Customer Search and Market Power: Some Laboratory Evidence

by Timothy Cason
Department of Economics, Purdue University

and

Daniel Friedman
Department of Economics, UCSC

June 1999

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 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 the number of observed prices (one or two), the search cost, and the number and kind of buyers. When there are more buyers (especially robot buyers), observed transaction prices conform remarkably closely to theory (competitive *Bertrand* prices when buyers observe two prices and monopoly *Diamond* prices when buyers observe only one price). With human subject buyers we observe less extreme prices, but outcomes are closer to theory than outcomes in previous laboratory experiments with similar environments.

Acknowledgements:

We thank the NSF for funding the work under grants SBR-9617917 and SBR-9709874, Brian Eaton for programming assistance, and Sujoy Chakravarty and Sharad Barkataki and especially Garrett Milam for research assistance. We are grateful to Peter Diamond, David Easley and Andrew Muller, and participants at the 1999 Economic Science Association conference for helpful suggestions. We retain responsibility for any errors.

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@cats.ucsc.edu

Customer Search and Market Power: Some Laboratory Evidence

1. Introduction

Industrial organization practitioners have come to realize that market power is subtle, not simply a matter of structural concentration ratios. What matters are strategic opportunities. This point in some sense goes back to Bertrand (1883), who obtained very a different equilibrium prediction than Cournot with identical structure but a different market institution. More recently, Diamond (1971) demonstrated that even a very small customer search cost can drastically alter strategic opportunities and completely reverse the Bertrand result.

Generations of skeptical students have learned and forgotten these theoretical points, and practitioners also generally ignore the potential impact of search costs. Laboratory evidence so far supports the skepticism. As we will explain later, previous laboratory studies have detected only minor effects from customer search costs.

Our laboratory investigation aims to give search costs a better shot at affecting the exercise of market power. Seemingly minor changes in the design greatly improve the predictive accuracy of a Diamond-style model. Indeed, actual outcomes in our more favorable treatments (large numbers of automated buyers) show that human sellers approximate the extreme theoretical predictions rather closely. That is, even when search costs are quite small, surplus goes almost entirely to sellers when the model says it should, and goes almost entirely to buyers when the model says it should. We conclude that the practical importance of customer search cost for market power may have been underestimated.

We begin in the next section by summarizing the relevant theory. Section 3 describes our experiment, which varies the buyer search cost, the number of prices seen in each search, and the number and nature of buyers. Section 4 presents the results, including an analysis of buyer search behavior as well as sellers' posted prices and transaction prices. Section 5 compares our results to those of previous laboratory studies, offers some interpretations and implications, and lists possible avenues for future research. Instructions to Subjects are attached as Appendix A.

2. Theoretical Considerations

Bertrand (1883) models a one-shot posted offer market for a homogeneous good with no seller capacity constraints and no buyer search costs. The model predicts an extremely competitive outcome: in Nash equilibrium (NE) all transactions are at a price equal to (the second lowest) marginal cost, so buyers receive the entire surplus (or virtually all, when sellers have different marginal cost). 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.

Diamond (1971) demonstrates an equally striking but opposite outcome when buyers have positive search costs, however small. Assuming that buyers observe seller prices one at a time, the unique NE is the monopoly price and sellers get the entire surplus. The intuition is that positive search costs provide the seller with "local" market power over buyers who are in the seller's "store," so cream *skimming*, or choosing a slightly higher than other sellers (rather than slightly lower as in Bertrand), is profitable up to the monopoly price.

At least two more recent papers link the extreme predictions of these two models. 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.

Our laboratory experiment is based on another other link, the noisy search model of Burdett and Judd (1983). The model assumes a continuum of identical sellers with zero cost for producing a homogeneous good, and a continuum of buyers with identical willingness to pay (say \$2.00) for a single indivisible unit of the good and with identical search cost $c \geq 0$. Each seller posts a single price at which he is prepared to sell as many units as buyers order. Each buyer initially sees the price of one randomly selected seller with probability $q \geq 0$ or two such prices with probability $1- q \geq 0$. The buyer can purchase at the (lower) observed price, or can pay c to resample (one or possibly two prices according to the same q) from the seller distribution.

It is a simple corollary of Burdett and Judd's results (or of the simple intuition described above) that one recovers the Diamond equilibrium in the boundary case $q=1$ (where each buyer sees one price in each sample) and recovers the Bertrand equilibrium in the boundary case $q=0$ (where each buyer sees two prices in each sample). Since the buyer demand is inelastic on the relevant margin (a single indivisible unit per buyer), there is no efficiency loss in the Diamond equilibrium. Moreover, there is no actual search in equilibrium since all sellers charge the same price and hence buyers have no incentive to bear search costs.

Thus the model has several extreme but testable predictions.

- With sample size 1 ($q=1$), the price is at the monopoly/Diamond level $p=\$2.00$, and sellers receive the entire surplus.
- With sample size 2 ($q=0$), the price is at the competitive/Bertrand level $p=\$0.01$, and buyers receive the entire surplus.
- These predictions are the same whether search costs are high or low, as long as they are positive.
- Buyers don't actually search even when search costs are low.

3. Laboratory Procedures

Conceptually, the posted offer market institution we use is a simplification of the Customer Market described in Cason and Friedman (1998). The laboratory implementation, however, is entirely new and uses Java applets for buyers and sellers logged into a www browser. As shown in Figure 1a, each period each seller uses a scroll bar to post price and then is randomly matched with a fraction of the buyers (possibly none if there are only a few buyers). As shown in Figure 1b, each buyer sees one or two of the sellers' current posted prices, and clicks the appropriate button to indicate her choice: accept a posted price she sees, search, or quit for that period. In the second case, 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, or time has expired (a rare event), the buyers and sellers review their own profit or loss and see all posted and transaction prices for the past period, as in Figure 1c. 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 A and were assigned randomly to computer stations. The experimenter read the instructions aloud, allowed subjects to ask questions, and gave subjects a short quiz and conducted two practice periods. Subject payments (including a $5.00 show-up payment) were calibrated to average about $20 for sessions that lasted about 90-100 minutes, including instruction time. Depending on the quality of decisions (and luck), actual payments ranged from about $7.00 to over $30.00.

Sellers had zero production cost and no capacity constraint. Buyers all 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 focuses on three treatment variables: search cost *c*, sample size (one or two prices) and buyer population. The search cost was controlled at two levels: 20 cents and 60 cents. Except in pilot sessions we did not bother with the zero search cost case since it replicates standard posted offer markets.

We varied the price sample size between one and two prices as a within-session treatment variable. About one-half of the sessions began with 20-30 periods of a single price per sample, followed by 20-30 periods of two prices per sample. The other sessions reversed the order of the sample size treatments. When a buyer searched, the new sample was drawn after replacement of the earlier samples.

The baseline buyer population is six human buyers. Pilot experiments suggested that increasing the number of human buyers didn't have strong enough effect to justify its cost. However it is cheap and theoretically interesting to replace the human buyers by computer algorithms (or "robots") that follow equilibrium reservation price strategies. This gives the theory (which focuses on sellers' price posting strategies) its best shot, since it allows sellers to respond to known, stable buyer behavior and eliminates possible concerns about fairness.

Our robot implementations were as straightforward as possible, with the number of robots controlled at 6, 12 and many. In these robot buyer sessions, we told sellers that there were no human buyers and we told them the exact number of robots. We also described the robot buyer strategies, which differed between the price sample size treatments. In the treatment with one price per sample, the robot buyers followed the equilibrium strategy to buy at any observed price less than or equal to their resale value of $2.00, and to incur the search cost to draw another price if the observed price exceeds $2.00. In the treatment with two prices per sample, in equilibrium the reservation price (above which the buyer searches) is zero; to avoid such trivialities we set the reservation price to ten cents instead of zero. In other words, in this two-price sample treatment robot buyers immediately accept the lowest price observed if it is below $0.11. If both prices exceed $0.10 the robot buyer incurs the search cost to draw another sample of two prices.

Recall that sellers are not capacity constrained, and each buyer has a resale value of $2.00 and demands one unit. Increasing the number of robot buyers from 6 to 12 to many therefore simply produces an outward shift in demand and does not affect the equilibrium predictions. We increased the number of buyers to reduce the sampling variance of sellers' sales quantities. For example, consider the robot buyer treatment with one price per sample, and suppose that all sellers post prices less than or equal to $2.00. Seller prices are drawn independently for display to buyers, and buyers do not search because prices do not exceed $2.00. The expected number of sales for each seller is therefore (number of buyers)/(number of sellers), with a variance of (number of sellers-1)(number of buyers)/(squared number of sellers). The ratio of the standard deviation to the expected value decreases as the number of buyers increases, so in effect the sellers observe less noise in the feedback regarding the relationship between their price and their sales quantities. For example, holding the number of sellers fixed at 6 as we do and increasing the number of buyers from 6 to 12 decreases the ratio of the standard deviation to the mean from 0.913 to 0.645. We think that such reduced noise will aid learning. In the treatment with "many" robots, we eliminated the sampling variance by replacing the random number of sales with its expected value.[1] This last treatment should promote the fastest learning.

