over the 8 datasets with 60-cent search costs is 0.058. These medians bracket the estimates of 0.092 and 0.119 provided by Capra et al.[9] In interpreting the results, one should bear in mind that the magnitude of l depends inversely on the scale of expected profits. Expected seller profits are substantially higher with higher search costs, which is a main explanation for the lower estimated l. The l estimates are uniformly lower with human buyers and are uniformly higher with 12 or many

[8] We fix the learning parameter r=0.75 because this is similar to the estimated r=0.72 reported in Capra et al. (1999b). Alternative estimates that fix r=1 indicate that the qualitative conclusions are not sensitive to changes in this range of r.

robot buyers, which suggests that the sampling variability introduced by small numbers of buyers and human buyers with unobserved reservation prices make seller price choices more noisy.

A simulation indicates that this learning model can explain some regularities in the data not captured in the NSE model. The simulated sellers begin each run with uniform beliefs regarding the price choices of other sellers, and they updated these beliefs (and expected profits) based on others' actual price choices using the geometric fictitious play rule. For this simulation we used l=0.1, r=0.75 and search cost=20 cents. We ran the simulations for 30 periods and simulated three sessions for *q*=1/3 and three sessions for *q*=2/3, in parallel to the actual data. The simulated price distributions, like the actual distributions, are less skewed than the NSE distributions. Over the last 10 periods of these simulated 20-cent search cost runs, for *q*=(1/3, 2/3), respectively, mean posted prices are (16.6, 43.6), median posted prices are (17, 45) and posted price skewness is (-0.067, -0.912). These summary statistics are closer to the observed values than the NSE predictions in the top half of Table 7. The simulation also leads to a positive bias in the cross-sectional correlation in posted prices, unlike the zero expected correlation of the NSE (Table 11). This is because the sellers' adaptive beliefs in this model are correlated. The observed correlations substantially exceed the simulated correlations, however.[10] The simulation, like the NSE, predicts zero price autocorrelation within sellers and therefore cannot generate the significantly positive autocorrelation noted in Table 10.

5. Discussion

Noisy Search Equilibrium (NSE) is a wonderful theoretical model, stark in its assumptions, crisp in its predictions, and intuitive in the role it assigns to buyer search in

---

[9] Capra et al. actually report estimates of an error parameter $m = 1/l$.

[10] In particular, the median observed cross-sectional price correlation is 0.51, but the median simulated cross-

promoting or preventing price dispersion. Our experiment provides the clearest evidence to date on its empirical validity, and the model does remarkably well, far better than we anticipated. NSE correctly predicts when the laboratory markets converge to a unified price and when dispersion persists. It also predicts rather well the high efficiency and the observed range of prices, the central tendency as well as a measure of dispersion. The NSE comparative statics indeed capture the impact of the buyer search parameters (sample size 2-*q* and search cost *c*). The predictions are especially accurate in the many robot buyer treatment, which eliminates extraneous (with respect to NSE) variability arising from sampling variance and buyer idiosyncrasies. Even the NSE predictions of human buyer behavior hold up well: our estimates of reservation prices account for the vast majority of actual choices and they shift with the search parameters in the predicted directions.

Good theories accurately predict large scale empirical regularities, but even the best theory will eventually break down when pressed hard enough. It is instructive to see where and how the breakdown occurs. We observed several interrelated departures from NSE predictions. Although it was much less severe than in previous studies such as Abrams et al. (2000), we also observed a bias in prices towards the center of the overall price range, i.e., towards equal splits of the 2.00 surplus. The bias was clearest with human buyers and with search parameters (q=1, q=2/3 and c=60 cents) that award little or no surplus to buyers in NSE. Actual price distributions are somewhat less skewed than in NSE. Perhaps more important and contrary to NSE, there is considerable price correlation across sellers and across time. Sellers balance many small price decreases (especially when they have few recent sales) with occasional large increases.

These anomalous empirical regularities open the way to new theory and new empirical work. On the theoretical side, one could extend the static model by incorporating seller risk

---

sectional price correlation is only 0.10. Recall that this correlation is zero in the NSE.

aversion and possibly explain the price bias in the cases that NSE awards little surplus to buyers. One could try to model price dynamics using more sophisticated learning models than we explored here. Autocorrelated posted prices and the dependence on recent sales, for example, might be captured by individual differences in learning parameters or perhaps by differing experience even in a reinforcement learning model. It will be interesting to see whether one gets Edgeworth-like price cycles from such models. One might also try to adapt the behavioral models recently used to explain ultimatum game data. Such models might account for the differences observed across sessions and the bias away from prices that result in highly asymmetric distributions of exchange surplus.

On the empirical side, one might try to see whether the NSE buyer search parameters can be mapped onto different search technologies available in various field markets (e.g., in the available internet search engines), and if so, whether they can account for the observed degree of price dispersion. New laboratory experiments also are in order. One could investigate the gap between human and robot buyers, using Selten's "strategy method" for buyers. That is, before seeing posted prices, each buyer would enter a reservation price that commits him to the corresponding accept/search decisions. (Two practical problems to consider: if the reservation price is below all the posted prices when $c>0$, this algorithm will bankrupt him. Also, some traders might try to use this as a signaling device.) One could use methods similar to ours to investigate the sequential search model in which the sample size is endogenous. Finally, an extension we plan to undertake is to examine "customer markets" in which buyers observe costlessly the price of the seller from whom they last purchased but have to pay a switch cost to find an alternative seller. Such markets are ubiquitous in the modern world but the theory is less well developed.

References

Abrams, Eric, Martin Sefton and Abdullah Yavas (2000), "An Experimental Comparison of Two Search Models," *Economic Theory*, forthcoming.

Baye, Michael and John Morgan (1999), "Information Gatekeepers and the Competitiveness of Homogeneous Product Markets," Indiana University (Kelley School) working paper.

Bertrand, Joseph (1883), "Theorie Mathematique de la Richesse Sociale," *Journal des Savants,* pp. 499-508.

Borenstein, Severin and Andrea Shepard (1996), "Dynamic Pricing in Retail Gasoline Markets," *RAND Journal of Economics* 27, pp. 429-451.

Brown, James and Robert Rosenthal (1990), "Testing the Minimax Hypothesis: A Re-examination of O'Neill's Game Experiment," *Econometrica* 58, pp. 1065-1081.

Brown Kruse, Jamie, Steven Rassenti, Stanley Reynolds and Vernon Smith (1994), "Bertrand-Edgeworth Competition in Experimental Markets," *Econometrica* 62, pp. 343-371.

Brynjolfsson, Erik and Michael Smith (1999), "Frictionless commerce? A Comparison of Internet and Conventional Retailers," MIT (Sloan) working paper.

Burdett, Kenneth and Kenneth Judd (1983), "Equilibrium Price Dispersion," *Econometrica* 51:4, pp. 955-969.

Capra, C. Monica, Jacob Goeree, Rosario Gomez, and Charles Holt (1999a), "Anomalous Behavior in a Traveler's Dilemma?," *American Economic Review* 89, pp. 678-690.

Capra, C. Monica, Jacob Goeree, Rosario Gomez, and Charles Holt (1999b), "Learning and Noisy Equilibrium Behavior in an Experimental Study of Imperfect Price Competition," University of Virginia working paper.

Cason, Timothy and Daniel Friedman (1999), "Customer Search and Market Power: Some

Laboratory Evidence," in *Advances in Applied Microeconomics,* Vol. 8, M. Baye (ed.), Stamford, CT: JAI Press, pp. 71-99.

Chen, Hsiao-Chi, James Friedman and Jacques-François Thisse (1997), "Boundedly Rational Nash Equilibrium: A Probabilistic Choice Approach," *Games and Economic Behavior* 18, pp. 32-54.

Davis, Douglas and Charles Holt (1996), "Consumer Search Costs and Market Performance," *Economic Inquiry* 34, 133-151.

Diamond, Peter (1971), "A Model of Price Adjustment," *Journal of Economic Theory* 3, pp. 156-168.

Edgeworth, F. Y. (1925), "The Pure Theory of Monopoly," in *Papers Relating to Political Economy*, vol. 1, by F. Y. Edgeworth (New York: Burt Franklin).

Gastwirth, Joseph L. (1976), "On Probabilistic Models of Consumer Search for Information," *Quarterly Review of Economics*, 90:1, 38-50.

Grether, David, Alan Schwartz and Louis Wilde (1988), "Uncertainty and Shopping Behavior: An Experimental Analysis," *Review of Economic Studies* 55, 323-342.

Hopkins, Ed and Robert Seymour (1999), "The Stability of Price Dispersion under Seller and Consumer Learning," Edinburgh University working paper.

O'Neill, Barry (1987), "Nonmetric Test of the Minimax Theory of Two-Person Zerosum Games," *Proceedings of the National Academy of Sciences*, U.S.A. 84, pp. 2106-2109.

McKelvey, Richard and Thomas Palfrey (1995), "Quantal Response Equilibria for Normal Form Games," *Games and Economic Behavior* 10, pp. 6-38.

Ruffle, Bradley (2000), "Some Factors Affecting Demand Withholding in Posted-Offer Markets," *Economic Theory*, forthcoming.

Salop, Steven and Joseph Stiglitz (1977), "Bargains and Ripoffs: A Model of Monopolistically Competitive Price Dispersion," *Review of Economic Studies* 44, pp. 493-510.

Stahl, Dale (1989), "Oligopolistic Pricing with Sequential Consumer Search," *American Economic Review* 79, pp. 700-712.

Stigler, George (1961), "The Economics of Information," *Journal of Political Economy* 69, pp. 213-225.

Tirole, Jean (1988), *The Theory of Industrial Organization*, Cambridge, MA: MIT Press.

Appendix A. Mathematical Details

Proposition 1. *Given search cost* $c>0$*, sample size parameter* $q \in (0,1)$ *and perceived posted price distribution F, there is a unique solution* $z=z^*(c,q,F)\geq 0$ to *equation (1). If the solution* z^* *is no greater than the willingness to pay 2.00, then it is optimal for each buyer to follow a reservation price strategy with* $p^*=z^*$.

Proof. The RHS of (1) can be re-written as $\int_0^z G(p\,|\,q)dp$; one integrates by parts, noting that the term $(z-p)G(p|\,q)$ is zero at $p=0$ and at $p=z$. Hence the derivative of the RHS is $G(z|\,q)$. It follows that the RHS is strictly increasing in z on (and above) the support of F (or G), i.e., at all points where its value is positive. The RHS is zero at $z=0$, and is unbounded as $z \rightarrow \infty$. The LHS of (1) is positive and constant, so the equation has a unique solution z^* by the Intermediate Value Theorem.

To see that there is an optimal reservation price strategy, note that the buyer's expected payoff when she decides to search is some value $V^S(c, q, F) \leq 2.00 - c$, and V^S is independent of the current lowest price p because there is no recall and no learning about the distribution. When she decides not to search her expected payoff is $V^N(p) = \max\{0, 2.00-p\}$. Since V^N is a continuous decreasing function of p with $V^N(0) = 2.00 > V^S$, we conclude that the set of prices p for which not searching is better than searching, i.e., for which $V^N(p) \geq V^S$, is an interval $[0, p^*]$.

To complete the proof, it remains only to show that $p^* = z^*$ whenever $p^* \leq 2.00$. If $p^*>2.00$, then we must have $V^S < 0$, i.e., the buyer will never find it worthwhile to search and will accept the lowest currently available price if it is below 2.00 and will otherwise quit. If $p^* \leq 2.00$ then the buyer is indifferent at posted price p^* between searching and not searching. But this indifference condition is exactly equation (1) with $z = p^*$. QED.

Proposition 2. *Given sample size parameter q Î (0,1) and identical reservation price strategies with reservation value p** ≤ 2.00 *for all buyers, there is a unique Nash equilibrium distribution F for sellers' posted prices. On its support interval [qp*/(2-q), p*], the distribution takes the value*

$$(2) \qquad F(p) = 1 + \left(1 - \frac{p^*}{p}\right)\frac{q}{2-2q} \, .$$

Proof. This result follows directly from Burdett and Judd (1983, Claim 2, p. 963). For completeness we reproduce here the derivation of equation (2). All sellers must have the same nonnegative profit in NE (nonnegative since a seller can ensure zero profit by posting a price above *p**). A seller posting price exactly at the upper endpoint *p** of the distribution sells only to the fraction *q* of buyers who have sample size 1, and thus earns *p*mq*, where *m* is the number of buyers per seller. Sellers choosing a lower price *p* trade off the larger sales volume (due to pricing below a rival when the buyer sees two prices) against smaller profit margin, and thus earns *pm(q + 2(1-q)(1-F(p))*. That is, the seller retains all of the *q* buyers who see one price but doesn't retain the fraction *F(p)* of the *2(1-q)* buyers who see a second price that is lower than *p*. One gets (2) by equating the two earnings expressions and solving for *F(p)*. QED.

Proposition 3. *For each sample size parameter q Î (0,1) and positive search cost c<2.00(1-q), there is a unique NSE* (*F, p**)*. The NSE reservation price is p* = c/(1-q) and hence the NSE price distribution is* $F(p) = 0.5(2-q)/(1-q) - 0.5[cq/(1-q)^2]\, p^{-1}$ *on the support interval [cq/((2-q)(1-q), c/(1-q)].*

Proof. Existence and uniqueness follow from Propositions 1 and 2. One obtains the value of *p** by inserting (2) into (1) with *p*=z,* and integrating by parts. Using the notation *a*=*zq*/(2-*q*) for the lower endpoint of the interval and *b*=*q*/(2-2*q*) for the coefficient of *(1 – p*/p)* in equation (2), we have

$$c = \int_0^z (z-p)dG(p) = (z-p)G(p)\,\Big|_0^z + \int_0^z G(p)dp$$

$$= 0 + \int_0^z qF(p) + (1-q)\left(1-(1-F(p))^2\right)dp$$

$$= \int_a^z q\left(1+b\left(1-\frac{z}{p}\right)\right) + (1-q)\left(1-b^2\left(1-\frac{z}{p}\right)\right)^2 dp$$

$$= \int_a^z 1+b\left(1-\frac{z}{p}\right)\left(q+(q-1)b\left(1-\frac{z}{p}\right)\right)dp$$

$$= \int_a^z 1+bq\left(1-\frac{z}{p}\right)\left(1-\frac{1}{2}\left(1-\frac{z}{p}\right)\right)dp \text{ (using } (q\text{-}1)b = -q/2)$$

$$= \int_a^z 1+\frac{bq}{2}\left(1-\frac{z^2}{p^2}\right)dp = [z-a]\left[1+\frac{bq}{2}\right]+\frac{bq}{2}z^2\left[\frac{1}{z}-\frac{1}{a}\right]$$

$$= z\left[1-\frac{q}{2-q}\right]\left[1+\frac{bq}{2}\right]+\frac{bq}{2}z^2\left[\frac{1}{z}-\frac{2-q}{zq}\right]. \text{ Hence}$$

$$\frac{c}{z} = \left[1-\frac{q}{2-q}\right]\left[1+\frac{bq}{2}\right]+\frac{bq}{2}\left[1-\frac{2-q}{q}\right] = \left[\frac{2(1-q)}{2-q}\right]\left[1+\frac{bq}{2}\right]-\frac{bq}{2}\left[\frac{2(1-q)}{q}\right]. \text{ Hence}$$

$$\frac{c}{z(1-q)} = \frac{2}{2-q}\left[1+\frac{bq}{2}\right]-b = \frac{1}{2-q}\left(2+bq-b(2-q)\right) = \frac{2+2bq-2b}{2-q}$$

$$= \frac{2\left(1-b(1-q)\right)}{2-q} = \frac{2\left(1-(q/2)\right)}{2-q} = \frac{2-q}{2-q} = 1. \qquad \therefore\ z = \frac{c}{1-q}.$$

To get the NSE price distribution F one simply inserts $p^* = c/(1\text{-}q)$ into equation (2). QED.

Corollary. *In NSE, buyers never search. For q Î (0,1) and c Î (0, 2.00(1-q)),*

1. *the NSE posted and transacted price densities, upper and lower endpoints, and mean and median are as listed in Table 1;*
2. *all sellers earn the same profit* $mqc/(1-q)$*, where* m *is the number of buyers per seller; and*
3. *the NSE demand for a seller posting price p is* $y = m\, s(p)$*, where the relative share is* $s(p) = q + 2(1-q)(1-F(p))$ *and F is the NSE price distribution given in Proposition 3.*

Proof. In equilibrium, buyers correctly perceive the price distribution and don't actually search since no posted price exceeds their reservation price p^*. As noted in the proof of Proposition 2, seller profit is p^*mq for all sellers; one obtains 2. by inserting the NSE reservation price $p^* = c/(1-q)$. The next to last sentence in the proof of Proposition 2 established 3. for an arbitrary distributions *F*.

