Econometrica, Vol. 65, No. 2 (March, 1997), 311–345

PRICE FORMATION IN SINGLE CALL MARKETS

BY TIMOTHY N. CASON AND DANIEL FRIEDMAN[1]

This paper reports a laboratory experiment designed to examine the price formation process in a simple market institution, the single call market. The experiment features random values and costs each period, so each period generates a new price formation observation. Other design features are intended to enhance the predictive power of the Bayesian Nash equilibrium (BNE) theory developed recently for this trading institution. We find that the data support several qualitative implications of the BNE, but that subjects' bid and ask behavior is not as responsive to changes in the pricing rule as the BNE predictions. Bids and asks tend to reveal more of the underlying values and costs than predicted, particularly when subjects are experienced. Nevertheless, observed trading efficiency falls below the BNE prediction. The results offer more support for the BNE when subjects compete against Nash "robot" opponents. A simple learning model accounts for several of the deviations from BNE.

KEYWORDS: Bayesian Nash equilibrium, experiments, auctions, learning.

1. INTRODUCTION

ECONOMISTS HAVE ALWAYS PUT the price system at center stage but still lack of a clear understanding of the price formation process. The traditional process theory is tâtonnement, which assumes away the main problems by positing a benevolent auctioneer and numerous small (or nonstrategic) traders. In most markets, active traders are not especially numerous and seem to be quite willing to exploit their private information. Benevolent auctioneers at best are an endangered species. How then do traders come to agree on transaction prices?

Laboratory experimentalists have studied price formation for several decades, and noticed right away that the price formation process depends crucially on the market institution. Chamberlin (1948) found that transaction prices and quantities had little tendency to converge to competitive equilibrium (CE) in his market institution of bilateral search. By contrast, Smith (1962) found rapid and reliable convergence to CE with repetition of the continuous double auction (CDA) institution. The general question of price formation thus resolves into three research questions. What are the relevant market institutions? What are the equilibrium properties of such institutions? And to what extent do human traders come to approximate the equilibrium outcomes?

[1]Financial support was provided by the National Science Foundation (SBR-9223830 and SBR-9223461). We gratefully acknowledge the programming support of Andrew Davis and Tai Farmer and the research assistance provided by Carl Plat. We received helpful comments from John Kagel, Steve Williams, Mark Van Boening, and seminar participants at Texas A&M, the University of Houston, Indiana University, and the Economic Science Association Fall 1994 meetings. We are also very grateful to the co-editor and two anonymous referees, who helped us sharpen our results and strengthen the presentation. We retain responsibility for remaining errors and omissions.

The current paper is a laboratory study of the single call market (SCM) trading institution, also known as the Clearinghouse or the two-sided sealed auction. The SCM institution collects buyers' and sellers' offers (bids and asks respectively) while the market is open, and then clears the offers at a uniform price when the market closes. The SCM is perhaps the simplest viable market institution and has the best developed equilibrium theory. The SCM is used extensively on organized exchanges for securities with insufficient trading volume to support continuous trading, and it is used on the NYSE and elsewhere to set daily opening prices.

Our laboratory study features a random values environment. Unlike the vast majority of market experiments that use stationary repetition, we induce value and cost parameters for buyers and sellers that are drawn independently each trading period from announced (uniform) distributions. Thus we get a new observation of the price formation process each period, rather than a single observation spread across many periods.

Modern theoretical analysis assumes each trader is fully aware of the strategic value of her own private information and knows the structure of other traders' strategies. Such analysis seeks to characterize the Bayesian Nash equilibria (BNE) of the incomplete information game defined by the trading institution and the trading environment. Vickrey (1961) first used the BNE approach to analyze one-sided auctions. Chatterjee and Samuelson (1983) extended the BNE approach to analyze two-sided markets with a single buyer and a single seller. Wilson (1985) and Gresik and Satterthwaite (1989) extended it yet again to analyze SCM markets. The most complete results to date appear in a series of articles by Satterthwaite, Williams and Rustichini, surveyed in Satterthwaite and Williams (1993), and our experiment is informed primarily by that work.