3.2 Design

In tests of the new user interface we varied both the number of buyers and sellers, but the main experiment holds the number of sellers constant at six. The buyer treatment (with one human plus three robot conditions) and the search cost treatment are fixed within each session but vary across sessions.

Table 1 lays out the design of our 18 sessions. It approximates a factorial design with 4 buyer conditions x 2 search cost conditions.[2] Each of the eight cells includes 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, which is randomized across treatments. There are a few minor irregularities, e.g., 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. Several statistical checks indicate that these irregularities had no effect on our conclusions.

4. Results

We begin with graphs of some individual sessions in Section 4.1, and then summarize the mean prices and market efficiency in Section 4.2. Section 4.3 presents some statistical tests of

[1] The formula is a bit more complicated when the sample size is two. Assuming again that all sellers post prices not exceeding the reservation value, the demand seen by the seller whose price is *r*-th highest among the six sellers is (number of buyers per seller) x 2 x (r-1)/5. Again there is sampling variance around this expected value except in the "many" robot buyer treatment.

[2] Not shown on this table are 60 or more subsequent periods in which some buyers observed a sample of two prices while other buyers observed a sample of only one price, within the same period. In these periods the independent likelihood of drawing two instead of one price was exogenous and common information, and this environment leads to price dispersion in equilibrium (Burdett and Judd, 1983). These additional periods are reported in a companion paper (Cason and Friedman, 1999).

the predictions of the Diamond and Bertrand models, and Section 4.4 analyzes buyer search behavior.

4.1 Overview of Some Representative Sessions

Figures 2 through 5 summarize some illustrative sessions, for each of the four buyer type and number treatments. The solid circles represent posted prices that result in transactions, and the open circles represent unaccepted price offers. Figure 2 summarizes inexperienced session PU7 with 6 human buyers and 60 cent search costs. In the initial run with one-price samples, prices begin a bit above the midpoint of the feasible price range ($1.00), and they slowly drift upward. By the last half of this run prices reach the $1.70 - $1.90 range, slightly below the Diamond prediction of $2.00. Following the shift to two-price samples prices fall steadily, and after about 10 periods they begin to approach the competitive prediction of 1 cent. Thereafter transaction prices remain around $0.20, although (as indicated by open circles) sellers sometimes unsuccessfully attempt to sell at higher prices.

Figures 3 and 4 present sessions with inexperienced human sellers and, respectively, 6 and 12 robot buyers. Figure 3 shows a UC session with 20 cent search costs. It displays a rather rapid and complete convergence to the Diamond equilibrium in the first run of 30 periods and then to the Bertrand equilibrium in the next run. Figure 4 shows a Purdue session with 12 robot buyers with a 60 cent search cost. Sellers in the first run with two-price samples always chose prices less than or equal to 5 cents after the fifth period. They adjusted prices quickly after the treatment switch to one-price samples, and the majority of transaction prices are within 15 cents of the $2.00 Diamond equilibrium. Significant variability remains throughout the second run, however.

Figure 5 summarizes 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. Sellers in the first run with one-price samples unerringly chose the Diamond/Monopoly equilibrium price of $2.00. They immediately dropped price in the second run with two-price samples almost to the $0.01 Bertrand equilibrium level. This is among the most striking realizations of extreme equilibria that we know of in experimental economics.

Impressions from the other 14 sessions are generally similar. The change of sample size from one to two (or two to one) always has a very strong impact in the theoretically predicted direction, but the exact equilibrium predictions seem more accurate with large numbers of robot buyers. We shall test these impressions formally in the following subsections.

4.2 Mean Prices and Efficiency

Overall impressions about pricing trends can be gleaned from Figures 6 and 7. With 60 cent search costs (Figure 6), mean posted prices for the one-price sample treatment are lower (and therefore are farther from the Diamond equilibrium prediction) as the number of buyers decreases. Mean prices are lowest with 6 human buyers. For the two-price sample treatment prices deviate noticeably from the $0.01 Bertrand equilibrium prediction only for the human buyers treatment. With 20 cent search costs (Figure 7), mean posted price for the one-price sample treatment can again be ordered by the number and type of buyers. With human buyers, mean price do not manage to rise above $1.20 cents. In the two-price sample treatment, however, mean prices are quite near the $0.01 equilibrium in all treatments, except for the first 4 and last 4 periods with 6 human buyers.

It is standard in laboratory markets to measure the efficiency of market outcomes using the fraction of potential gains from trade (or surplus) actually realized by subjects. Each transaction in this environment leads to a potential $2.00 surplus, which will be dissipated when buyers search or fail to transact. Sellers almost never priced above robot buyers' reservation price, so efficiency was virtually 100 percent in all robot buyer sessions. Table 2 shows that efficiency was usually high in the human buyer sessions as well, averaging in excess of 98 percent in the two-price sample treatment. Efficiency was lower, but still averaged over 92 percent, in the one-price sample treatment. In section 4.4 below we shall analyze how the unrealized surplus in this treatment arose from buyer searches and demand withholding.

4.3 Hypothesis Tests

Standard hypothesis tests of the exact Diamond and Bertrand pricing equilibria are not very informative. This is because predicted prices are at the extreme boundaries of the range of feasible prices, so any deviations from the predicted prices are necessarily on one side of the relevant equilibrium. Consequently, the assumption that errors are distributed about the equilibrium with a mean of zero (under the null hypothesis) is incorrect. This subsection therefore focuses on comparative static predictions of the equilibrium models.

Support for the comparative static prediction regarding the sample size observed by buyers (one price versus two prices) is quite obvious from Figures 2 through 7. Prices are clearly greater in the one-price sample treatment in all four buyer number/type conditions. With robot buyers virtually no overlap exists between the price distributions for the different price sample size treatments. With 6 human buyers the distributions in the two price sample size treatments overlap somewhat, particularly with 20 cent search costs. Nevertheless, we can easily reject the null hypothesis that prices are drawn from the same distribution for the one-price and two-price sample treatments. For example, a Wilcoxon signed-rank test on the paired final prices in each run for each seller indicates that prices are significantly higher at the end of the one-price sample treatment than at the end of the two-price sample treatment. The Wilcoxon *p*-value never exceeds 0.001 in any of the (4 buyer number/type * 2 search cost) = 8 treatments.

Another striking implication of the theoretical model is that it predicts no effect for the search cost treatment. The intuitive alternative hypothesis is that prices are higher for higher search costs. The mean prices shown in Table 3 indicate very little difference in average posted prices for the two search cost treatments, except for the inexperienced human buyer sessions. Indeed, for the robot buyer treatments with one-price samples, mean prices are often higher for the lower (20 cent) search cost.

To test formally whether the prices are different in the two search cost treatments we must account for the fact that offer prices are not independent; in particular, the same sellers within a session post multiple prices across periods. We therefore employ a random-effects error structure, with the seller as the random effect. Using this error structure we reject the null hypothesis of no difference in prices in the two search cost treatments in the inexperienced 6 human buyers treatment, both for the two-price sample (t=2.31) and for the one-price sample (t=5.73) treatments. In both cases mean prices are about 50 percent higher with 60 cent search costs than with 20 cent search costs. In the other datasets for the two-price sample treatment the data do not reject this null hypothesis of no price difference.[3] For the one-price sample treatment, when prices are significantly different in the robot buyers data they are higher in the low search cost treatment than in the high search cost treatment. We therefore conclude that prices increase with search costs only for inexperienced sessions with human buyers.

The mean prices in Table 3 do not account for changes in prices across periods, and Figures 6 and 7 suggest that time trends could be important, at least for early periods in runs with human buyers. Table 4 presents estimates of a random-effects regression for mean transaction

[3] The likelihood function for the random effects model failed to converge for the estimates with two-price samples for 600 robot buyers and inexperienced subjects. The difference in mean prices (6.6 cents) is not economically

price by period; we suppress for brevity the qualitatively similar results for mean posted price. Individual sessions represent the random effect, and the explanatory variables include dummy variables for the treatments and their interactions, as well as a front-loaded time trend represented by 1/(period in run).

Prices are statistically significantly higher in the one-price sample treatment, and the point estimate indicates an economically huge increase of about 130 cents when buyers observe one rather than two prices in each sample. The human buyer interaction estimate indicates a partial but significant offset of about 35 cents with human buyers, again consistent with impressions from previous figures and tables. The search cost=60 cents dummy is not significant, consistent with theory and earlier tests. The interaction of search costs with the one-price sample treatment dummy is significantly negative, but it is small in size, has the wrong sign, and is offset by the (insignificant) search cost=60 cents dummy. The dummy variable for whether one-price samples were provided in the previous run is intended to capture potential hysteresis effects. The estimate is significantly positive, indicating that prices are higher in a two-price sample treatment that followed a run with one-price samples, compared to a two-price sample treatment that opened the session. The coefficient for 1/period is not significantly different from zero, which indicates no systematic time trend across periods once these other factors are accounted for.

4.4 Buyer Search Decisions

In equilibrium all sellers post identical prices, so buyers have no incentive to search. In the experiment, of course, sellers often post different prices, and buyers may find searching worthwhile. Our human buyers in runs with two price samples typically buy at the lowest price without searching: of the 760 initial samples of two prices drawn by human buyers, in only 21

---

significant, however, even if it turned out to be statistically significant.

cases (3 percent) do buyers search. Nearly all of these cases—17 of them—are in the 20 cent search, inexperienced data.