Thus we need only verify the formulas in the Table. The upper and lower endpoints of the support intervals are, of course, the same as in Proposition 3. The posted price density $f(p)$ is the derivative of the distribution function $F(p)$ obtained in Proposition 3, and is zero off the support interval. The transaction price density is obtained by multiplying the posted price density $f(p)$ by the relative customer share $s(q) = q + 2(1-q)(1-F(p))$. Inserting equation (2) for $F(p)$, one has $s(q) = q + 2(1-q)(-1)(1 - p^*/p)q/(2 - 2q) = qp^*/p = [cq/(1-q)]\, p^{-1}$, and the expression for the transaction price density follows directly. The means and medians are obtained by the following tedious but direct calculations.

Mean posted price:

$$m^p = \int_a^z pf(p)dp = \int_a^z \frac{cq}{2(1-q)^2} p^{-2} pdp = \frac{cq}{2(1-q)^2}\int_a^z p^{-1}dp = \frac{cq}{2(1-q)^2}\left[\ln(z) - \ln(a)\right]$$

$$= \frac{cq}{2(1-q)^2}\ln(z/a) = \frac{cq}{2(1-q)^2}\ln(\frac{2}{q} - 1).$$

(The last step follows from substituting in $z=c/(1-q)$ and $a=zq/(2-q)$ and simplifying.)

Median posted price is the solution p^m to

$½ \equiv F(p^m) = 1+½(1-z/p^m)q/(1-q) \Leftrightarrow 1 = (z/p^m-1)q/(1-q) \Leftrightarrow z/p^m-1 = (1-q)/q$

$\Leftrightarrow z/p^m = 1/q \Rightarrow p^m = qz = qc/(1-q)$.

Mean transacted price:

$$m^T = \int_a^z pf^T(p)dp = \int_a^z \frac{c^2q^2}{2(1-q)^3}\, p^{-3}pdp = \frac{c^2q^2}{2(1-q)^3}\int_a^z p^{-2}dp = \frac{c^2q^2}{2(1-q)^3}\left[\frac{-1}{z} - \frac{-1}{a}\right]$$

$$= \frac{c^2q^2}{2(1-q)^3}\left[\frac{2(1-q)^2}{cq}\right] = \frac{qc}{1-q}$$

(The last steps substitute in $z=c/(1-q)$ and $a=zq/(2-q)$.)

Median transacted price is the solution p^M to the following, where $k = \dfrac{c^2q^2}{2(1-q)^3}$:

$$½ = F(p^M) = \int_a^{p^M} kx^{-3}dx = \frac{k}{2}\left(\frac{1}{a^2} - \frac{1}{(p^M)^2}\right) \Leftrightarrow \frac{1}{a^2} - \frac{1}{(p^M)^2} = \frac{1}{k}$$

$$\Leftrightarrow \frac{1}{(p^M)^2} = \frac{1}{a^2} - \frac{1}{k} = \left[\frac{(2-q)(1-q)}{cq}\right]^2 - \frac{2(1-q)^3}{c^2q^2} = \frac{(1-q)^2(4-4q+q^2-2+2q)}{c^2q^2}$$

$$\Leftrightarrow \left(p^M\right)^2 = \frac{c^2q^2}{(1-q)^2(2-2q+q^2)} \Leftrightarrow p^M = \frac{cq}{(1-q)\sqrt{2-2q+q^2}}\,.$$

Appendix B

INSTRUCTIONS TO TRADERS

Market with Search (DIA)

January 18, 1999

I. General

1. This is an experiment in the economics of market decision making. The National Science Foundation and other research organizations have provided funds for the conduct of this experiment. The instructions are straightforward, and if you follow them carefully you can earn a CONSIDERABLE AMOUNT OF MONEY which will be PAID TO YOU IN CASH at the end of the experiment.

2. In this experiment we create a simulated market. As a BUYER or SELLER in this market, you can use your computer to purchase or sell units of the good. Remember that the information on your computer screen is PRIVATE. To insure the best results for yourself and complete data for the experimenters, DO NOT TALK with other market participants while trade is in progress, and DO NOT DISCUSS your information with others at any point during the experiment.

3. Your computer screen will tell you whether you are a buyer or a seller and will display useful information about buying and selling opportunities.

4. Each time for buying and selling is called a TRADING PERIOD and will usually last two minutes or less. At the start of each period, sellers POST PRICES, i.e., each seller enters a price for his or her units. Then some of the prices are shown to each buyer as explained below. Each buyer decides whether or not to search for other prices and whether to buy a single unit at the posted price. Then the trading period is over.

5. At the end of the trading period, all units are "consumed" and your profits for that period are computed as explained below. The computer screen will display your profits for that period and your total profits over all periods so far. Other information on last period’s trading activity may also be displayed. Then the new trading period will begin. Everyone has new opportunities to buy or sell each period; old units do not carry over into the new period. At least 40 trading periods are scheduled in most experiments.

6. At the end of the last trading period, you will be paid in cash with your total profits converted at the rate written at the end of your instructions, plus a $5.00 participation fee. For example, if your total profits for all 50 trading periods were $24.36 and the conversion rate (written at the back of your instructions) were 0.5, then you would take home 0.5 x 24.36 + 5.00 = $17.18 in cash. All buyers have the same conversion rate and all sellers have the same conversion rate, but these two conversion rates may differ.

7. Important information will be written on the board at the beginning of the experiment. The information may include: buyer search cost, probability buyers see one vs. two prices, and the number of buyers and sellers.

II. Sources of Profit

1. Each buyer can purchase a single unit of the good each period. All buyers have a value of $2.00 for the unit. A buyer who purchases a unit at price p earns PROFIT = $2.00 - p that period. For example, a buyer would earn a profit of $2.00 - $0.94 = $1.06 if she purchases a unit at price $0.94. Note that she would earn a negative profit (lose money) if she paid a price above her value of $2.00. Buyers who don't buy a unit automatically earn a profit of zero that period.

2. Each seller can sell several units of the good each period. All sellers have zero costs in this experiment. A seller who posts price p earns PROFIT = p - 0 on every unit sold that period. For example, a seller posting a price of $0.94 would earn a per unit profit of $0.94 - 0; his profit for that period would be $0.94 if he sells 1 unit, $1.88 if he sells 2 units, etc. Sellers who don't sell any units automatically earn a profit of zero that period.

III. How to Buy or Sell

1. If you are a seller, you post your price each period using the scroll bar in the Seller window. You adjust the price, shown at the left of the scroll bar, by clicking on the arrows at either end of the scroll bar. See figure 1 below. Then click on the OK button to post the price you chose. The computer will then wait for all other sellers to post.

Figure 1

2. After all sellers have posted prices, then each buyer's computer screen initially displays a seller's posted price. A buyer can only purchase units from a seller whose price appears on her screen. The computer **randomly and independently** determines which sellers' prices are displayed on each buyer's screen **each period**. Which price(s) a buyer sees does not depend on the actions of any buyer or seller, and different buyers are likely to observe different sellers' prices. Each period every seller's price is equally likely to be shown to each buyer.

3. In some periods each buyer may be shown the price posted by a second seller, again randomly determined. The chance of seeing a second price is the same for all buyers, and is determined independently for each buyer. Therefore, in some periods some buyers will see two prices as in Figure 2a and the other buyers see only one price, as in Figure 2b.

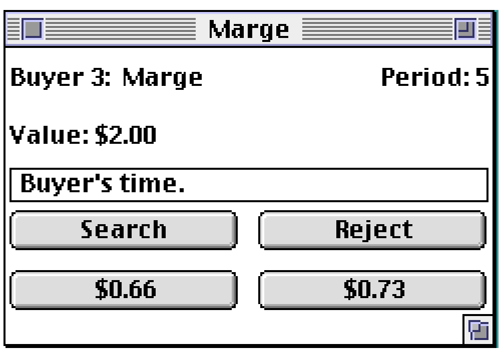

Figure 2a

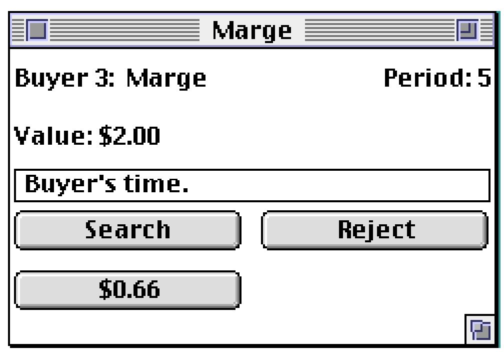

Figure 2b

4. The probability that any buyer observes only one price can take on four possible values. The experimenter will announce the probability that buyers observe only one price in each period and it is also displayed on the screen in the Experiment Description window shown in Figure 3.

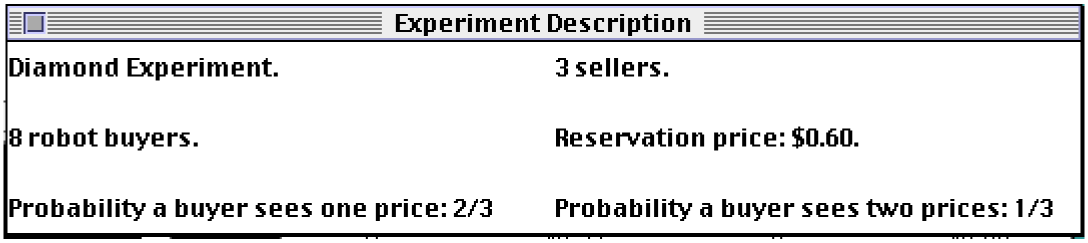

Figure 3

The table below describes the four possible cases:

Likelihood of One Price vs. Two
All buyers see two prices with certainty. (Probability of seeing one price is 0).
All buyers see one price with probability 1/3, two prices with probability 2/3.
All buyers see one price with probability 2/3, two prices with probability 1/3.
All buyers see one price with certainty. (Probability of seeing one price is 1).

5. Buyers who want to see other posted prices can always pay a fee and search. The fee is announced at the beginning of the experiment, and must be paid on every search. For example, a buyer who searches twice when the fee is $0.25 will have $0.50 deducted from her profits that period. When a buyer searches, the prices they had seen go back into the pool of seller prices and a new sample of prices is drawn using the exact same rules the original sample was drawn. Neither the number of prices nor the identity of the seller(s) from whom the prices sampled in a search come have any relationship to earlier prices observed. All samples are independent.

6. To make a purchase, the buyer clicks on the price at which she wants to buy, see figure 4. If a buyer chooses to search and draw a another sample she clicks on the **Search** button. A buyer who chooses not to buy this period clicks on **Reject** button and is finished for that period, and will earn zero profit.

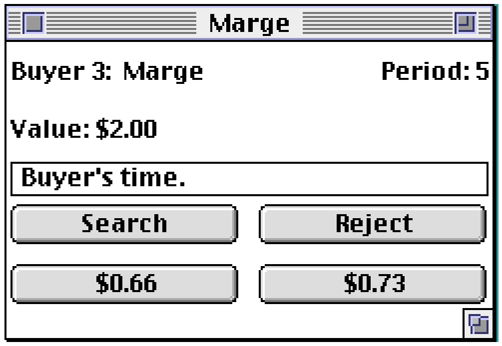

Figure 4

IV. At the End of the Period

1. When all buyers have finished, every trader's screen will summarize the period's activities as follows. A History window will appear summarizing your actions in the previous five periods as well as your Total Profit for all periods. Figure 5 below shows a buyer's history window, seller windows are identical except they read "Shares Sold" and there is no search cost column.

Marge History

Period	Price	Shares Bought	Search Cost	Profit
1	$0.18	1	$0.00	$1.82
2	$0.58	1	$0.60	$0.82
3	$0.71	1	$0.00	$1.29
4	$0.52	1	$0.00	$1.48
5	$0.66	1	$0.00	$1.34
				Total Profit: $6.75

Figure 5

2. In addition, your own profit calculations for the period are shown in detail in the Interim window on your screen. Figure 6a shows an example of a seller's interim screen and 6b below shows that of a buyer. For a seller each line, labeled Sale, represents one unit sold in the period. Cost for a seller reflects production cost (in this case zero). For a buyer, each line shows a search or a purchase/not purchase decision. Cost for a buyer is either search cost (as in the first row) or the price of the good (shown in the second row). The Net Profit column lists your profit on each transaction, which is totaled at the bottom.

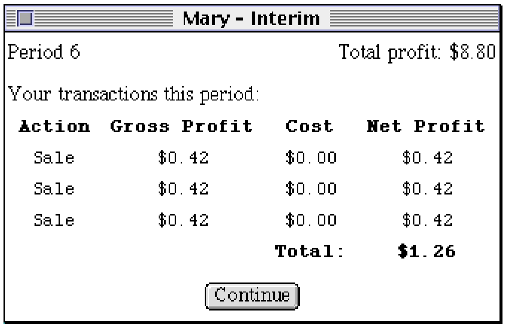

Mary - Interim

Period 6 Total profit: $8.80

Your transactions this period:

Action	Gross Profit	Cost	Net Profit
Sale	$0.42	$0.00	$0.42
Sale	$0.42	$0.00	$0.42
Sale	$0.42	$0.00	$0.42
		Total:	**$1.26**

Figure 6a (seller)

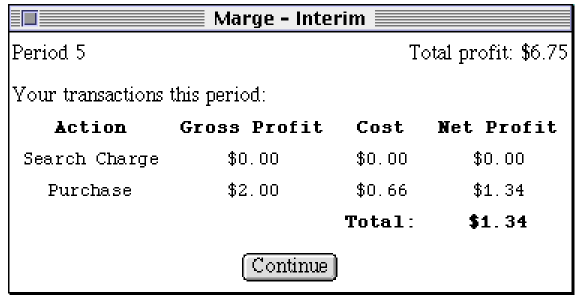

Marge - Interim

Period 5 Total profit: $6.75

Your transactions this period:

Action	Gross Profit	Cost	Net Profit
Search Charge	$0.00	$0.00	$0.00
Purchase	$2.00	$0.66	$1.34
		Total:	**$1.34**

Figure 6b (buyer)

3. A third window may appear which summarizes all market activity in the period, see Figure 7. (Whether or not this window appears in today's experiment will be announced.)

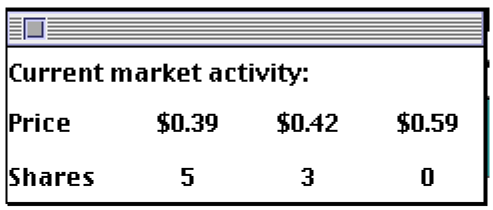

Current market activity:

Price	**$0.39**	**$0.42**	**$0.59**
Shares	**5**	**3**	**0**

Figure 7

In this example there were three sellers. One seller sold 5 units at $0.39 each, another sold 3 units at $0.42 each, and the third sold 0 units at $0.59 each.

4. When you are finished viewing this information, click on the "CONTINUE" button in the Interim window. Once all traders have continued or when time expires another period begins as explained in III. 1. above.

V. Frequently Asked Questions

1. Q: Is this some psychology experiment with an agenda you haven’t told us?

   A: No. It is an economics experiment. If we do anything deceptive, or don’t pay you cash as described, then you can complain to the campus Human Subjects Committee and we will be in serious trouble. These instructions are on the level and our interest is in seeing how people make decisions in market situations.

2. Q: The price scroll bar seems to keep scrolling. Is it stuck?

   A: No, but if you **double click** on the arrow at either end of the scroll bar, the price will continue to scroll as long as your pointer is positioned over that arrow, **even after you release the mouse button. Be very sure that the price has stopped on the price you want to post before clicking the OK button.**

3. Q: If a buyer searches but decides not to buy in that period, does she still pay the search cost(s)?

A: Yes, the search cost is deducted from a buyer's profit whether she ends up buying or not. This could mean a negative profit in a period where she searches but does not buy.

4. Q: Does everyone face the same cost of searching?

   A: Yes, search costs are identical for all buyers and announced by the experimenter and are posted on the blackboard.

5. Q: When a buyer searches, how many prices does she see?

   A: This depends on the probability, call it q, (which is set by the experimenter, announced to all players and posted on the board and her screen) of a buyer seeing one price vs. two. A buyer search reveals one price with probability q or two with probability (1-q), as explained in section III. 3. above. For example, if q = 1/3, a buyer who searches has a 1 in 3 chance of observing a single price and a 2 in 3 chance of observing two prices. If the buyer sees two prices, each will always be from different sellers.

6. Q: If I saw two prices one time does that mean I will see two again if I search?

   A: Not necessarily. The number of prices observed in a search is independent of the number observed from earlier observations or searches. A buyer who sees one price initially has the same chance as one who initially sees two of observing a single price again when he searches.

7. Q: If I search, will I see prices different from those I've already seen this period? Can I still buy at a price I saw before I searched?