Several previous laboratory studies are relevant. Smith et al. (1982) compares performance of the continuous double auction market institution (CDA) to several variants of the SCM in a simple stationary repetitive environment. Price formation was more rapid and reliable in the CDA but a multiple-unit, recontracting version of the SCM had equivalent allocational efficiency. Friedman (1993) makes similar comparisons in laboratory asset markets with new private information each period. Contrary to his prior expectations, he found that price formation is almost as reliable in the SCM as in the CDA.

The most relevant previous work is Kagel and Vogt (1993) and Kagel (1994), who also use a random values environment informed by the BNE theory. They find that the CDA institution is overall slightly more efficient than the SCM. They focus on the impact of trader numbers in the SCM, and report that efficiency increases with increasing numbers of traders, but much less rapidly than theory suggests. By contrast, our experiment focuses on variations in the trading institution and on the role of trader experience, given a constant moderate number of buyers and sellers. Our data analysis employs "payoff space metrics" and Selten and Buchta's (1994) learning direction theory as well as more standard techniques.

We begin in the next section with a brief summary of the BNE theory of the SCM, emphasizing its testable implications. In Section 3 we describe our laboratory procedures. The results are collected in Section 4. We find that the data support several qualitative implications of the BNE, but that subjects' bid and ask behavior is not as responsive to changes in the pricing rule as the theory predicts. Bids and asks tend to reveal more of the underlying values and costs than predicted, particularly when subjects are experienced. Nevertheless, observed trading efficiency falls below the BNE prediction. Although the theory's performance does not improve with increased experience, the results offer more support for the BNE when subjects compete against Nash "robot" opponents.

The reader might find the following interpretation useful in working through the results.[2] Some traders respond correctly to changes in the pricing rule, but the stronger tendency is for traders to underreveal their true values and costs at first, even when the pricing rule makes full revelation a dominant (but not entirely transparent) strategy. Over time, traders learn to reveal more fully. Some traders may learn to make approximate best responses to the empirical distribution of bids and asks. But most traders learn to reveal more than in BNE, perhaps because the learning process is biased. A simple learning model suggests that traders respond much more strongly to "missed" trades due to underrevelation than to adverse pricing due to overrevelation. Section 5 offers some other interpretations, conjectures, and suggestions for further work.

2. THEORY

The trading environment involves m buyers and m sellers. Buyer i's payoff (or profit or surplus) is $v_i - p$ if she purchases a single indivisible unit of the good at price p and is 0 otherwise (e.g., if she does not transact). Similarly seller j's payoff is $p - c_j$ if he sells a unit at price p and is 0 otherwise. The values v_i and costs c_j are privately drawn from known distributions. In our experiment $m = 4$ and the values and costs come from independent draws from the uniform distribution on $[0.00, 4.99]$.

The single call market (SCM) trading institution solicits a bid (or highest acceptable purchase price for a single unit) b_i from each buyer i and an ask (or lowest acceptable sale price) a_j from each seller j. The demand revealed in $\{b_i\}$ and the supply revealed in $\{a_j\}$ then are cleared at a uniform equilibrium price p^*. With indivisible units, there often is an interval $[p_l, p_u]$ of market clearing prices, in which case the chosen price is $(1-k)p_l + kp_u$ where k in $[0,1]$ is a specified parameter.[3]

[2] This interpretation owes much to the suggestions of John Kagel, a co-editor, and the anonymous referees.

[3] Satterthwaite and Williams (1993) express p^* more compactly as $(1-k)s(m) + ks(m+1)$ where $s(i)$ is the ith smallest offer in the combined set $\{b_i\} \cup \{a_j\}$.