Searches are more common when human buyers sample only one price at a time. Of the 820 such initial samples, we observe 163 cases in which buyers search (and 39 cases in which they refuse to purchase). A substantial majority of searches occur in the 20 cent search cost treatment. Figures 8 and 9 display the distribution of prices leading to a purchase, prices leading to a buyer search and prices leading to a buyer rejection for the 20 cent search cost, one-price sample treatment. The upper panel of each figure displays the raw frequency, while the lower panel displays the search and rejection rate for each price range. The rate that the buyers search clearly increases with the price, especially for the experienced data shown in Figure 9. For the inexperienced data shown in Figure 8, search rates also rise with the offer price but the picture is clouded by the bimodal price distribution.[4]

For comparison to the upper panels of Figures 8 and 9, Figure 10 displays the raw frequency of purchase, search and rejection prices for the 60 cent search cost treatment with one-price samples. Search rates are substantially lower in this treatment, which jibes with the finding in Table 3 that in the (inexperienced) human buyer treatment, prices are higher with higher search costs.

5. Discussion

Our findings can be summarized as follows. In a laboratory setting where anonymous sellers post prices and where buyers must pay a search cost to sample individual sellers' prices, we find striking confirmation of a simple equilibrium model derived from Diamond (1971).

[4] The two modes reflect differences across sessions; mean offer prices were about 80 cents in UC5 and about 120-140 cents in PU5. The figure suggests that buyers have session-specific expectations about the prices offered by

Posted prices indeed are near the Diamond/monopoly price when the sample size is one and near the Bertrand/competitive price when the sample size is two. The predictions are especially accurate when we use robot buyers who play equilibrium strategies and when we eliminate sampling variance by using large numbers of such robots. Two other striking predictions are also supported: actual posted prices indeed are insensitive to the level of search cost, and actual buyer search is relatively rare. The main qualification is that the predictions are less accurate for inexperienced human buyers.

Our study was planned and (except for the last few sessions) executed independently of Abrams, Sefton and Yavas (1999), who also study a posted offer laboratory market with search. Their setup is quite similar to ours, with the following major exceptions: (a) buyers don't observe the distribution of posted or transacted prices in the previous period, (b) buyers do see sellers' identification numbers throughout each session, and (c) no robot buyers are used.[5]

They are able to detect an effect of the price sample size on posted and transacted prices, but it is relatively small, and in all treatments their prices are concentrated in the second quartile (corresponding to 0.50-1.00 in our parameterization). Their results (and their design features a and b) are consistent with earlier work by Davis and Holt (1996). By contrast, in our experiment prices are concentrated in the top quartile (1.50-2.00) for the one-price sample treatment and in the bottom quartile (0-0.50) for the two-price sample treatment, consistent with theory. Grether, Schwartz and Wilde (1988) obtained results closer to ours in three of four short one-price sample runs with anonymous sellers and lagged information on price distribution. Therefore it seems reasonable to conjecture that the equilibrium model of Diamond and Bertrand will do better

other sellers. Recall that all subjects learned all posted prices at the end of every period (cf. Figure 1c).

[5] Less important, we believe, are other differences such as (d) sellers in their market can change price when buyers search, (e) buyers can search at most once per period, (f) only one 25 period run per session, and (g) only 8 human buyers (with 8 human sellers) are used. Of course, their parameterization is also slightly different, with buyer values

when sellers do not have the opportunity to form reputations and when buyers have a better idea of the distribution of prices.

Two directions suggest themselves for follow up work. First, it might be interesting to look at the theory and empirics of fractional values of the sample size parameter q, i.e., when some buyers sample two prices and others sample one. We report some preliminary results along that line in Cason and Friedman (1999).

More importantly for understanding market power, we hope experimentalists will pursue our conjecture on when the equilibrium model predicts well. We showed that even small customer search costs can have a huge effect, and can transfer market power almost entirely from buyers to sellers. Hence it can no longer be taken for granted that buyer search is practically of little importance. In most field settings, buyers are often able to get good information on price distributions but sellers are able to form reputations to some degree. The impact of buyer search in such settings deserves further investigation.

---

of 120 instead of our 2.00 and search costs of 6 and 18 instead of our .20 and .60.

References

Abrams, Eric, Martin Sefton and Abdullah Yavas (1999), "An Experimental Comparison of Two Search Models," Manuscript (Penn State College of Business).

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

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

Cason, Timothy and Daniel Friedman (1998), "A Laboratory Study of Customer Markets," Manuscript (University of California at Santa Cruz).

Cason, Timothy and Daniel Friedman (1999), "Customer Search and Price Dispersion: A Laboratory Study," Manuscript (University of California at Santa Cruz).

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

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

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

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

Table 1:

Summary of Laboratory Sessions

Session Name	Search Cost	Number of Sellers	Number of Buyers

Table 2:

Overall Trading Efficiency, Including Loss due to Search Costs, for Human Buyer Sessions

Price sample size treatment	Experience Level	Search Cost=20 cents Overall Efficiency	Search Cost=60 cents Overall Efficiency
Two prices	Inexperienced	0.989	0.995
Two prices	Experienced	0.983	1.000
One price	Inexperienced	0.920	0.925
One price	Experienced	0.927	0.951

Table 3:

Mean Prices by Search Cost

Buyer Number and Type	Experience Level	Search Cost=20 cents Mean Price in cents (Std. Error)	Search Cost=60 cents Mean Price in cents (Std. Error)
Two-price sample size treatment			
6 Human Buyers	Inexperienced	39.2 (2.22)	60.2 (2.63)
6 Human Buyers	Experienced	85.0 (1.57)	88.2 (3.40)
6 Robot Buyers	Inexperienced	4.4 (0.12)	5.3 (1.83)
12 Robot Buyers	Inexperienced	3.9 (0.11)	4.2 (0.54)
600 Robot Buyers	Inexperienced	1.9 (0.14)	8.5 (2.23)
600 Robot Buyers	Experienced	3.2 (0.17)	2.4 (0.12)
One-price sample size treatment			
6 Human Buyers	Inexperienced	107.1 (1.79)	151.9 (1.51)
6 Human Buyers	Experienced	140.6 (1.33)	137.8 (2.73)
6 Robot Buyers	Inexperienced	186.6 (1.29)	154.8 (2.59)
12 Robot Buyers	Inexperienced	189.1 (0.79)	159.5 (1.77)
600 Robot Buyers	Inexperienced	198.1 (0.68)	188.0 (1.24)
600 Robot Buyers	Experienced	200.0 (0)	200.0 (0)

Table 4:

Mean Transaction Price Per Period Regression Results

Note: All models estimated using a random effects error structure, with the session as the random effect.

Variable	Estimate	Std. Error
Constant	-15.857*	6.839
Lagged Mean Transaction Price	0.347**	0.012
1/(period in run)	-0.046	1.950
One-price sample size treatment dummy	130.570**	2.674
Human buyer dummy*two-price sample size dummy	37.298**	7.774
Human buyer dummy*one-price sample size dummy	-32.347**	7.769
12 Robot buyers dummy	4.767	8.607
600 Robot buyers dummy	6.875	8.604
Search cost=60 cents dummy	4.515	5.722
Search cost=60 cents dummy*one-price sample size dummy	-5.937**	5.712
Experience dummy	9.241	6.854
Last run was one-price sample size treatment dummy	12.946**	1.399
Observations	782	
R^2	0.96	

Note: Baseline (omitted dummies) condition is 6 robot buyers, two-price sample size, search cost=20 cents, inexperienced, no previous run.
**Denotes significantly different from 0 at 1 percent.
*Denotes significantly different from 0 at 5 percent.