   A: Not necessarily. Once a buyer decides to search, the price(s) he saw initially or in a previous search go back into the pool from which the new price observations are randomly drawn. It **is** possible for a buyer to observe the same seller's price again in a search, though there is no way of knowing whether it is the same seller's price or another seller posting an identical price. You can only buy at a price that is currently shown on your screen.

Table 1:
Noisy Search Equilibrium Price Distributions

Prices	density (p)	Support endpoints lower (a)	upper (z)	mean	median
Posted	$\frac{cq}{2(1-q)^2} p^{-2}$	$\frac{cq}{(1-q)(2-q)}$	$\frac{c}{1-q}$	$\frac{cq}{2(1-q)^2}\ln\left(\frac{2}{q}-1\right)$	$\frac{cq}{1-q}$
Trans-acted	$\frac{c^2q^2}{2(1-q)^3} p^{-3}$	$\frac{cq}{(1-q)(2-q)}$	$\frac{c}{1-q}$	$\frac{cq}{1-q}$	$\frac{cq}{(1-q)\sqrt{2-2q+q^2}}$

Note: for $q\in(0, 1)$ and $c\in(0, 2\text{-}2q)$

Table 2:
Testable Predictions

- Buyers never search.
- Prices are unified for q=0 and q=1, but are dispersed for q=1/3 and q=2/3.
- With more robots, learning should be faster and less biased and therefore the NSE predictions should be more accurate.
- Noisy Search Equilibrium (NSE) prices in dollars:

	q=0*	q=1/3	q=2/3	q=1
[Lower bound, Reservation Price] = [a, z]				
c=0.20	0.01, 0.01	0.06, 0.30	0.30, 0.60	2.00, 2.00
c=0.60	0.01, 0.01	0.18, 0.90	0.90, 1.80	2.00, 2.00
Mean transaction = Median posted price				
c=0.20	0.01	0.10	0.40	2.00
c=0.60	0.01	0.30	1.20	2.00

*In the experiment, the lowest price sellers can post is one cent, not zero.

Table 3:
Summary of Laboratory Sessions

Session Name	Search Cost	Number of Sellers	Number of Buyers	Human Buyers?	Number of Periods	Experienced?	*q*-sequence
UC5	20 cents	6	6	Yes	100	No	1, 0, 2/3, 1/3
PU5	20 cents	6	6	Yes	81	No	0, 1, 1/3, 2/3
PU8x	20 cents	6	6	Yes	131	Yes	1, 0, 2/3, 1/3
PU7	60 cents	6	6	Yes	98	No	1, 0, 2/3, 1/3
UC8	60 cents	6	6	Yes	100	No	0, 1, 1/3, 2/3
UC7x	60 cents	5	5	Yes	100	Yes	0, 1, 1/3, 2/3
PU4	20 cents	6	6	No	100	No	1, 0, 2/3, 1/3
UC11	20 cents	6	6	No	140	No	1, 0, 2/3, 1/3
PU3	60 cents	6	6	No	100	No	1, 0, 2/3, 1/3
UC10x	60 cents	6	6	No	140	Yes	0, 1, 1/3, 2/3
UC2	20 cents	6	12	No	80	No	1, 0, 2/3, 1/3
UC1	20 cents	5	10	No	80	No	0, 1, 1/3, 2/3
PU2	60 cents	6	12	No	100	No	0, 1, 1/3, 2/3
PU6	60 cents	6	12	No	100	No	1, 0, 2/3, 1/3
UC9x	20 cents	6	600	No	140	Yes	1, 0, 1/3, 2/3, 1/3, 2/3
PU9	20 cents	6	600	No	100	No	0, 1, 1/3, 2/3
UC6	60 cents	6	45	No	100	No	1, 0, 2/3, 1/3
UC12x	60 cents	6	600	No	150	Yes	0, 1, 1/3, 2/3, 1/3, 2/3

Table 4:

Trading Efficiency in Human Buyer Sessions

q treatment	Experience Level	Search Cost = 20 cents		Search Cost =60 cents	
		Net	Gross	Net	Gross
0	Inexperienced	98.9	99.6	99.5	100
0	Experienced	98.3	98.3	100	100
1/3	Inexperienced	99.2	100	99.2	99.7
1/3	Experienced	99.5	99.6	99.6	100
2/3	Inexperienced	97.8	99.2	97.2	98.0
2/3	Experienced	97.2	97.8	100	100
1	Inexperienced	92.0	95.3	92.5	95.0
1	Experienced	92.7	95.0	95.1	98.0

Notes: Net trading efficiency is gross efficiency (the actual gains from trade) minus the incurred search costs. Efficiencies are expressed as a percentage of maximum possible gains from trade.

Table 5:
Mean Seller Profit

Panel A: *q*=1/3 and 2/3

	Search Cost			
	20 cents		60 cents	
Buyer Population	*q*=1/3	*q*=2/3	*q*=1/3	*q*=2/3
NSE Prediction	**10**	**40**	**30**	**120**
6 Human Buyers	42.4 (2.2)	79.7 (3.8)	75.5 (3.6)	100.9 (4.6)
6 Robot Buyers	18.1 (0.9)	52.4 (2.5)	42.2 (2.4)	119.9 (6.2)
12 Robot Buyers	13.9 (0.7)	44.8 (1.9)	40.4 (1.7)	124.1 (4.4)
Many Robot Buyers	16.1 (0.3)	44.7 (0.4)	50.3 (1.1)	141.8 (1.8)
All Buyers	24.8 (0.8)	56.6 (1.4)	53.6 (1.3)	121.4 (2.3)

Panel B: *q*=0 and 1

	Search Cost			
	20 cents		60 cents	
Buyer Population	*q*=0	*q*=1	*q*=0	*q*=1
NSE Prediction	**1**	**200**	**1**	**200**
6 Human Buyers	50.9[a] (3.5)	110.5 (5.4)	56.4[a] (4.0)	138.2 (7.1)
6 Robot Buyers	3.7 (0.3)	185.7 (9.3)	3.3 (0.3)	168.6 (9.2)
12 Robot Buyers	3.3 (0.2)	188.9 (7.4)	3.1 (0.2)	160.0 (7.6)
Many Robot Buyers	1.9 (0.1)	199.0 (0.3)	3.0 (0.2)	194.0 (3.1)
All Buyers	20.0[b] (1.4)	159.8 (3.5)	18.5[b] (1.4)	162.9 (3.8)

Notes: *q* is the probability that buyers observe one price instead of two prices. Entries are mean seller profit per buyer in cents. Standard errors are shown in parentheses.
[a]Median profit for 6 human buyers for *q*=0 is 5.0 when search cost is 20 cents and is 34.0 when search cost is 60 cents.
[b]Median profit across buyer number/type treatments for *q*=0 is 2.4 when search cost is 20 cents and is 2.0 when search cost is 60 cents.

Table 6
Comparisons of Mean Transaction Prices

q	9 sessions with 20-cent search cost Mean ± Std. Error	9 sessions with 60-cent search cost Mean ± Std. Error	Mean Paired Difference for 60 cent – 20 cent *Wilcoxon p- Value*
0	17.4 ± 9.5	21.2 ± 9.3	**NSE = 0** 3.7 ± 4.9 *p=0.73*
1/3	25.1 ± 4.9	54.0 ± 6.6	**NSE = 20** 28.9 ± 3.0 *p<0.01*
Mean Paired Difference for (q=1/3) – (q=0) *Wilcoxon p- Value*	**NSE = 9** 7.7 ± 7.7 *p=0.16*	**NSE = 29** 32.8 ± 6.0 *p<0.01*	
2/3	60.1 ± 7.6	118.6 ± 7.5	**NSE = 80** 58.5 ± 12.0 *p<0.01*
Mean Paired Difference for (q=2/3) – (q=1/3) *Wilcoxon p- Value*	**NSE = 30** 35.0 ± 3.6 *p<0.01*	**NSE = 90** 64.6 ± 10.2 *p<0.01*	
1	164.3 ± 14.5	163.7 ± 8.1	**NSE = 0** -0.7 ± 12.1 *p=0.55*
Mean Paired Difference for (q=1) – (q=2/3) *Wilcoxon p- Value*	**NSE = 160** 104.2 ± 20.3 *p<0.01*	**NSE = 80** 45.1 ± 7.6 *p<0.01*	

Notes: NSE denotes noisy search equilibrium predictions in cents, shown in **bold**. Mean transaction prices and price differences ± standard errors are shown to the nearest tenth of a cent. Wilcoxon p-values, shown in *italics*, are for the test that the paired difference is significantly different from zero.

Table 7:
Mean, Median and Skewness of Posted Prices for the Final 10 Periods of Each Treatment Run

	Mean Posted Price	Median Posted Price	Skewness of Posted Price
q=1/3, Search Cost=20 Cents			
NSE Prediction	**12.1**	**10**	**1.191**
6 Human Buyers	35.8	30	2.363
6 Robot Buyers	17.4	15	0.604
12 Robot Buyers	13.0	12.5	0.929
Many Robot Buyers	18.3	18	-0.030
q=2/3, Search Cost=20 Cents			
NSE Prediction	**41.6**	**40**	**0.486**
6 Human Buyers	84.7	78.5	0.687
6 Robot Buyers	51.6	54	-0.924
12 Robot Buyers	44.2	42	0.481
Many Robot Buyers	47.0	45.5	-0.007
q=1/3, Search Cost=60 Cents			
NSE Prediction	**36.2**	**30**	**1.191**
6 Human Buyers	72.3	70	0.969
6 Robot Buyers	40.8	34.5	0.893
12 Robot Buyers	44.3	41.5	0.583
Many Robot Buyers	54.9	52	0.330
q=2/3, Search Cost=60 Cents			
NSE Prediction	**124.8**	**120**	**0.486**
6 Human Buyers	98.3	99.5	-0.310
6 Robot Buyers	106.0	100.5	0.694
12 Robot Buyers	120.5	120	0.184
Many Robot Buyers	142.2	139.5	-0.131

Table 8:
Standard Deviation of Posted Prices over the Final 10 Periods of Each Run

Panel A: *q*=1/3 and 2/3

	Search Cost			
	20 cents		60 cents	
Buyer Population	*q*=1/3	*q*=2/3	*q*=1/3	*q*=2/3
NSE Prediction	**5.9**	**8.4**	**17.6**	**25.2**
6 Human Buyers	9.0	11.4	8.0	16.0
6 Robot Buyers	5.9	3.9	19.7	23.0
12 Robot Buyers	3.6	7.7	18.6	19.1
Many Robot Buyers	4.6	7.0	14.8	26.4

Panel B: *q*=0 and 1

	Search Cost			
	20 cents		60 cents	
Buyer Population	*q*=0	*q*=1	*q*=0	*q*=1
NSE Prediction	**0**	**0**	**0**	**0**
6 Human Buyers	12.0	7.5	14.0	8.9
6 Robot Buyers	1.3	14.8	3.1	17.1
12 Robot Buyers	1.2	3.8	0.7	8.7
Many Robot Buyers	1.1	0	0.9	2.1

Notes: Entries are the median, over the final 10 periods of each run in the given treatment, of price dispersion measured as the standard deviation of posted prices in cents in a given period.

Table 9:
Human Buyer Reservation Price Estimates (in cents)

	q=1/3 Reservation Price (violations)	*q*=2/3 Reservation Price (violations)
Search Cost=20 Cents **NSE Reservation Price**	**30**	**60**
Session UC5 (inexperienced)		
Buyer 1	51 (none)	84 (none)
Buyer 2	51 (none)	73 (2 A, 2 B)
Buyer 3	110 (none)	79 (none)
Buyer 4	69 (none)	75 (1 A, 1 B)
Buyer 5	53 (none)	85 (none)
Buyer 6	51 (none)	75 (none)
Session Mean (Median)	64 (52)	79 (77)
Session PU5 (inexperienced)		
Buyer 1	70 (2 A, 4 B)	111 (3 A, 3 B)
Buyer 2	94 (none)	116 (1 A, 1 B)
Buyer 3	141 (none)	120 (none)
Buyer 4	109 (1 A, 1 B)	125 (none)
Buyer 5	120 (none)	123 (none)
Buyer 6	123 (1 A, 1 B)	119 (1 A, 2 B)
Session Mean (Median)	110 (115)	119 (120)
Session PU8x (experienced)		
Buyer 1	54 (1 A, 0 B)	110 (0 A, 1 B)
Buyer 2	51 (none)	100 (1 A, 0 B)
Buyer 3	54 (none)	104 (2 A, 2 B)
Buyer 4	83 (none)	103 (1 A, 1 B)
Buyer 5	51 (1 A, 1 B)	104 (1 A, 1 B)
Buyer 6	53 (none)	109 (1 A, 1 B)
Session Mean (Median)	58 (54)	105 (104)
Search Cost=60 Cents **NSE Reservation Price**	**90**	**180**
Session PU7 (inexperienced)		
Buyer 1	107 (none)	120 (none)
Buyer 2	135 (none)	118 (none)
Buyer 3	104 (none)	120 (none)
Buyer 4	127 (none)	115 (none)
Buyer 5	85 (none)	110 (1 A, 1 B)
Buyer 6	101 (none)	125 (none)
Session Mean (Median)	110 (106)	118 (119)
Session UC8 (inexperienced)		
Buyer 1	125 (none)	175 (none)
Buyer 2	123 (1 A, 1 B)	135 (1 A, 1 B)
Buyer 3	117 (1 A, 0 B)	167 (none)
Buyer 4	130 (none)	179 (none)
Buyer 5	135 (none)	143 (1 A, 1 B)
Buyer 6	130 (none)	137 (none)
Session Mean (Median)	127 (128)	156 (155)
Session UC7x (experienced)		
Buyer 1	111 (none)	133 (none)
Buyer 2	150 (none)	124 (none)
Buyer 3	100 (1 A, 1 B)	144 (none)
Buyer 4	125 (none)	132 (none)
Buyer 5	125 (none)	144 (none)
Session Mean (Median)	122 (125)	135 (133)

Table 10:
Price Autocorrelation within Sellers

Rejection Rates of Mixed Strategy NSE—95% Confidence Level

Search Cost

Buyer Population	20 cents		60 cents		Row Totals
	q=1/3	*q*=2/3	*q*=1/3	*q*=2/3	
95th percentile of simulation	0.241	0.238	0.265	0.219	
6 Human Buyers	17/18	16/18	16/17	13/17	62/70
6 Robot Buyers	9/12	4/12	11/12	9/12	33/48
12 Robot Buyers	10/11	2/11	10/12	7/12	29/46
Many Robot Buyers	11/12	9/12	12/12	8/12	40/48
Column Totals	47/53	31/53	49/53	37/53	164/212

q is the probability that buyers observe one price instead of two prices.

Table 11:
Price Correlation Across Sellers, within a Period

Rejection Rates of Mixed Strategy NSE—95% Confidence Level

Search Cost

Buyer Population	20 cents		60 cents		Row Totals
	q=1/3	*q*=2/3	*q*=1/3	*q*=2/3	
95th percentile of simulation	0.289	0.286	0.337	0.300	
6 Human Buyers	17/18	13/18	13/17	13/17	56/70
6 Robot Buyers	9/12	3/12	8/12	8/12	28/48
12 Robot Buyers	10/11	8/11	9/12	10/12	37/46
Many Robot Buyers	11/12	10/12	8/12	9/12	38/48
Column Totals	47/53	34/53	38/53	40/53	159/212

q is the probability that buyers observe one price instead of two prices.

Table 12:
Frequency of Posted Price Increases and Price Decreases by (Normalized) Number Sold in the Previous Period

Panel A: *q*=0

	Units Sold Last Period (Divided by Number of Buyers Per Seller)						
Relative to last period, price this period…	0	1	2	3	4	5 & 6	Total
Decreases	549	225	65	29	12	0	880
Increases	138	201	113	36	4	2	494
Stays Unchanged	267	378	166	49	13	7	880

Panel B: *q*=1/3

	Units Sold Last Period (Divided by Number of Buyers Per Seller)						
Relative to last period, price this period…	0	1	2	3	4	5 & 6	Total
Decreases	827	503	127	29	15	0	1501
Increases	231	429	156	55	8	1	880
Stays Unchanged	331	377	167	40	15	1	931

Panel C: *q*=2/3

	Units Sold Last Period (Divided by Number of Buyers Per Seller)						
Relative to last period, price this period…	0	1	2	3	4	5 & 6	Total
Decreases	677	439	113	29	9	2	1269
Increases	319	493	115	32	3	0	962
Stays Unchanged	309	424	160	51	11	0	955

Panel D: *q*=1

	Units Sold Last Period (Divided by Number of Buyers Per Seller)						
Relative to last period, price this period…	0	1	2	3	4	5 & 6	Total
Decreases	262	226	84	23	4	0	599
Increases	258	368	109	20	4	1	760
Stays Unchanged	181	563	117	30	4	0	895

Table 13:
Noisy Logit Best Response Model Estimates for *l* (*r*=0.75)

	Search Cost			
	20 cents		60 cents	
Buyer Population	q=1/3	q=2/3	q=1/3	q=2/3
6 Human Buyers	0.092 (0.004)	0.054 (0.006)	0.031 (0.002)	0.017 (0.003)
6 Robot Buyers	0.178 (0.011)	0.099 (0.006)	0.065 (0.005)	0.037 (0.003)
12 Robot Buyers	0.344 (0.026)	0.206 (0.017)	0.084 (0.006)	0.050 (0.004)
Many Robot Buyers	0.204 (0.015)	0.167 (0.011)	0.076 (0.005)	0.062 (0.004)
Pooled Over Buyer Number/Types	0.126 (0.004)	0.102 (0.004)	0.052 (0.002)	0.039 (0.002)

q is the probability that buyers observe one price instead of two prices. Standard Errors are shown in parentheses. All estimates are different from zero at the 99% significance level.