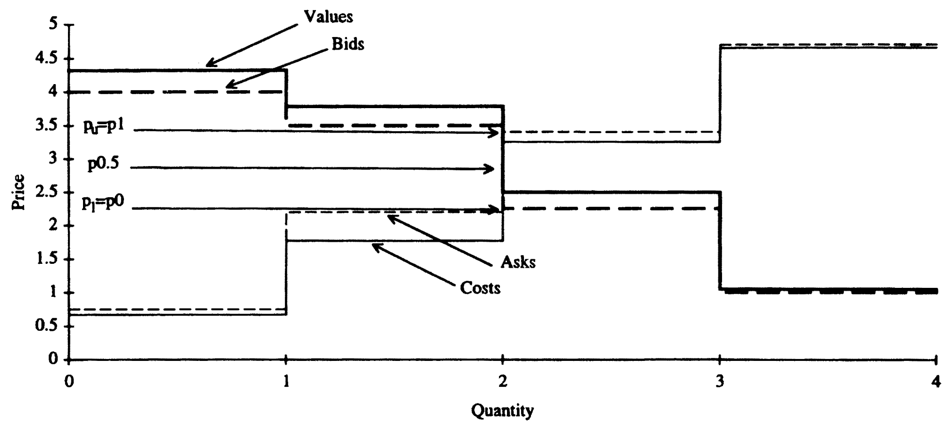

FIGURE 1.—Example price determination for three price rules.

Figure 1 illustrates the pricing rules for a specific draw of values and costs, shown with solid lines. Traders reveal demand and supply through their bids and asks, shown with dotted lines. The interval of market clearing prices is $[p_l, p_u]$. In the experiment we examine the three pricing rules denoted $p1$, $p0.5$, and $p0$ in Figure 1. These prices correspond respectively to the upper end ($k = 1.0$), midpoint ($k = 0.5$), and lower end ($k = 0$) of the market clearing price interval.

To what extent should traders reveal their true values and costs? In general, each buyer optimally reduces her bid below value (and each seller increases his ask above cost) to the point that (i) the marginal loss from the reduced probability of transacting just matches (ii) the marginal gain conditional on transacting. Satterthwaite and Williams (1989) observe that effect (ii) is absent in the SCM for sellers when $k = 1$, because then p^* is always set by a buyer or by a nontransacting seller. The same holds for buyers when $k = 0$. Hence full revelation in these cases ($a_j = c_j$ when $k = 1$ and $b_i = v_i$ when $k = 0$) is a dominant strategy, analogous to the dominant truth-telling incentives in the one-sided second-price (Vickrey) auction.

Other theoretical predictions arise from the Bayesian Nash equilibrium (BNE) of the model. Figure 2 illustrates predicted risk neutral bid and ask functions and suggests five qualitative hypotheses for the pricing rule treatments $k = 0$, $k = 0.5$, and $k = 1$.[4] The first two hypotheses concern the qualitative impact of k on the bid functions $B_k(v)$ and $A_k(c)$.

[4] For $k = 0$ and $k = 1$ the functions graphed in Figure 2 are the unique smooth BNE (Williams (1991)). For $k = 0.5$ the structure of BNE is more complex. Rustichini et al. (1994) compute numerically a family of asymmetric smooth equilibria that is entirely contained in an open epsilon neighborhood of the symmetric bid and ask functions graphed in Figure 2 for $k = 0.5$. [Eyeball interpolation of Figure 5 in Satterthwaite and Williams (1993) suggests that epsilon is less than 5 cents for our parameterization.] The hypotheses listed below collapse the family of asymmetric equilibria to the symmetric functions graphed in our Figure 2. This greatly simplifies the exposition and analysis, and should not affect the conclusions since epsilon = 5 cents is negligible relative to the behavioral variability in the data. Appendix A contains a separate justification for Hypotheses 1 and 2.