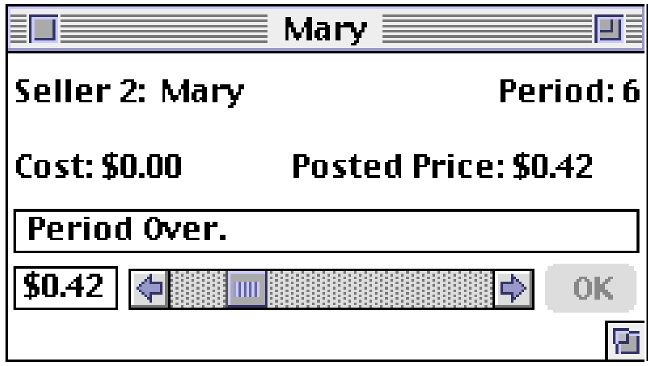

Figure 1a: Seller Window for Posting Prices

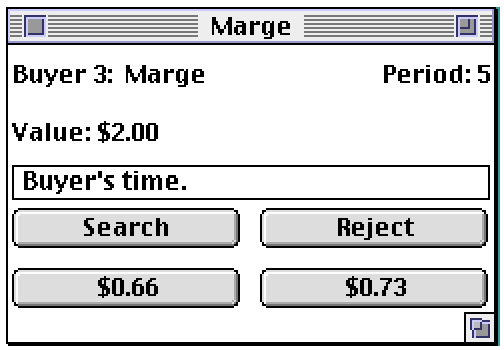

Figure 1b: Buyer Window for Purchase, Search or Reject Decision

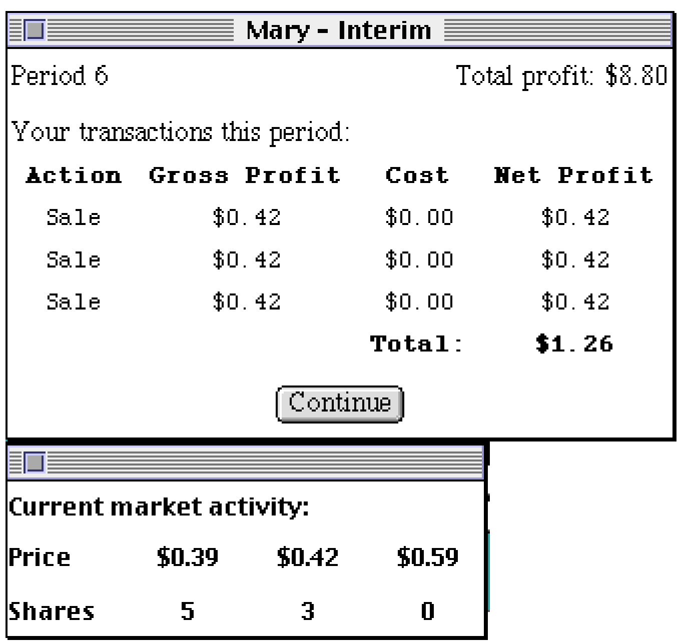

Figure 1c: Interim Screens with Profit Summary and All Posted Prices

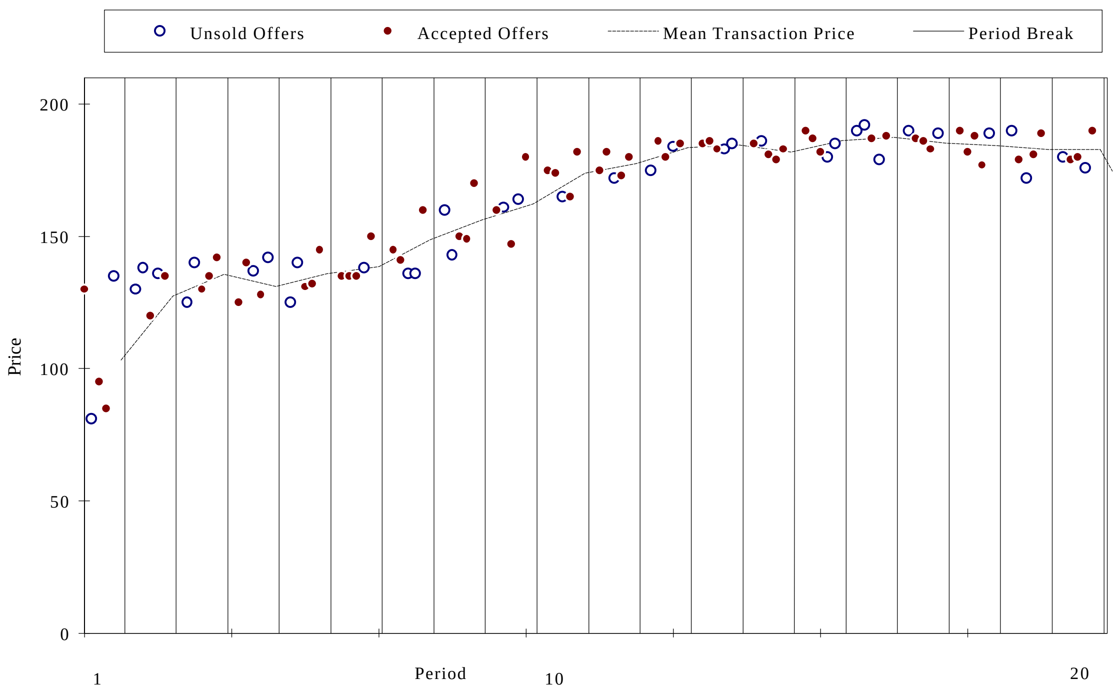