Figure 1: Posted Price Equilibrium Density Functions

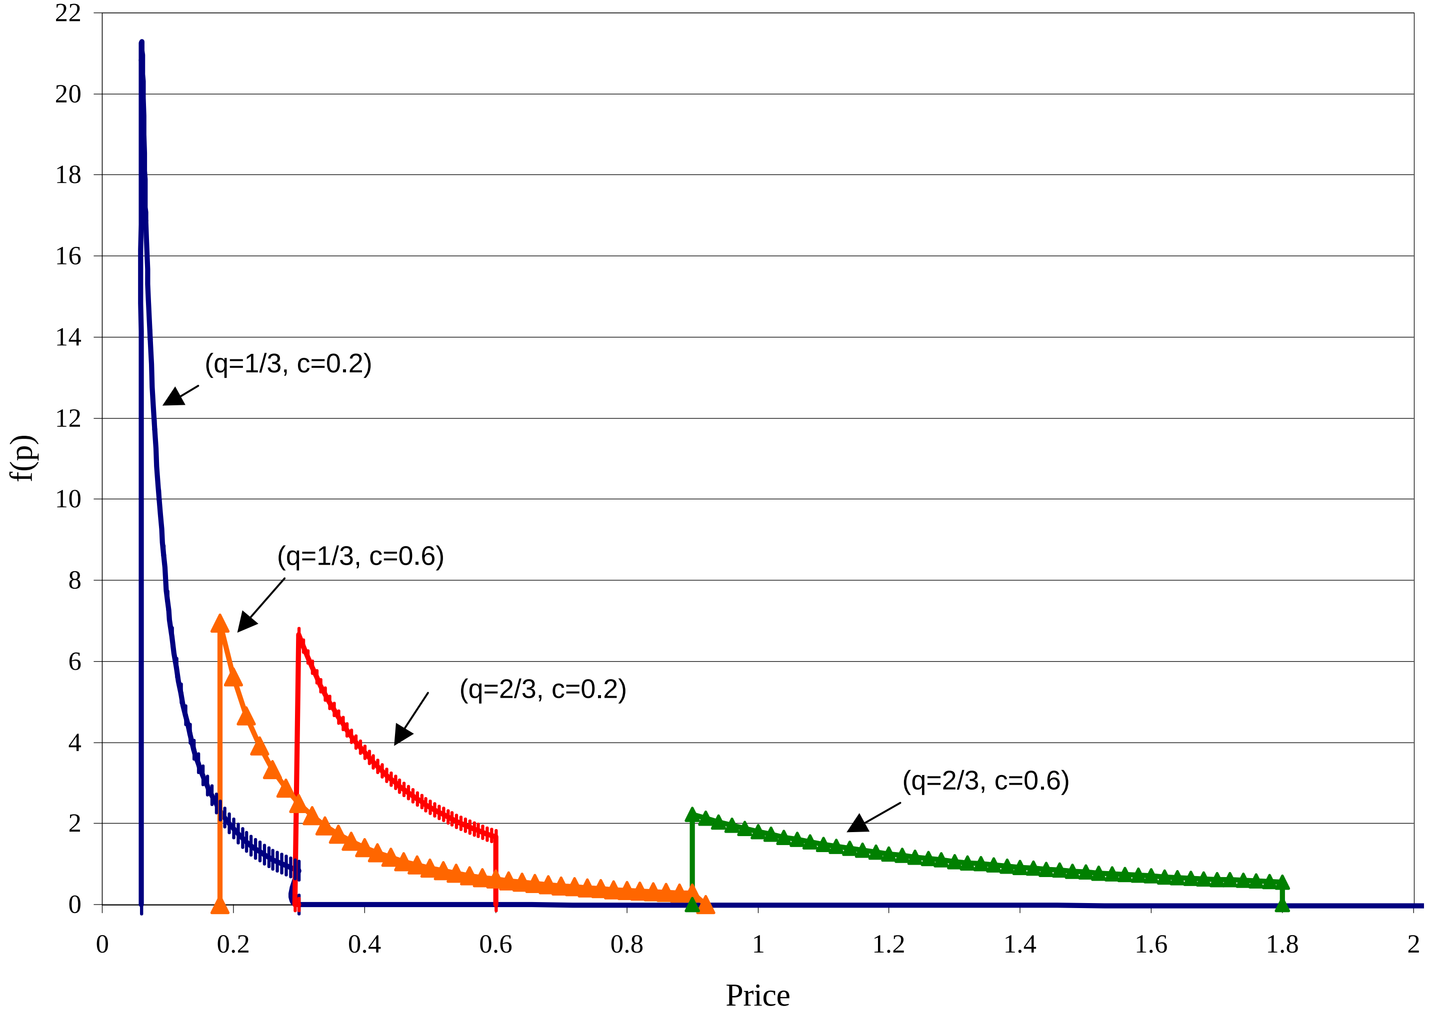

Figure 2A: Seller Window for Posting Prices

Figure 2B: Buyer Window for Purchase, Search or Reject Decision

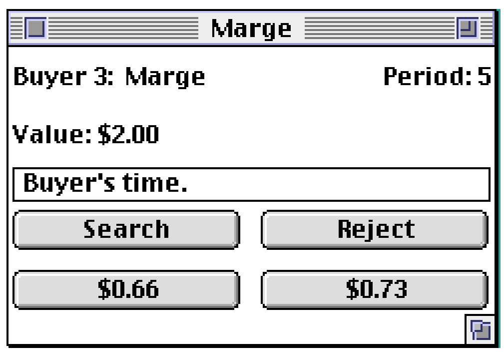

Figure 2C: Interim Screens with Profit Summary and All Posted Prices

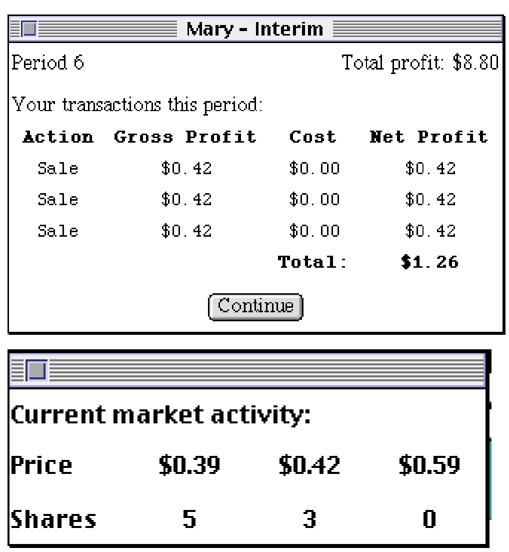

Figure 3: Prices in Session UC9x

(Experienced, Many Robot Buyers, c= 20 cents , q=2/3)

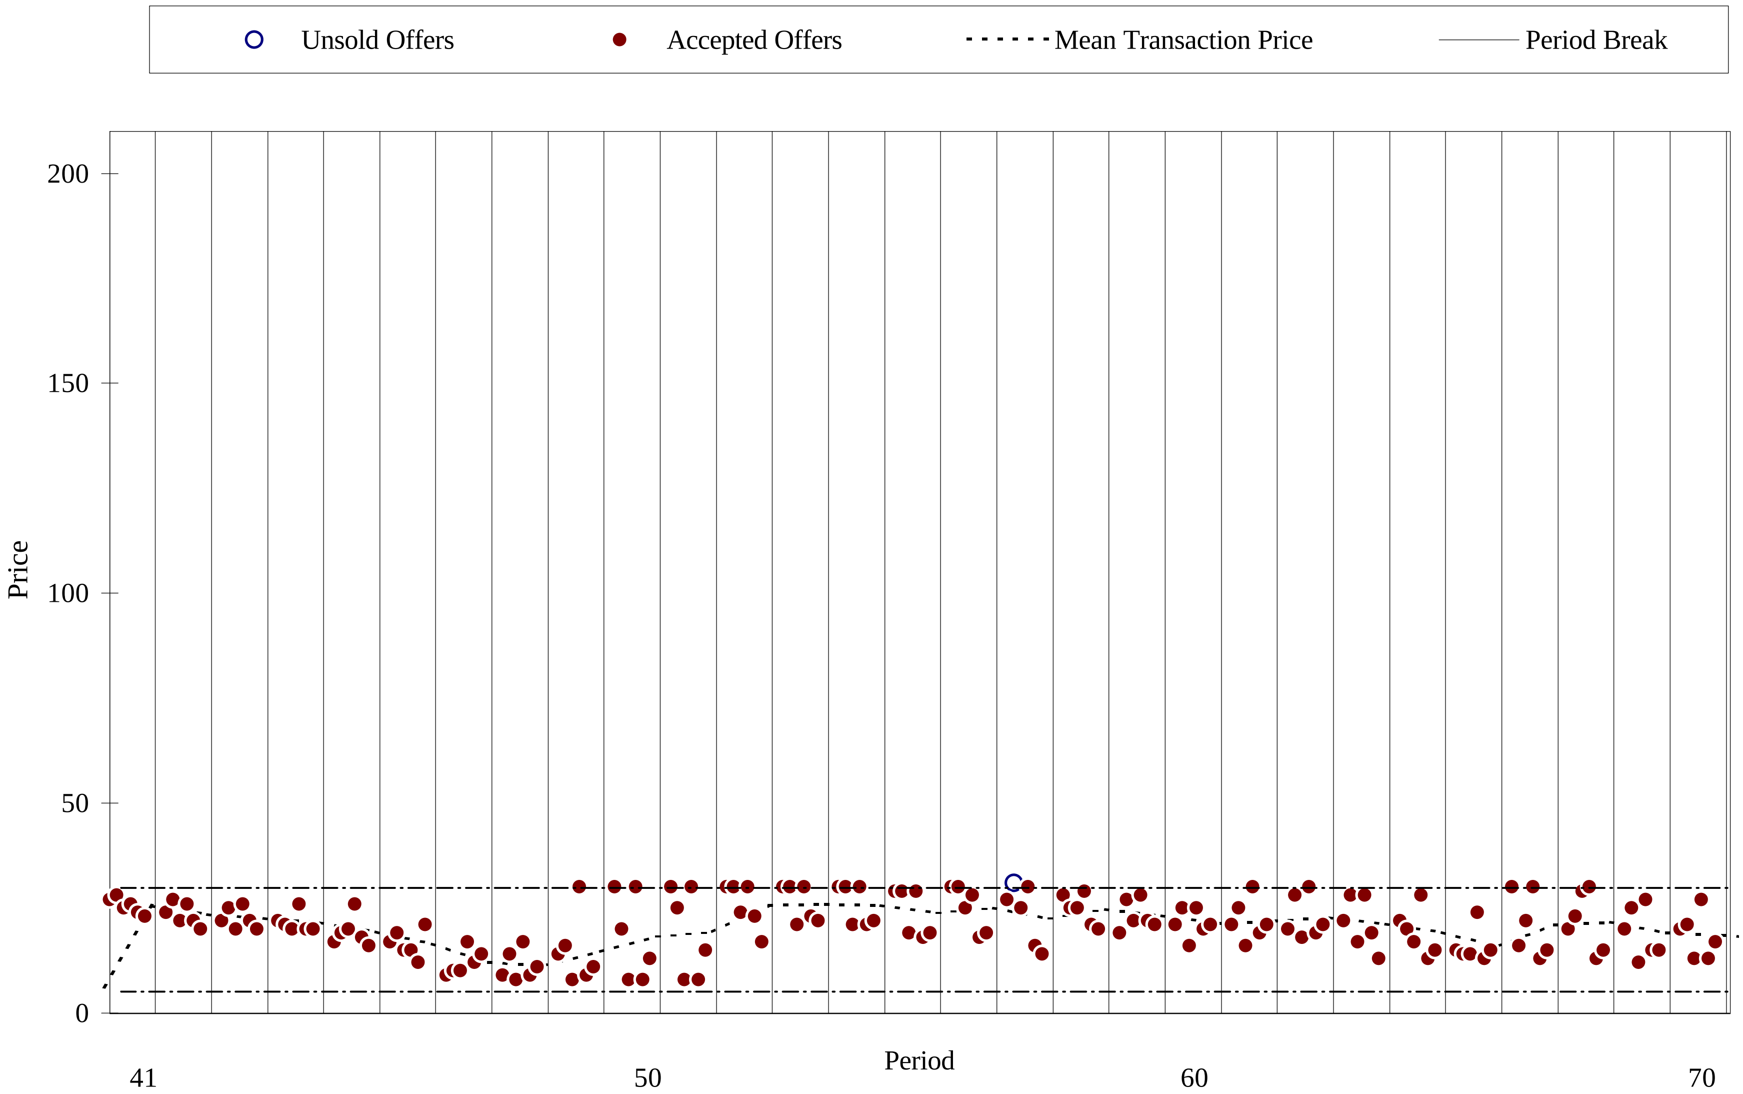

Figure 4: Prices in Session PU2
(Inexperienced, 12 Robot Buyers, c= 60 cents , q=2/3)