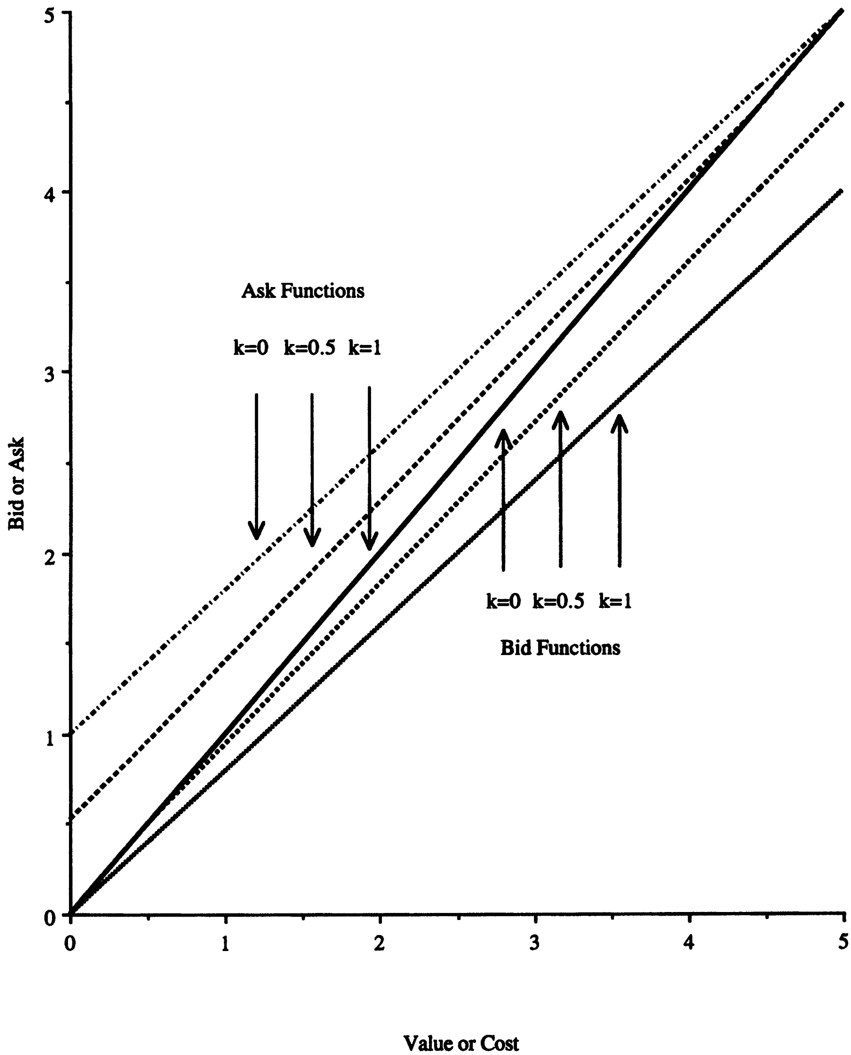

FIGURE 2.—Approximate risk neutral BNE bid and ask functions for the three k treatments (4 buyers, 4 sellers).

HYPOTHESIS 1 (Shifting Bid Functions): *For given* $v >$ \$0.52, *actual bid prices are decreasing in* k, *i.e.*, $B_0(v) > B_{0.5}(v) > B_1(v)$.

HYPOTHESIS 2 (Shifting Ask Functions): *For given* $c <$ \$4.47, *actual ask prices are decreasing in* k, *i.e.*, $A_0(c) > A_{0.5}(c) > A_1(c)$.

For $k = 0.5$ the equilibrium bid function is $B(v_i) = v_i$ for $v_i \leq 0.52$, and the equilibrium ask function is $A(c_j) = c_j$ for $c_j \geq 4.47$, so the strict inequalities in

Hypotheses 1 and 2 for $k = 0.5$ do not apply for these very low values and very high costs.

For the remaining three hypotheses, it is convenient to use ratios that summarize the degree of value (and cost) revelation in actual bid b_i (or ask a_j) given the pricing rule k and the actual value v_i (or cost c_j): $VUR_i^k \equiv (v_i - b_i)/v_i$ is the Value Underrevelation Ratio and $CUR_j^k \equiv (a_j - c_j)/(4.99 - c_j)$ is the Cost Underrevelation Ratio. These ratios represent the fraction of value discounted in a bid, and the cost "mark-up" relative to the highest possible cost of \$4.99. These ratios are constant for risk neutral subjects in the BNE of the $k = 0$ and $k = 1$ treatments because the BNE bid and ask functions are linear, and are approximately constant in the $k = 0.5$ treatment.