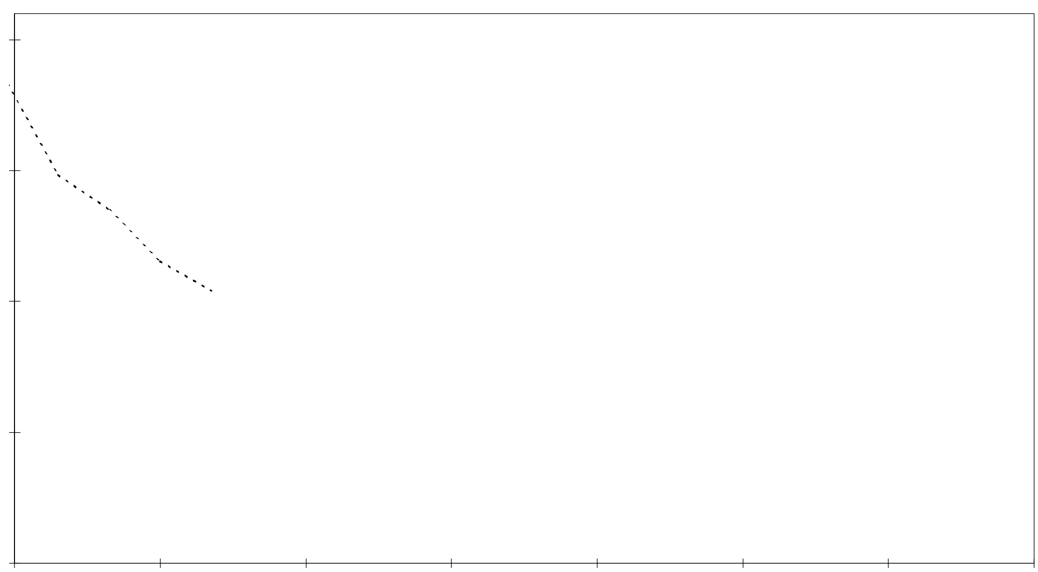

Figure 2: Session PU7

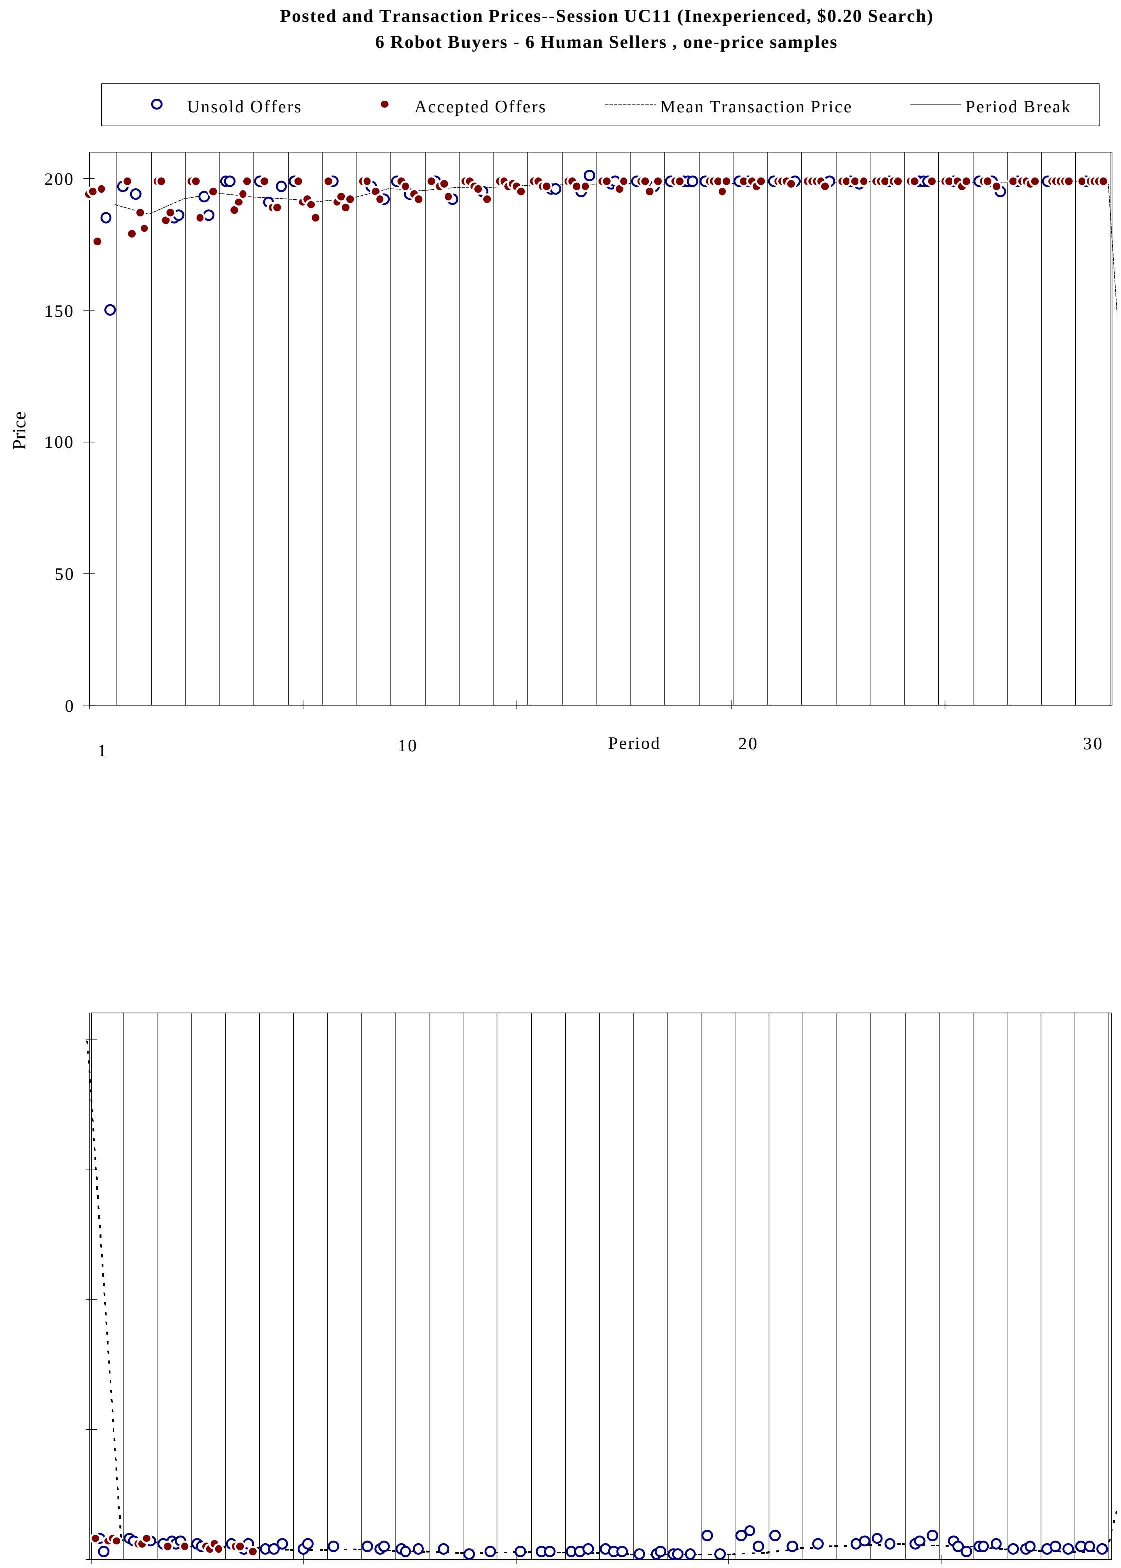

Figure 3: Session UC11

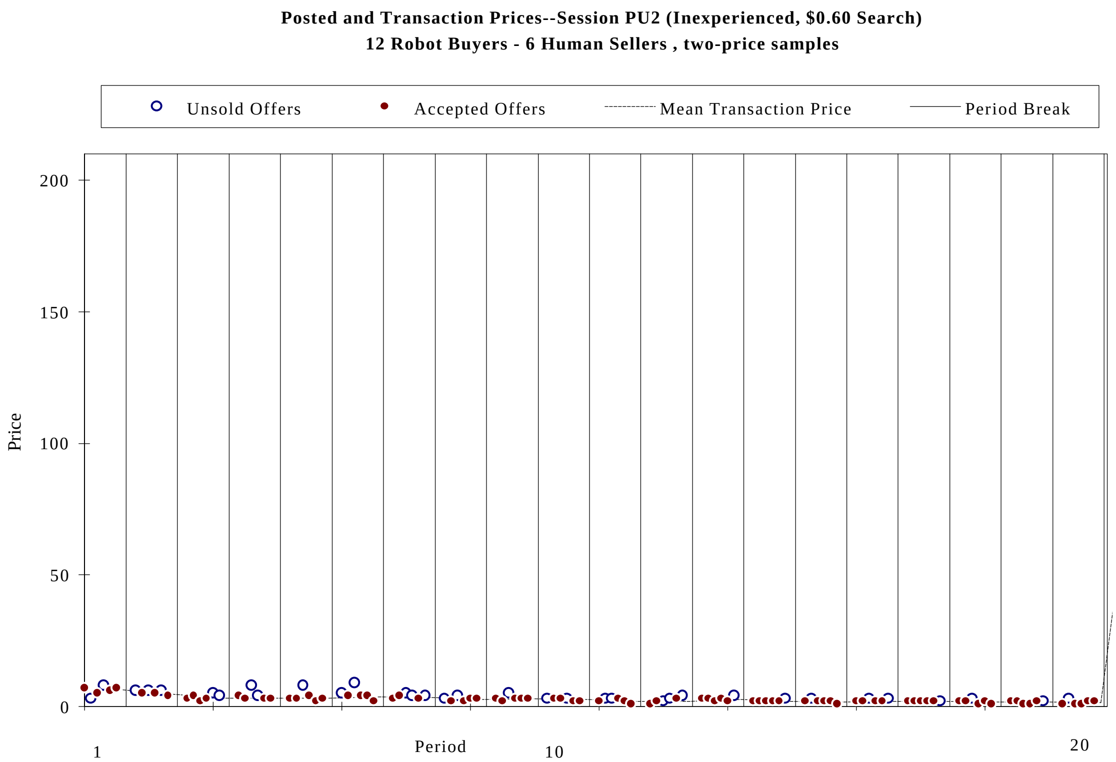

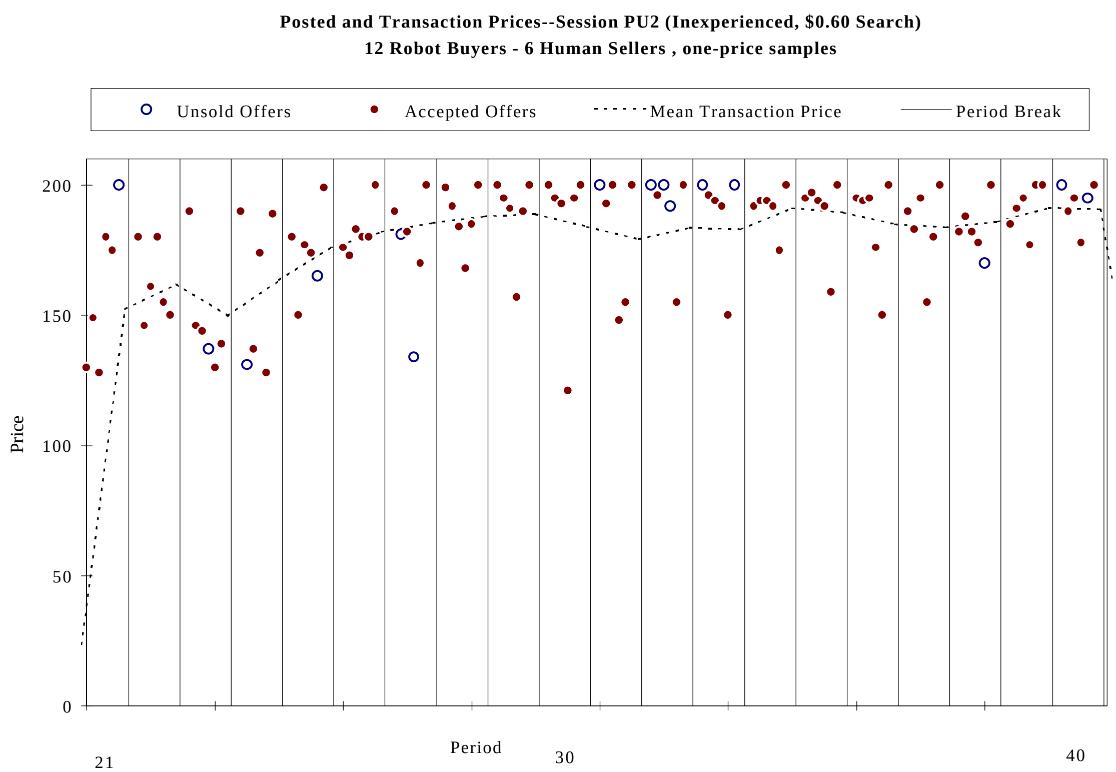

Figure 4: Session PU2

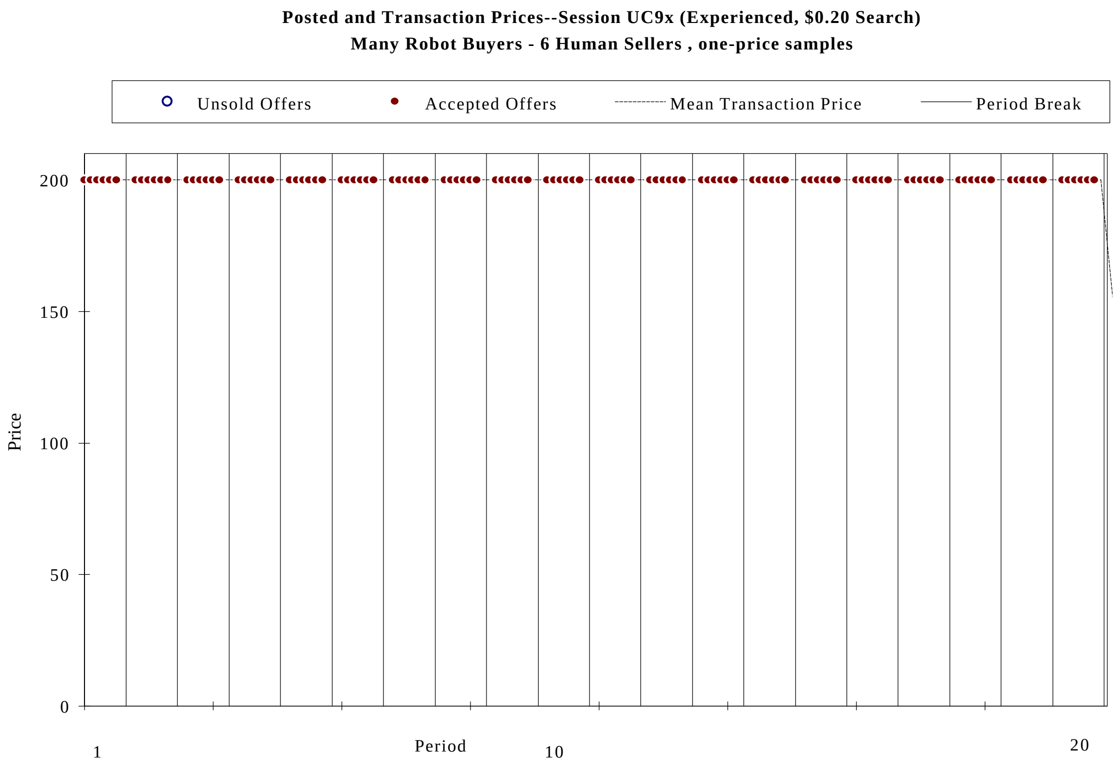

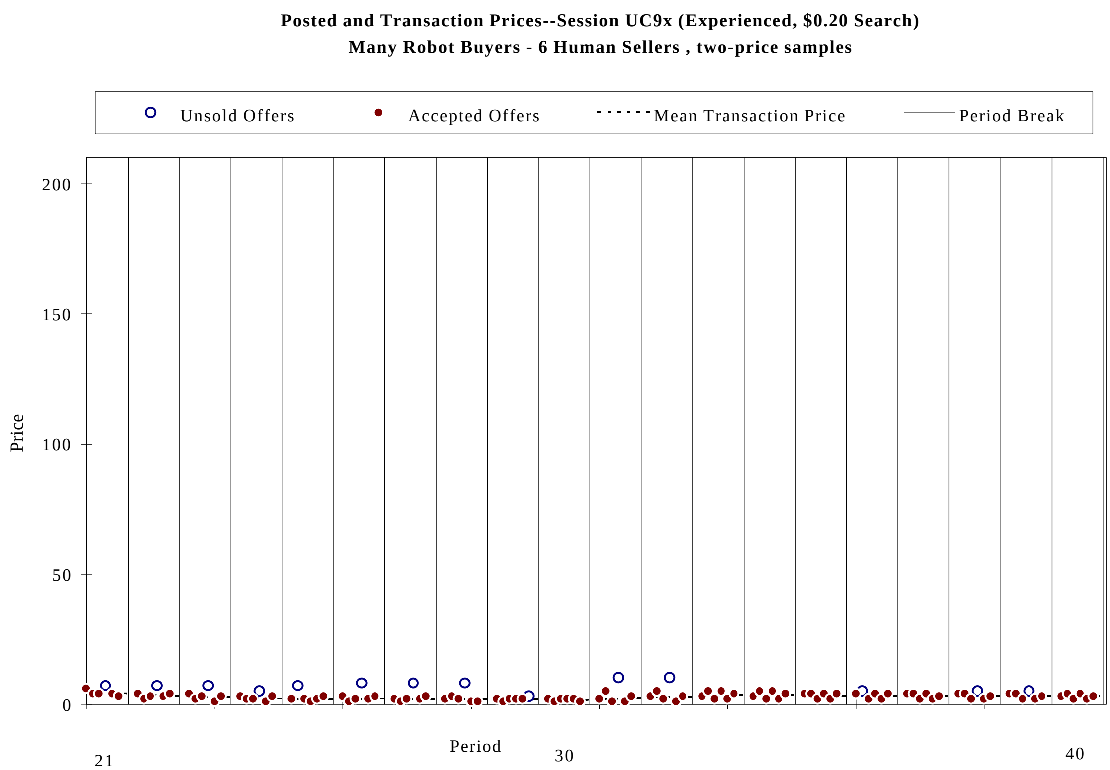