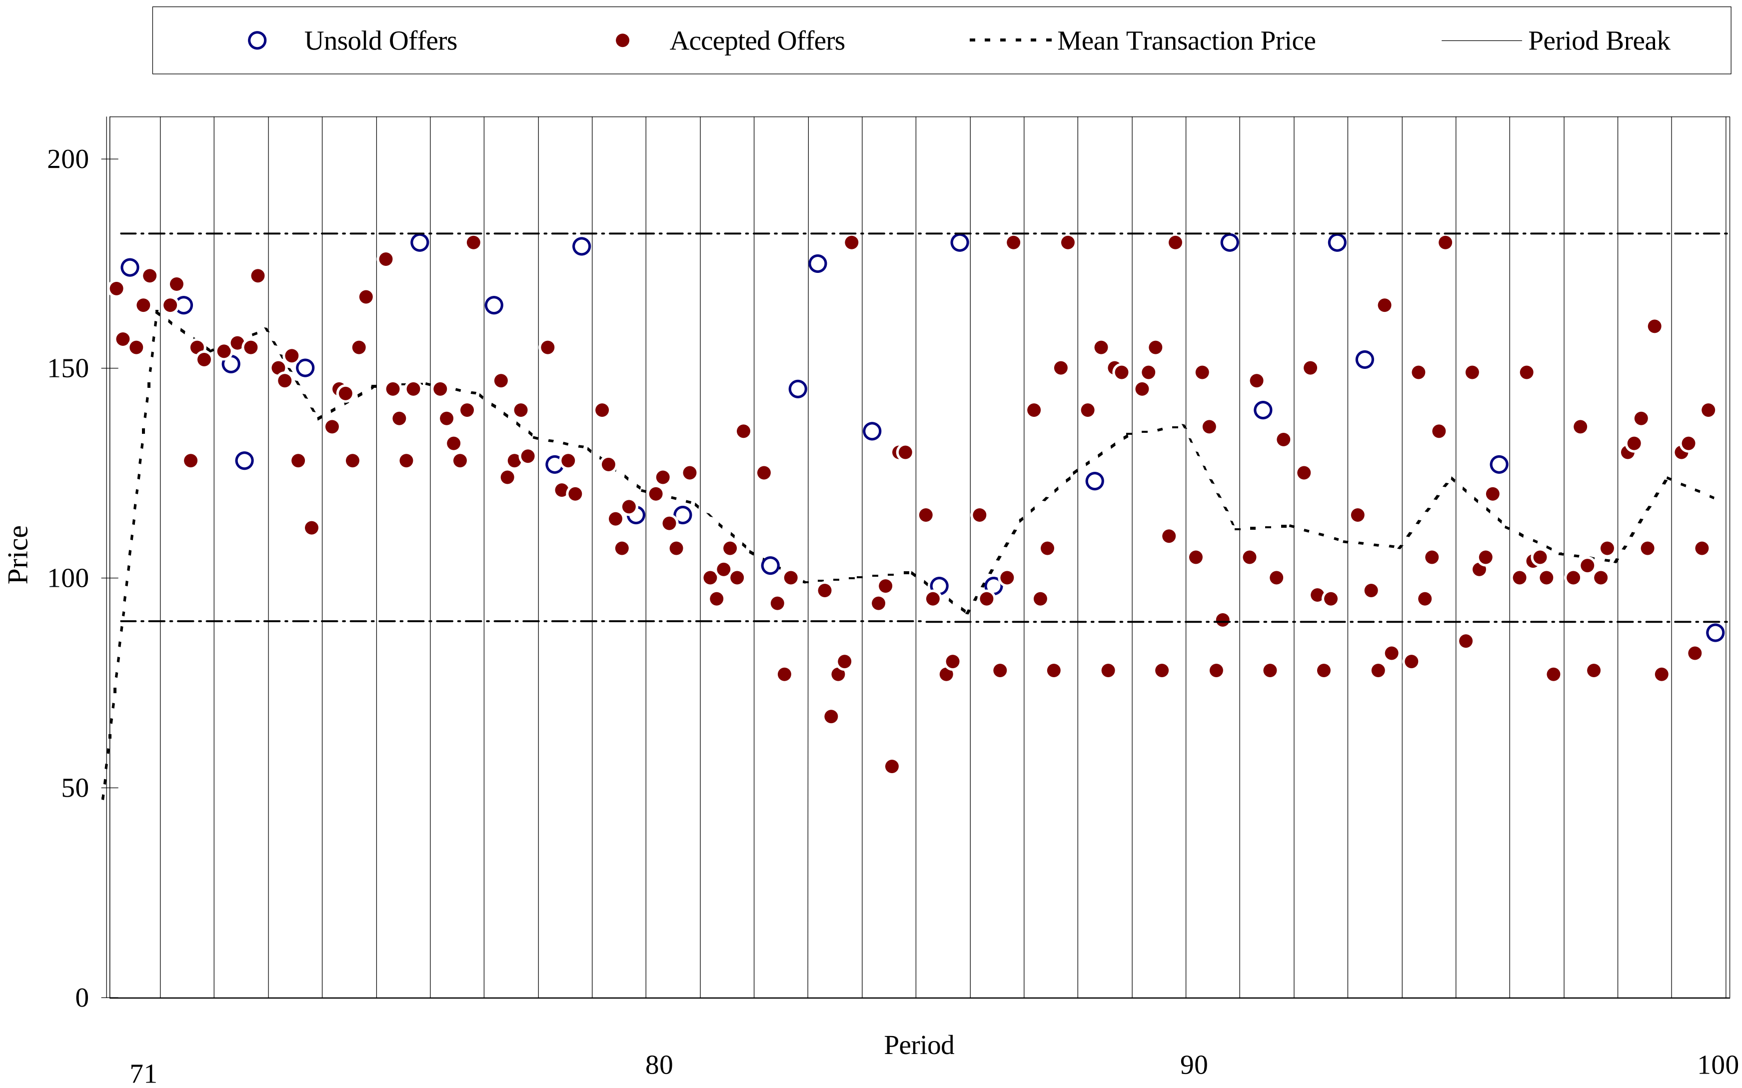

Figure 5: Prices inSession PU7

(Inexperienced 6 Human Buyers, c= 60 cents , q=1/3)