Figure 5: Session UC9x

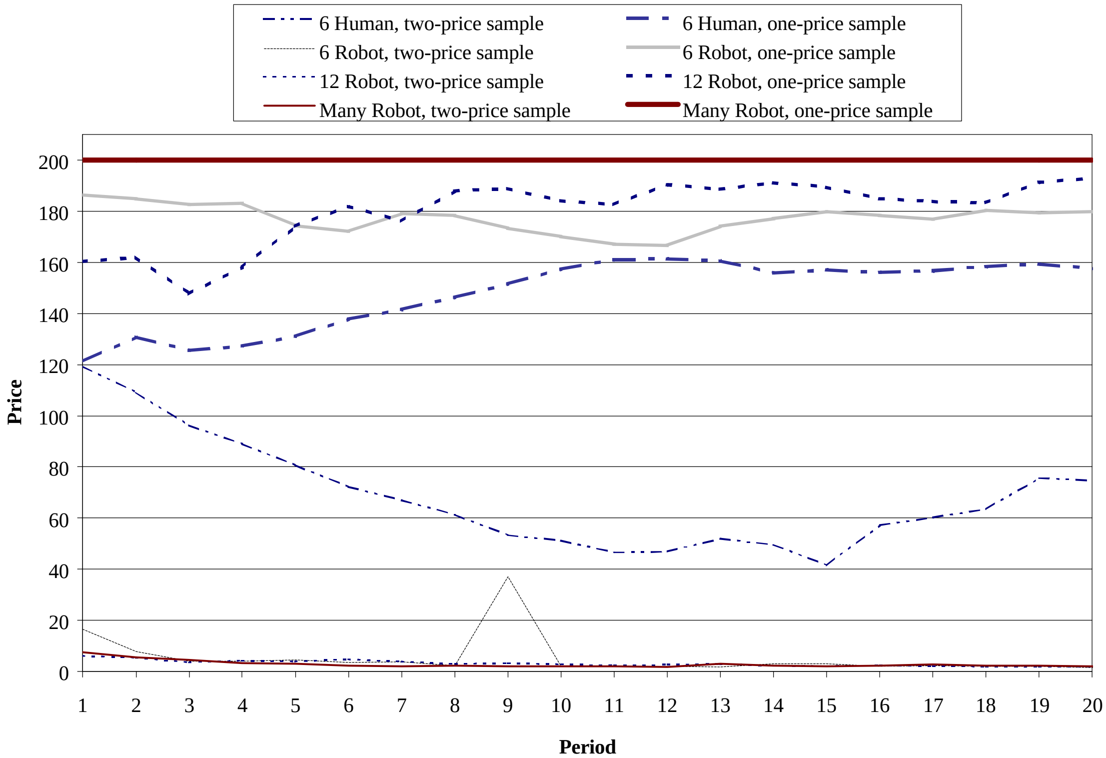

Figure 6: Mean Offer Prices for 60-cent search

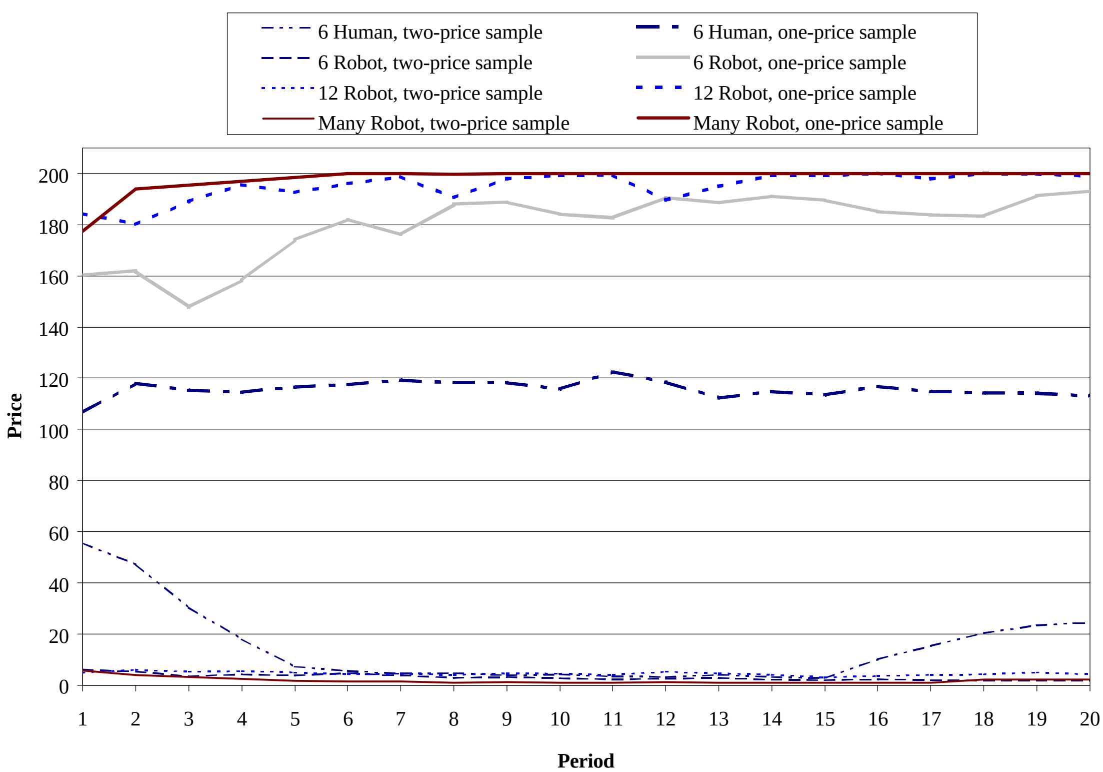

Figure 7: Mean Offer Prices for 20-cent search

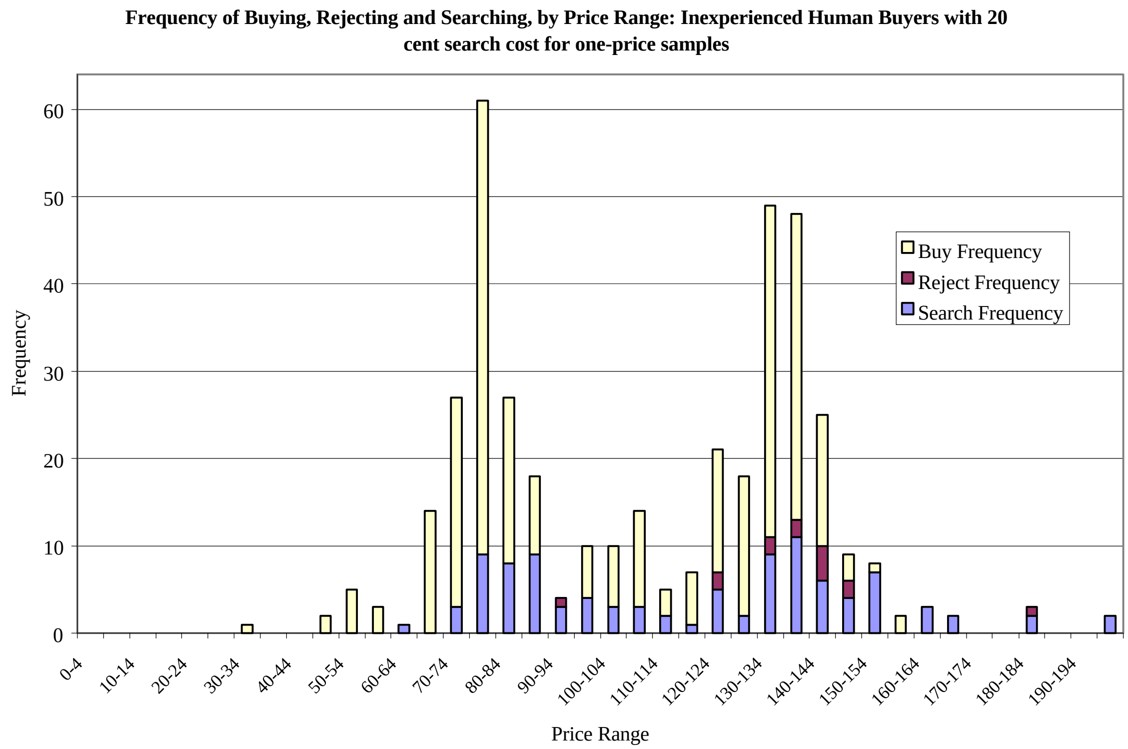

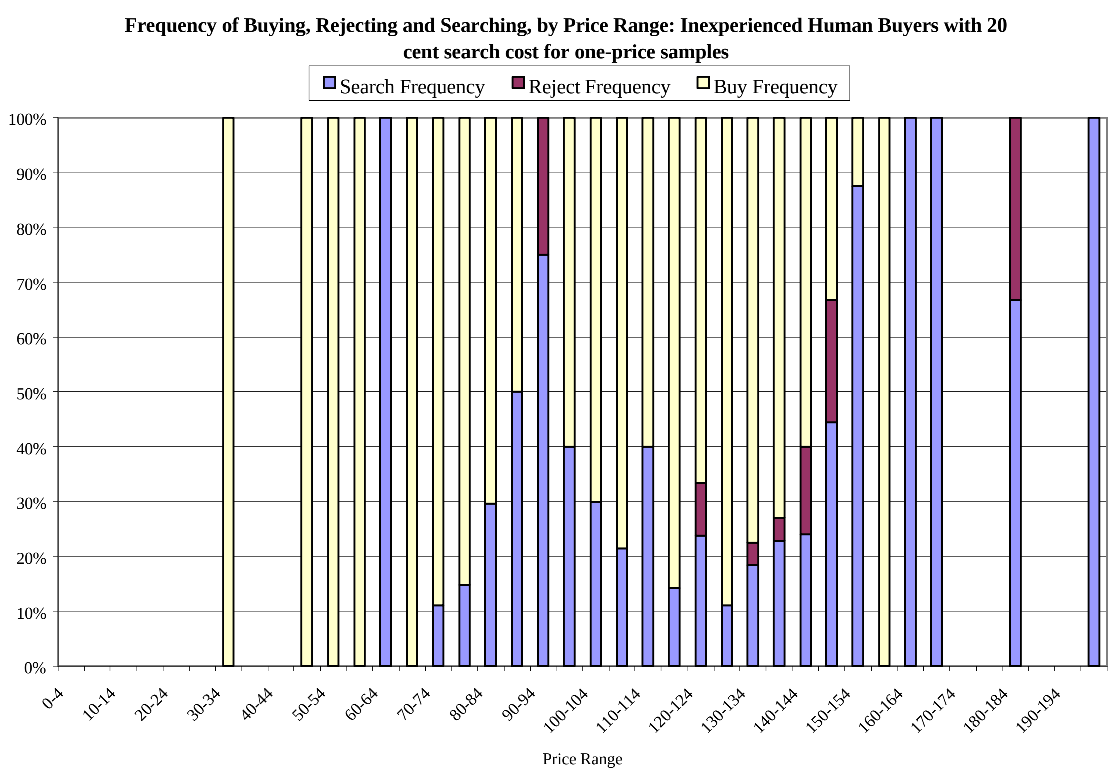