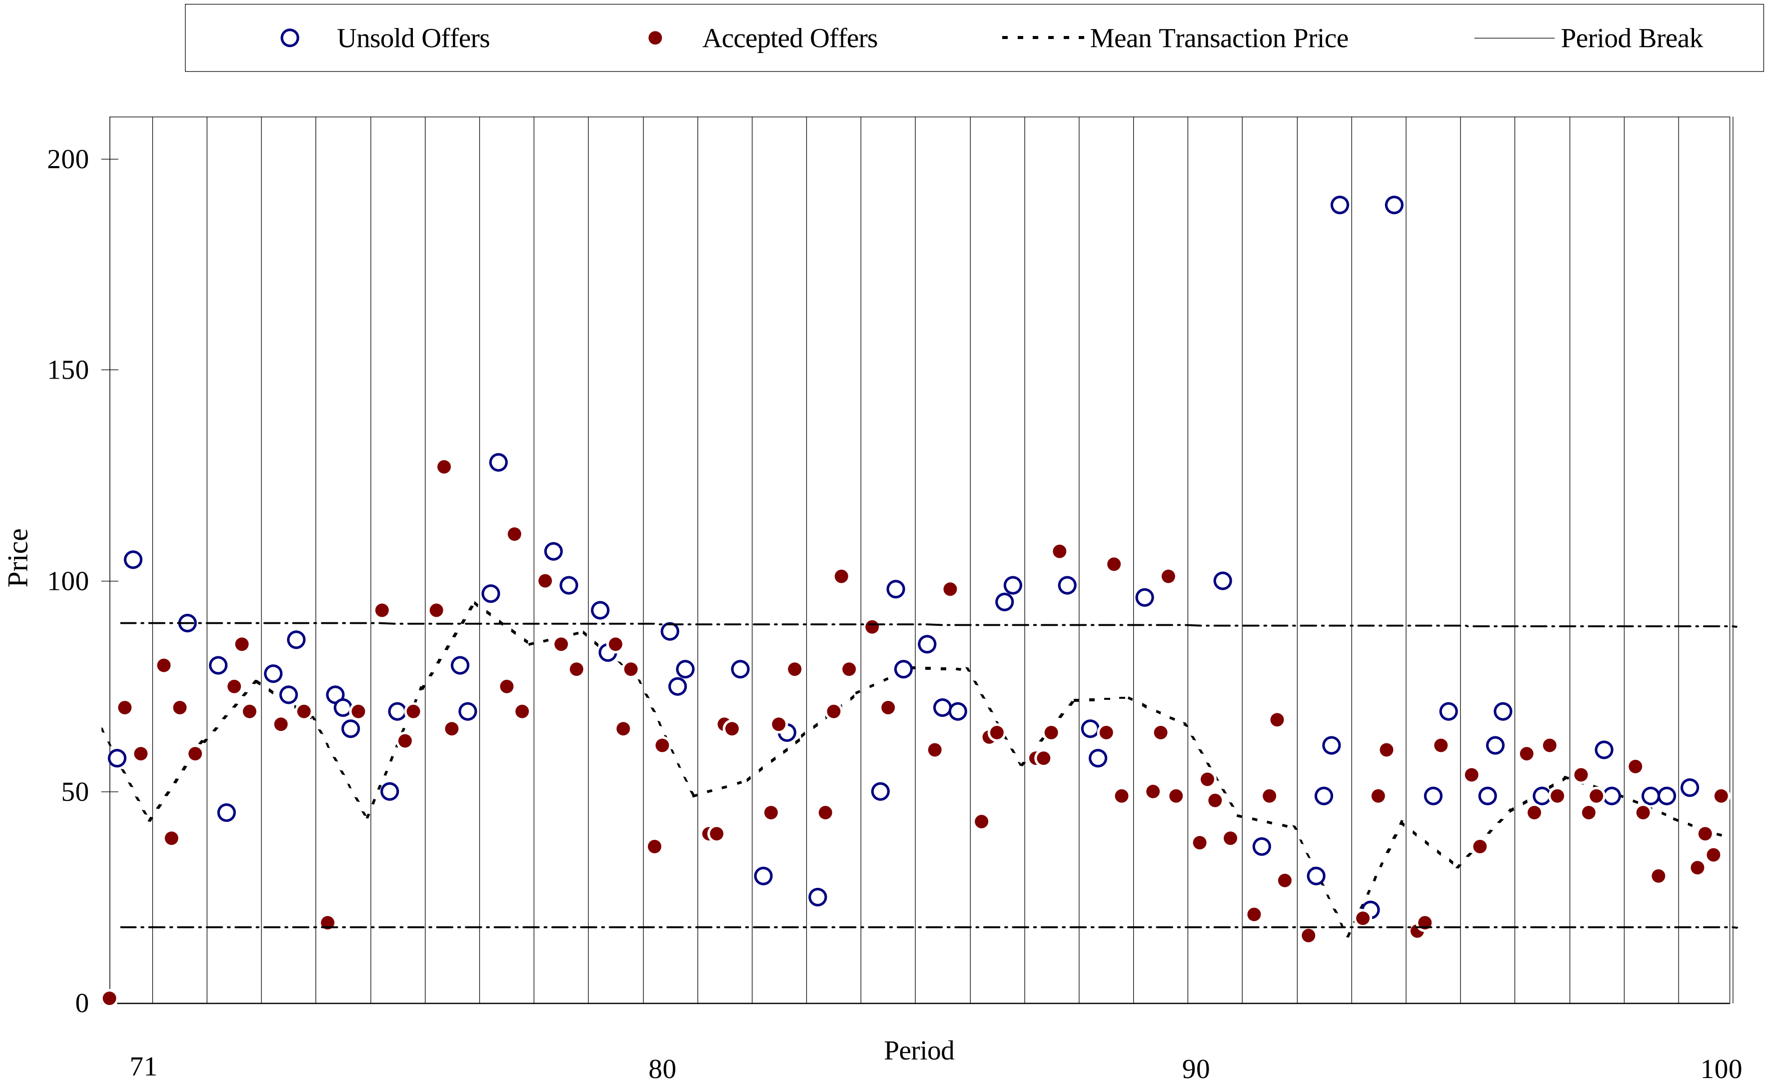

Figure 6 - Mean Transaction Price

Panel A: 6 human buyer sessions

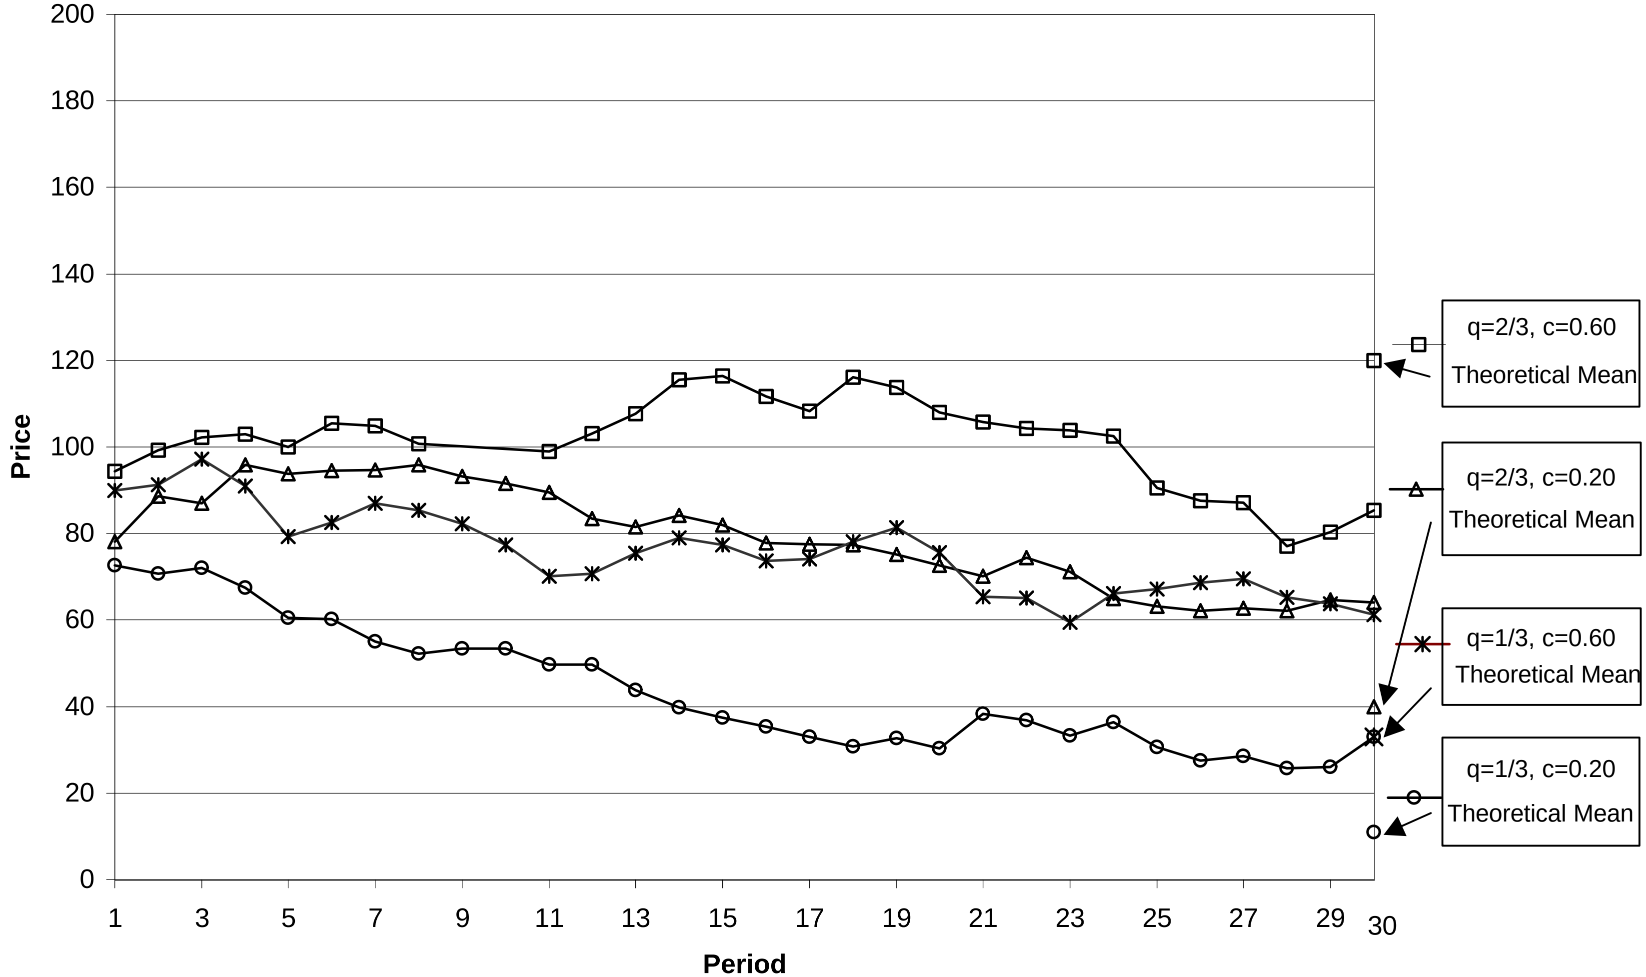

Figure 6 - Mean Transaction Price

Panel B: Many robot buyer sessions

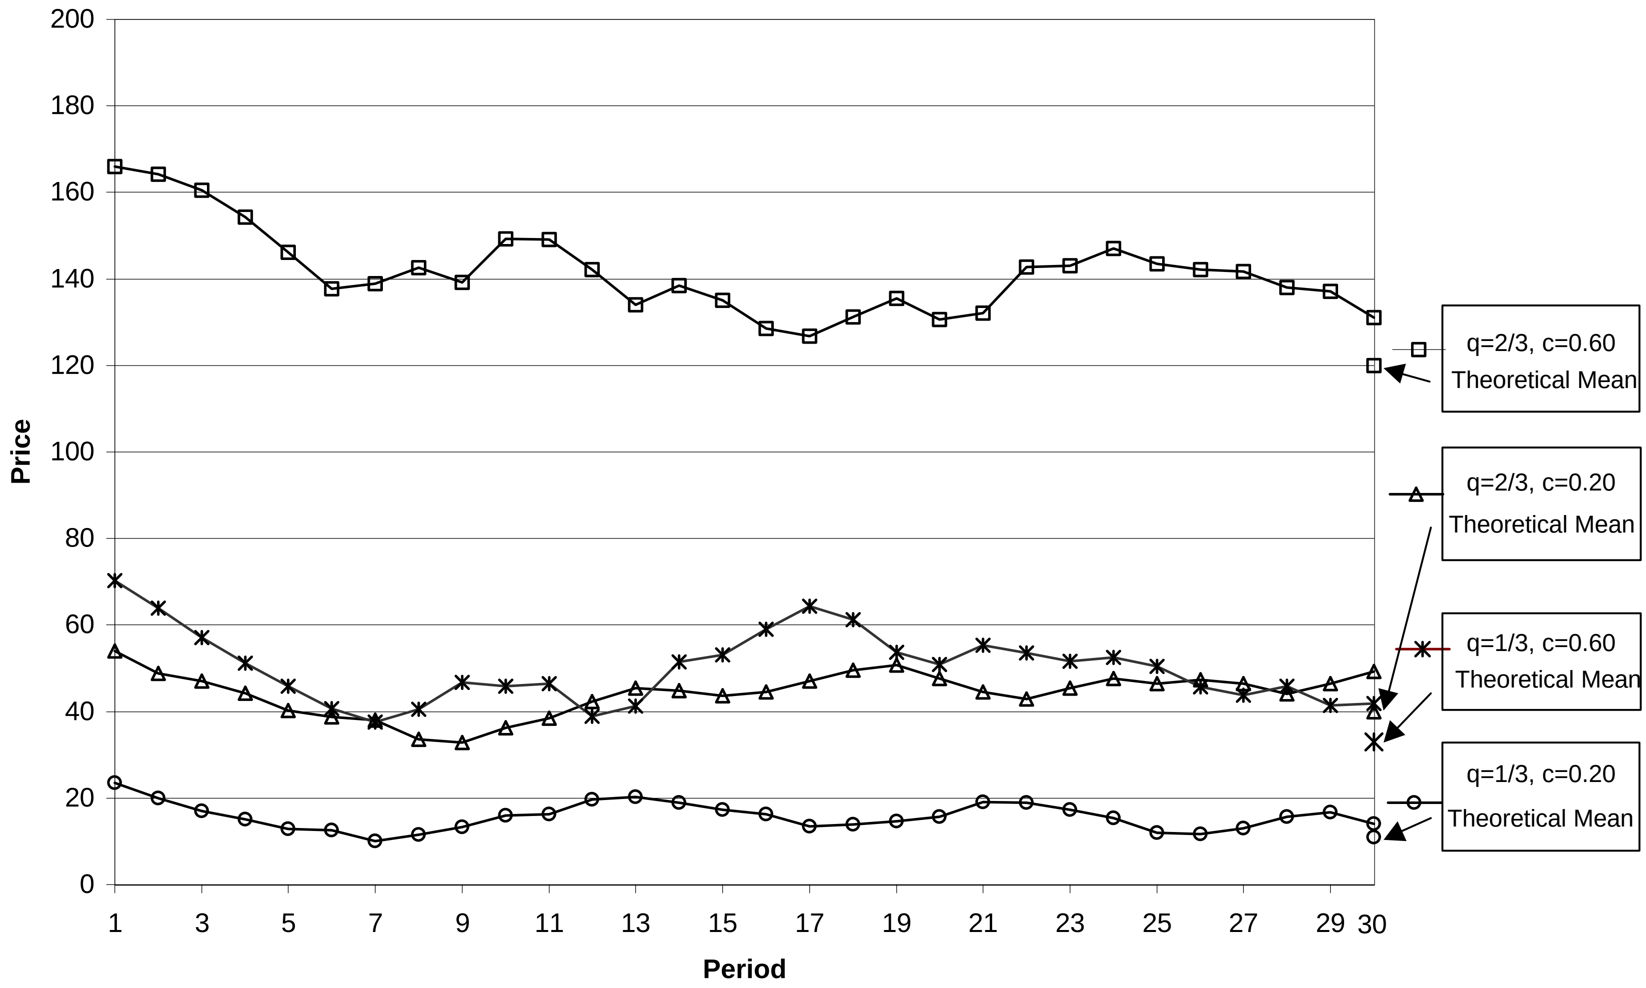

Figure 6 - Mean Transaction Price

Panel C: 6 robot buyer sessions

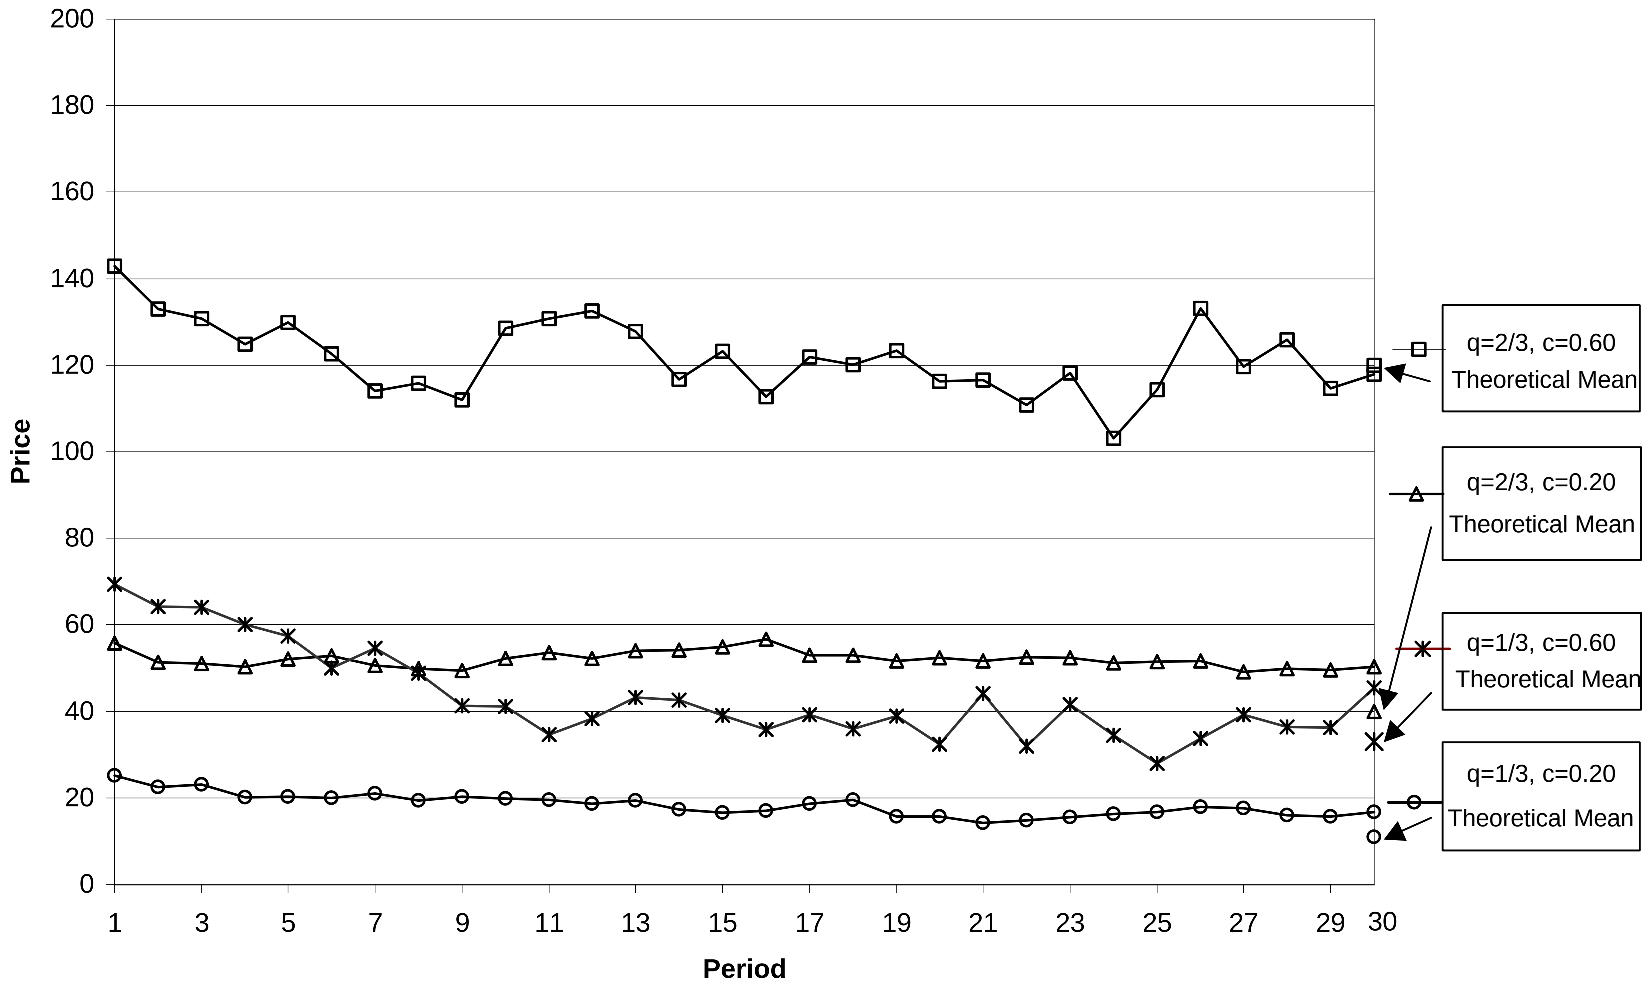

Figure 6 - Mean Transaction Price

Panel D: 12 robot buyer sessions

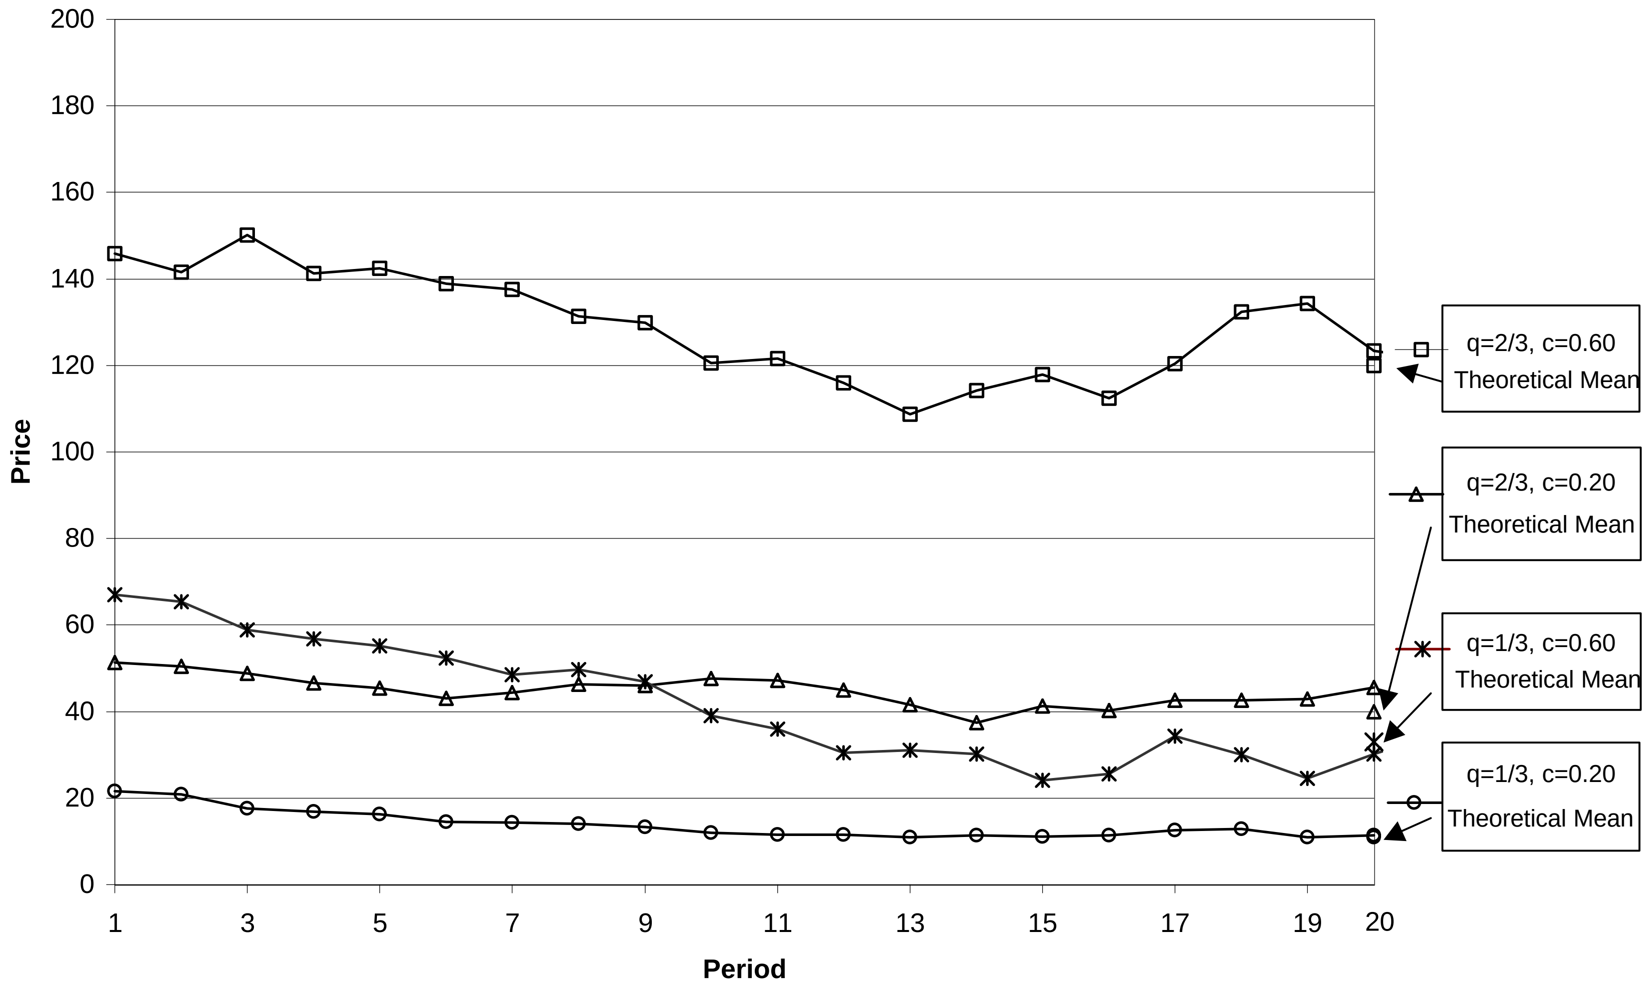

Figure 7 - Price Frequency: 6 Human Buyer Sessions
Panel A: q=1/3, 20-cent search cost

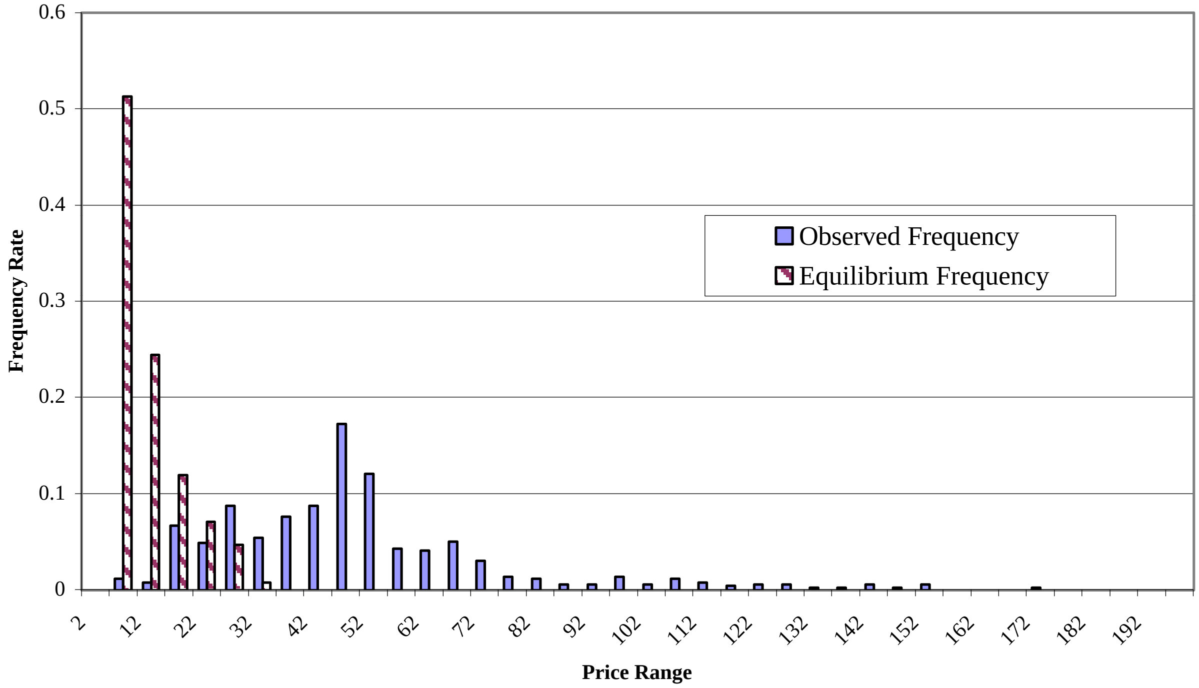

Panel B: q=1/3, 60 cent search cost

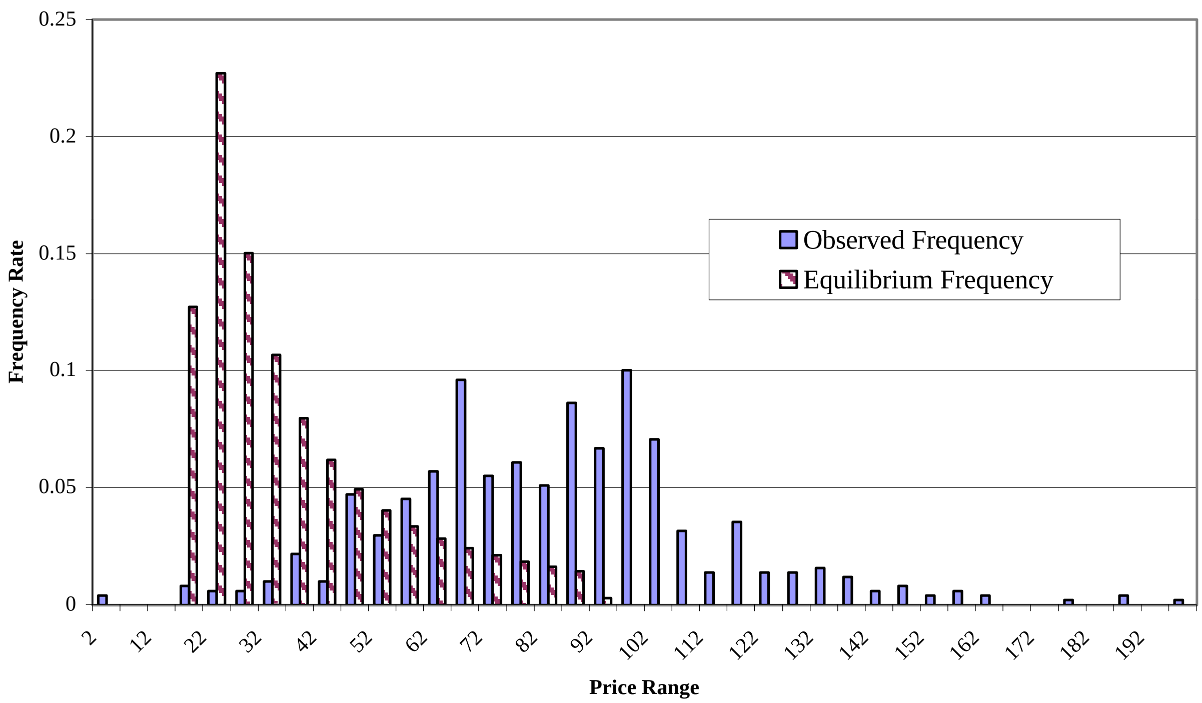

Figure 7 (continued) - Price Frequency: 6 Human Buyer Sessions
Panel C: q=2/3, 20 cent search cost