Figure 8: Distribution of Accepted Prices, and Prices Leading to Search and Rejection, for Inexperienced 20-cent search with one-price samples (Human Buyers)

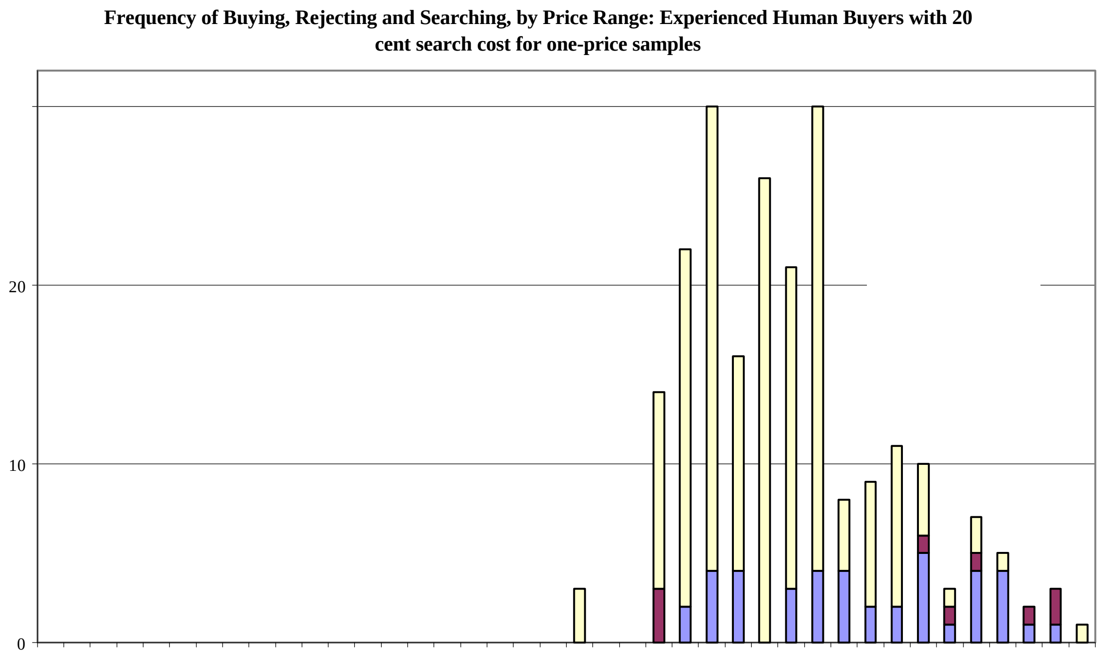

Figure 9: Distribution of Accepted Prices, and Prices Leading to Search and Rejection, for Experienced 20-cent search with one-price samples (Human Buyers)