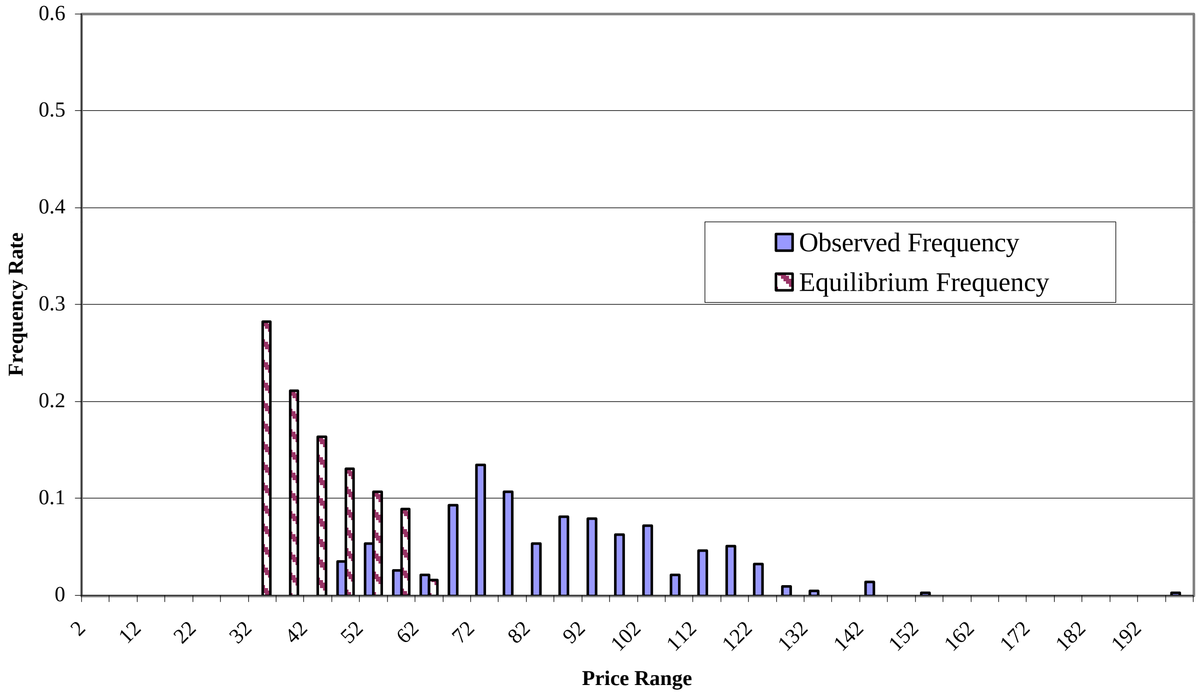

Panel D: q=2/3, 60-cent search

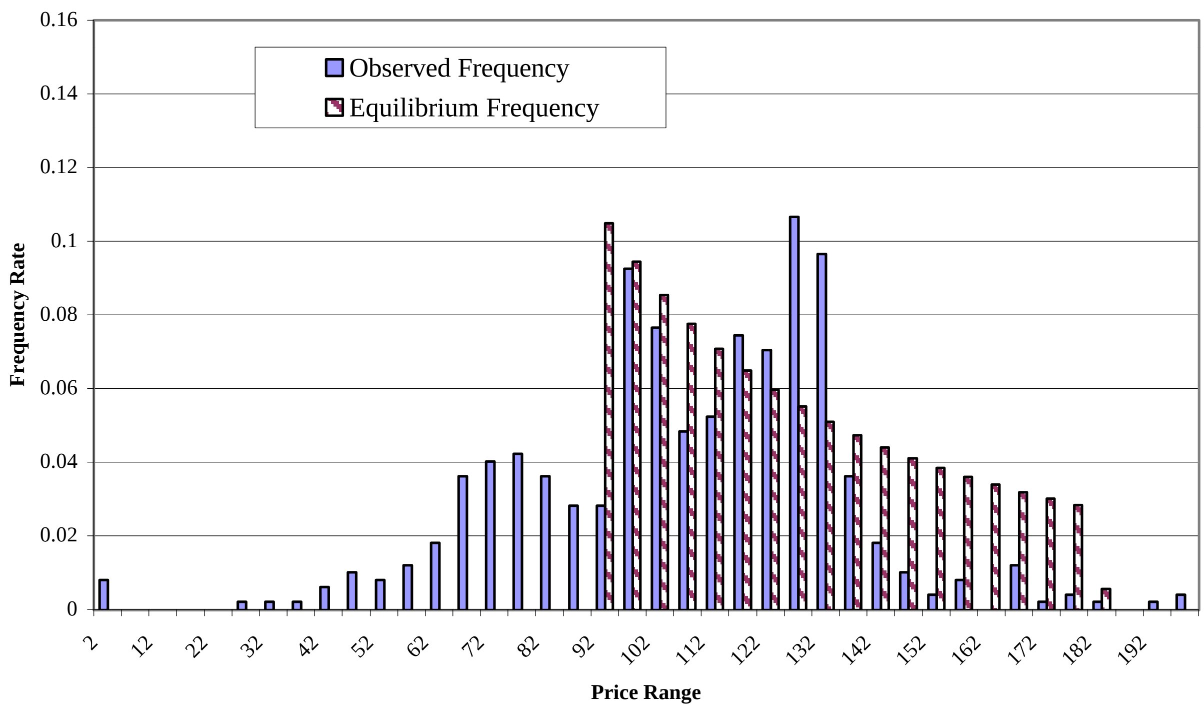

Figure 8 - Price Frequency: Many Robot Buyer Sessions

Panel A: q=1/3, 20 cent search cost

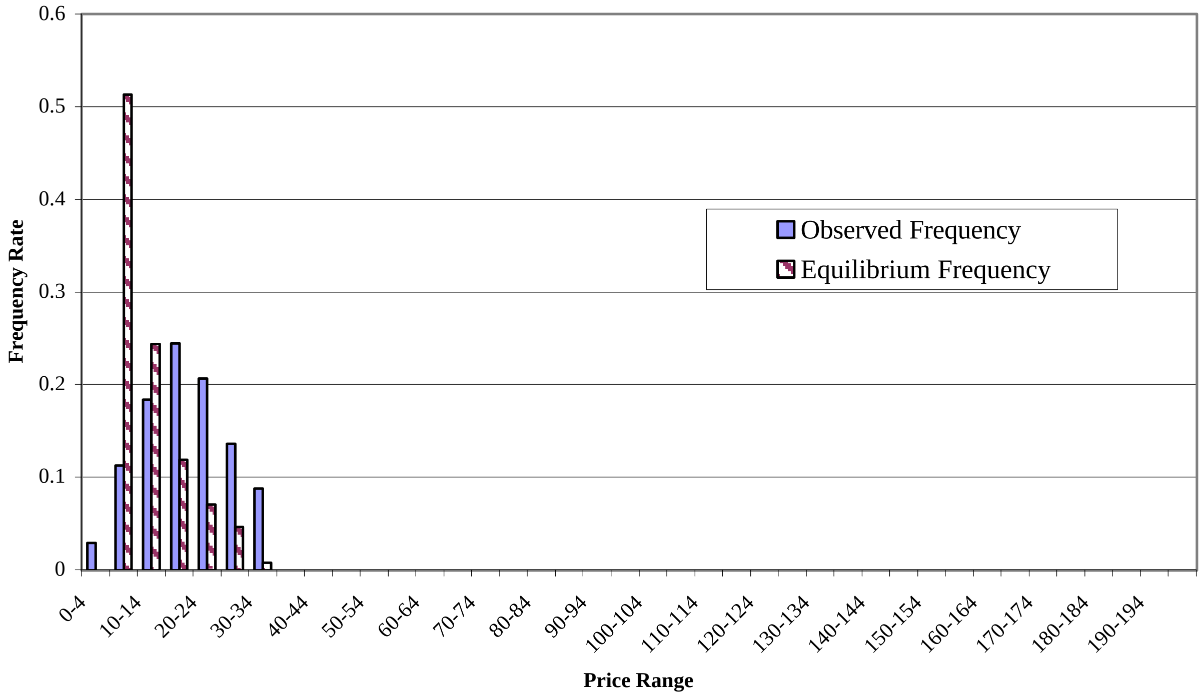

Panel B: q=1/3, 20 cent search cost

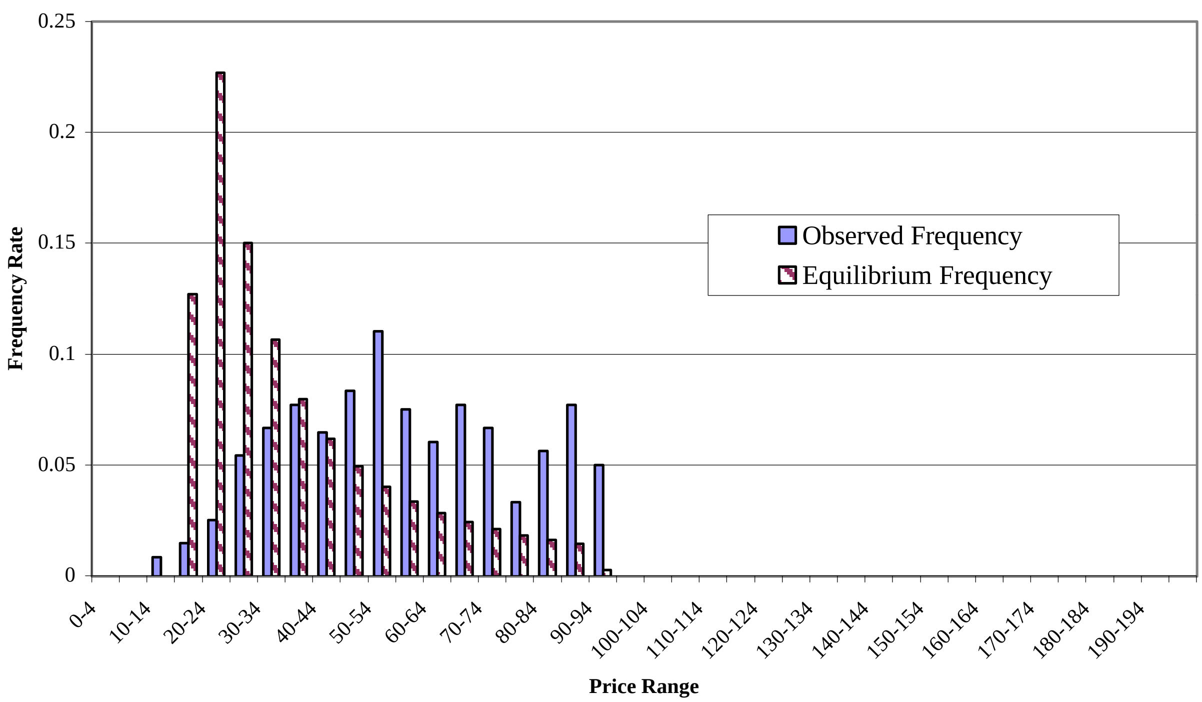

Figure 8 (continued) - Price Frequency: Many Robot Buyer Sessions
Panel C: q=2/3, 20 cent search cost

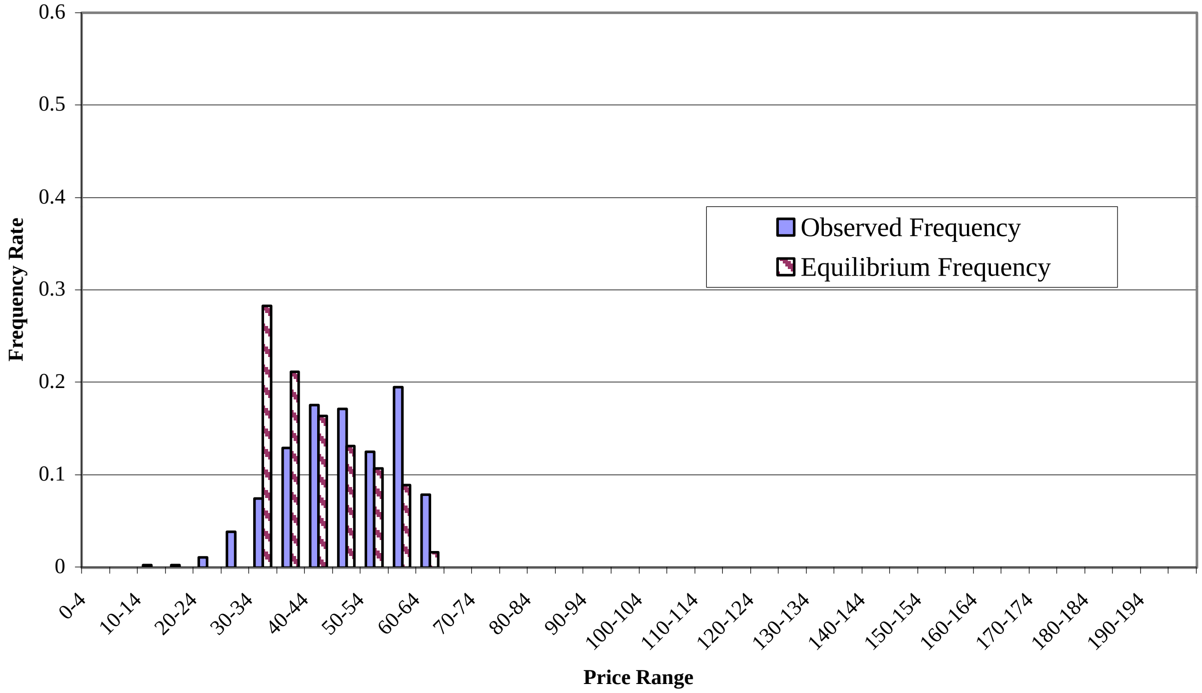

Panel D: q=2/3, 60 cent search cost

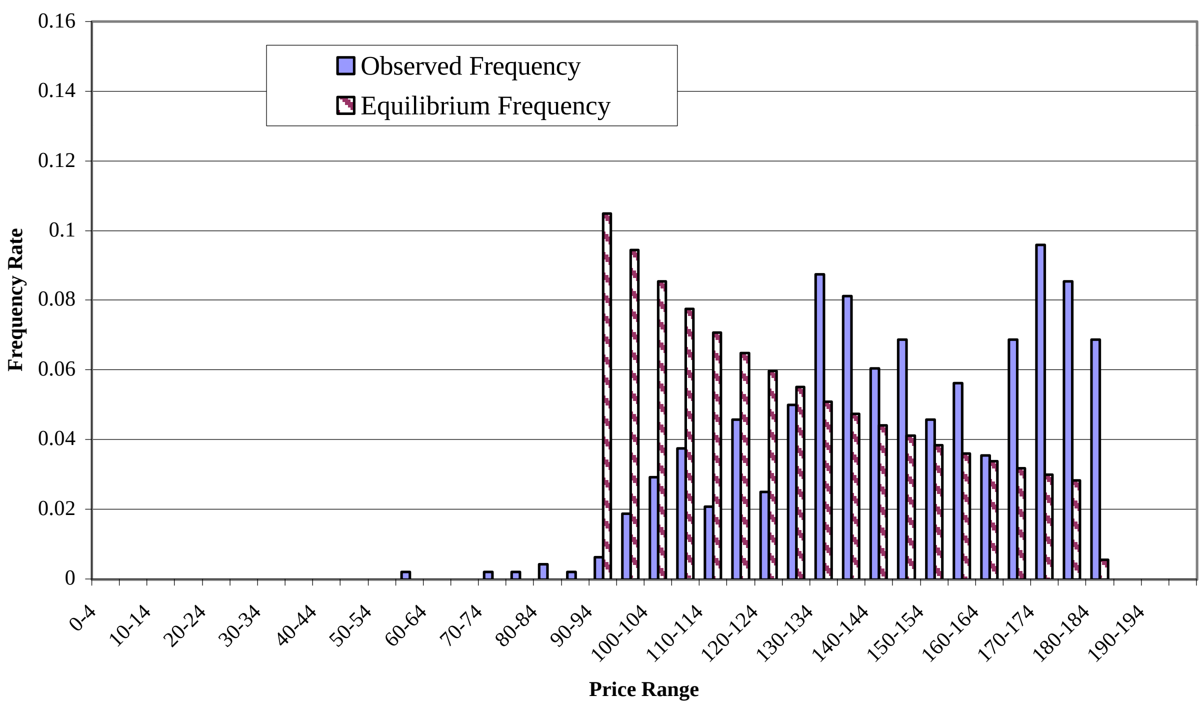

Figure 9:

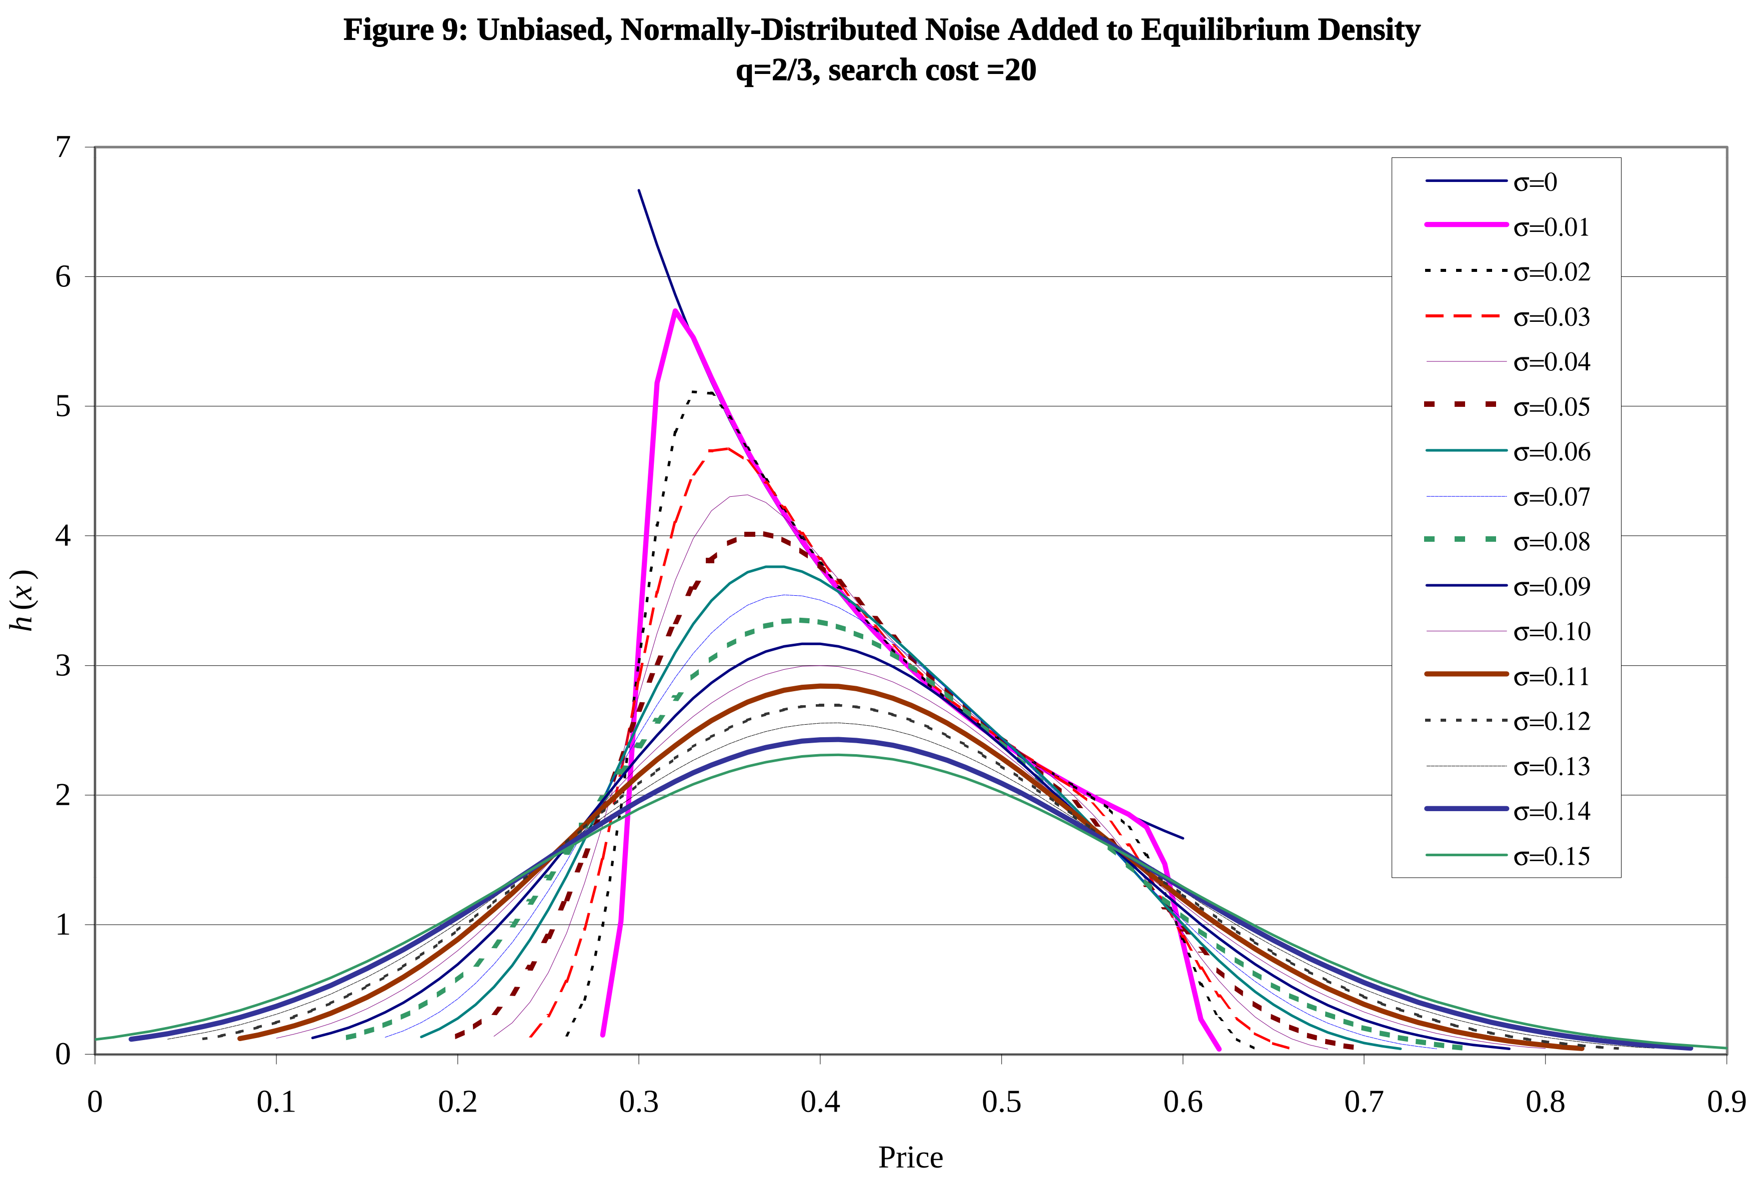
Figure 9: Unbiased, Normally-Distributed Noise Added to Equilibrium Density
q=2/3, search cost =20
7
6
5
4
3
2
1
0
h(x)
0
0.1
0.2
0.3
0.4
0.5
0.6
0.7
0.8
0.9
Price
σ=0
σ=0.01
σ=0.02
σ=0.03
σ=0.04
σ=0.05
σ=0.06
σ=0.07
σ=0.08
σ=0.09
σ=0.10
σ=0.11
σ=0.12
σ=0.13
σ=0.14
σ=0.15

Figure 10:

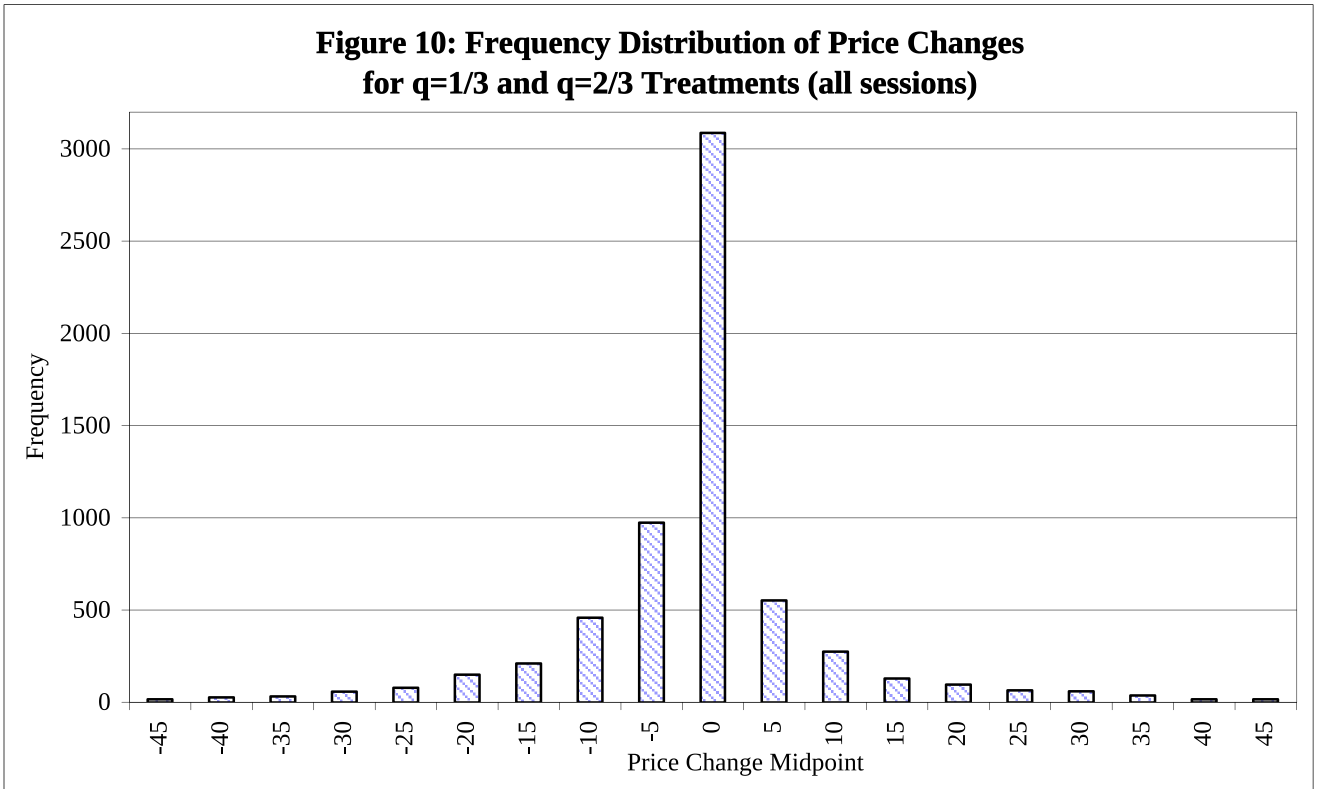
Figure 10: Frequency Distribution of Price Changes for q=1/3 and q=2/3 Treatments (all sessions)
Frequency
3000
2500
2000
1500
1000
500
0
-45 -40 -35 -30 -25 -20 -15 -10 -5 0 5 10 15 20 25 30 35 40 45
Price Change Midpoint

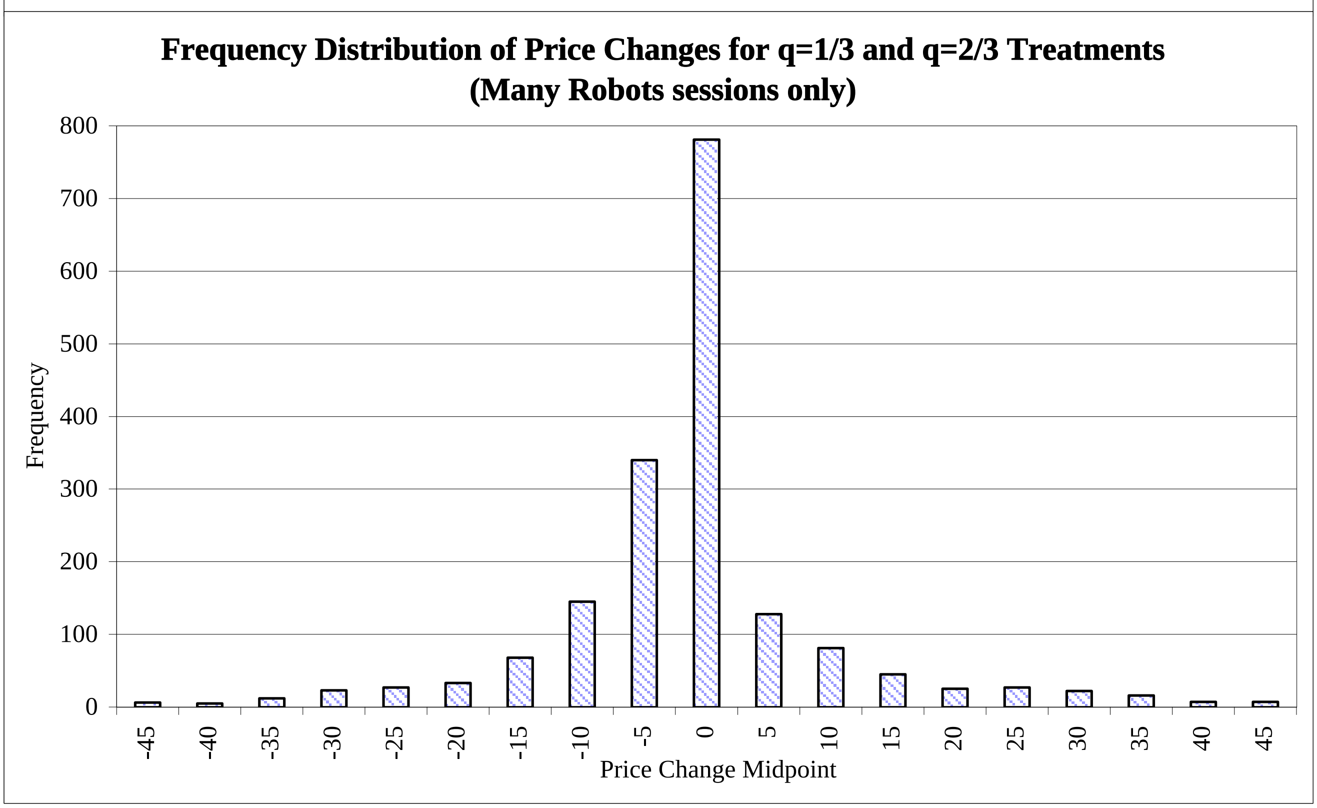
Frequency Distribution of Price Changes for q=1/3 and q=2/3 Treatments (Many Robots sessions only)
Frequency
800
700
600
500
400
300
200
100
0
-45 -40 -35 -30 -25 -20 -15 -10 -5 0 5 10 15 20 25 30 35 40 45
Price Change Midpoint