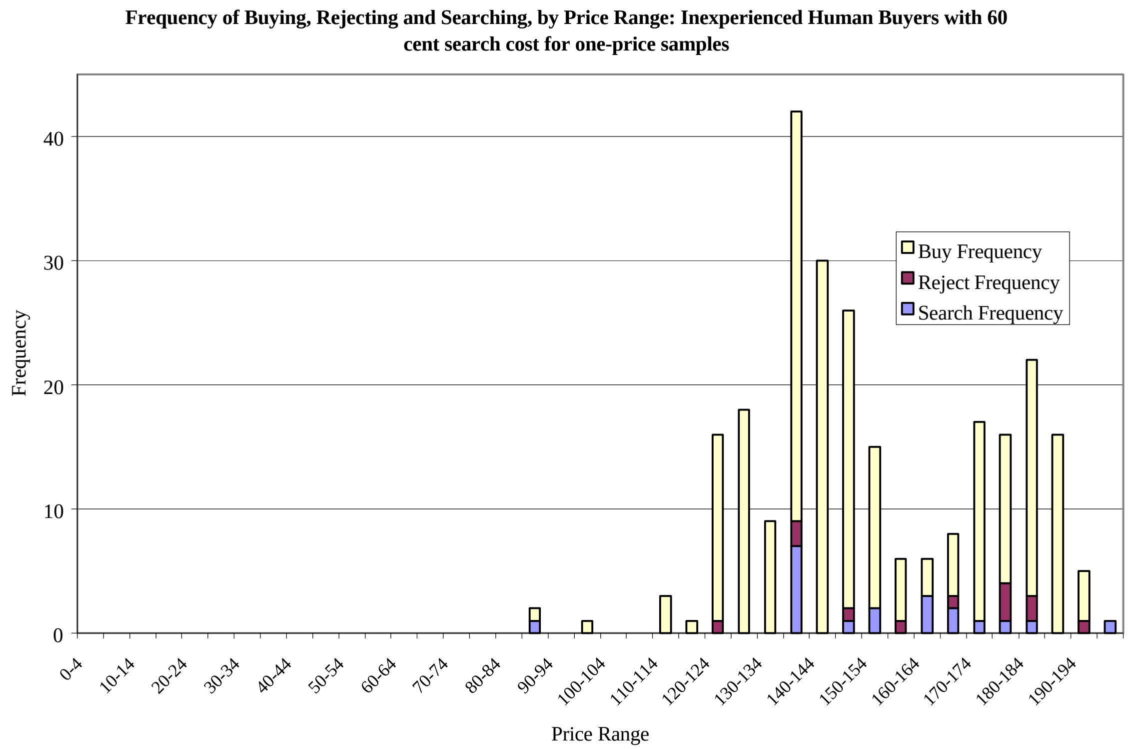

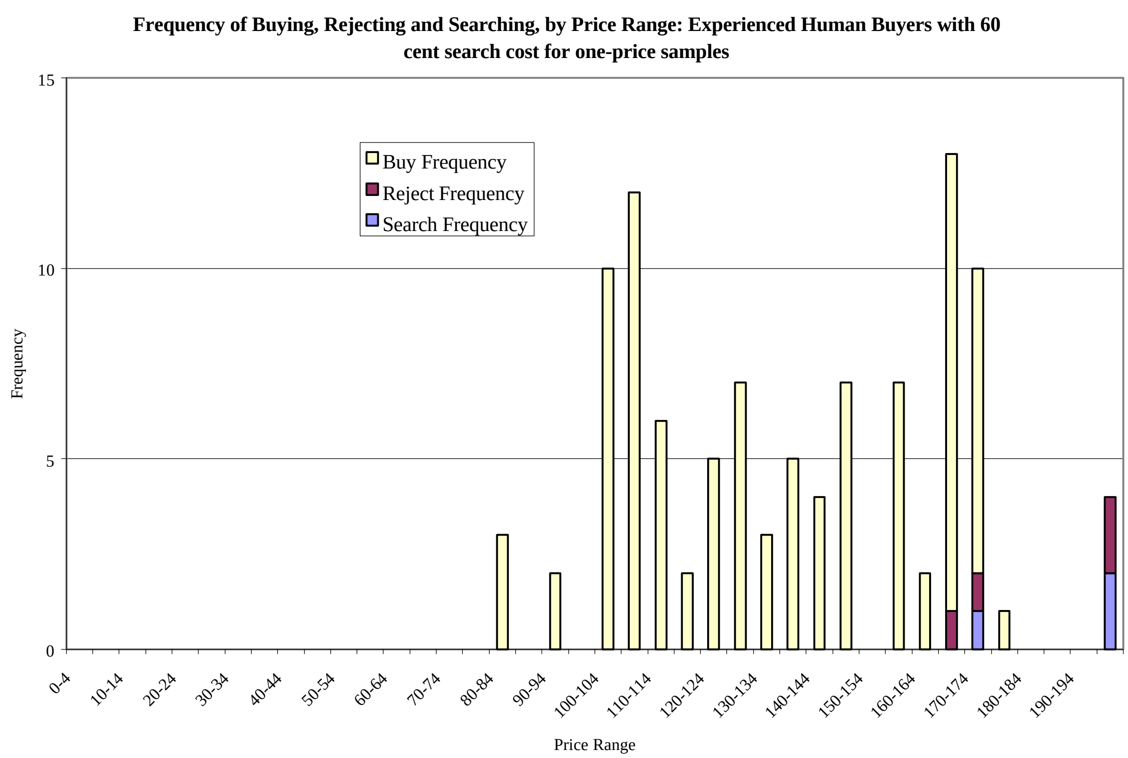

Figure 10: Distribution of Accepted Prices, and Prices Leading to Search and Rejection, for 60-cent search with one-price samples (Human Buyers)

Appendix A

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.

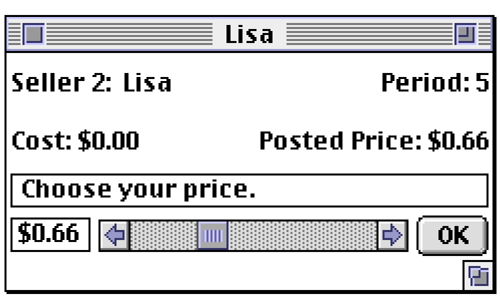

Figure 1

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.

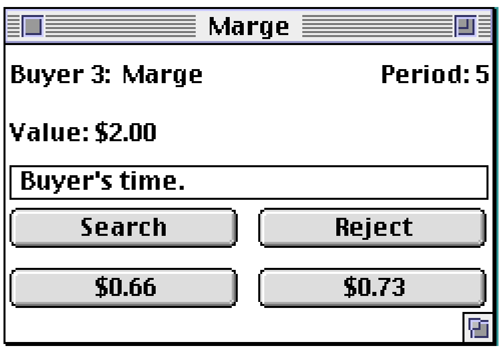

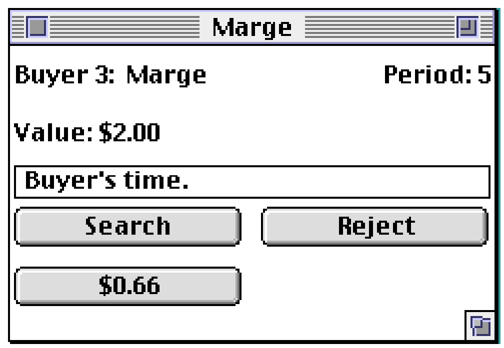

Figure 2a **Figure 2b**

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.

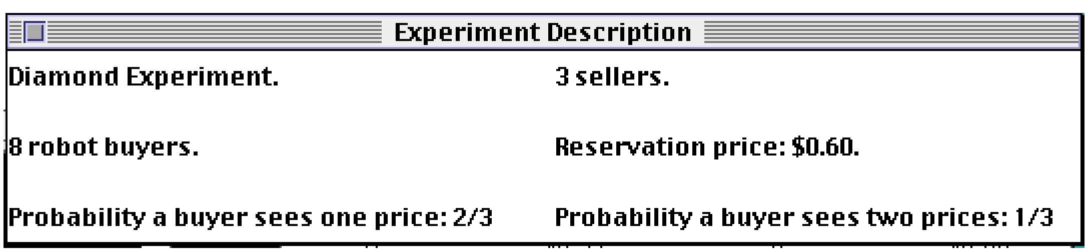

Figure 3

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.

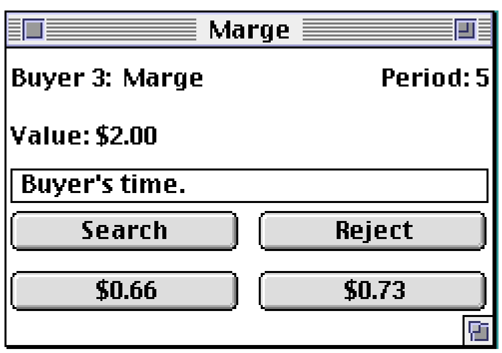

Figure 4

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

Continue

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

Continue

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.