HYPOTHESIS 3 (Symmetric Strategies for $k = 0.5$): *$\{VUR_i^{0.5}\}$ and $\{CUR_j^{0.5}\}$ have the same distribution.*

HYPOTHESIS 4 (Symmetric Strategies for Bids in $k = 0$ and Asks for $k = 1$): *$\{VUR_i^0\}$ and $\{CUR_j^1\}$ have the same distribution.*

HYPOTHESIS 5 (Symmetric Strategies for Bids in $k = 1$ and Asks for $k = 0$): *$\{VUR_i^1\}$ and $\{CUR_j^0\}$ have the same distribution.*

Hypotheses 3–5 concern the symmetry of the predicted bid and ask functions, arising from the symmetric decision problems for buyers and sellers in this trading institution. Full revelation is the dominant strategy for $k = 0$ buyers and $k = 1$ sellers; Hypothesis 5 is the weaker assertion that revelation, whether full or partial, will be the same in these conditions. The BNE bid and ask functions reveal least for $k = 0$ sellers and $k = 1$ buyers; Hypothesis 4 is the weaker assertion that the degree of underrevelation will be the same in these conditions. Hypothesis 3 for the $k = 0.5$ pricing rule is that the actual bid function and ask function will exhibit the same degree of underrevelation, whether or not that degree is the same as in the symmetric functions in Figure 2.

Two qualifications are in order, regarding risk aversion and regarding learning. The predictions for buyers' bids when $k = 0$ and for sellers' asks when $k = 1$ are based on dominant strategies and therefore hold regardless of trader risk preferences. The other functions in Figure 2 are based on the risk neutral BNE. Bids above and asks below these functions may be consistent with risk aversion and BNE (Rustichini, Satterthwaite and Williams (1990)). Since risk aversion may play a role, the qualitative Hypotheses 1–5 are useful supplements to tests of the risk neutral functions graphed in Figure 2. As long as traders' risk preferences influence behavior symmetrically for the buyer and seller roles (as implied by the symmetry of the institution), the distribution of the underrevelation ratios should be equal because subjects were assigned randomly to each role.

Satterthwaite and Williams (1993) point out that traders might have difficulty in learning BNE strategies: subjects must learn to best-respond against opponents who are also learning and revising strategies, so opponents' strategies may be noisy and unstable. Moreover, subjects each period observe only one point on opponents' bid and ask functions; they never observe the whole functions. Following a suggestion by Satterthwaite and Williams (1993), we employ an experimental treatment to mitigate the learning problems. In some sessions each subject competes against 7 "robot" opponents that are programmed to play equilibrium strategies. So in this treatment the subject faces a stationary problem which may aid learning.

Appendix A collects some mathematical details for this BNE derivation to make the present paper more self-contained, and includes a couple of minor theoretical extensions. In particular, we show that the bid and ask ordering summarized in Hypotheses 1 and 2 holds for a range of trader beliefs, not just BNE beliefs. We also derive best responses to fully revealing (or, "truth-telling") opponent strategies, which we use in Section 4.3 to evaluate a *mimicking hypothesis* for the robot opponent treatment.

A final methodological comment may be useful before presenting our experiment. The BNE theory we examine is internally consistent and prescribes self-interested behavior that seems well within ordinary human capacities. In principle it should be possible to find some laboratory environment, some subject pool, training procedures, and experimental design that will produce behavior closely approximating the BNE predictions. But our experimental procedures are not the result of a long search to find conditions that are sufficiently favorable to verify the BNE theory. Rather we begin with standard laboratory procedures and modify them in a few ways that we believe a priori will be favorable to the theory; e.g., using random values and costs and sometimes using robot traders. Our goal is not definitively to accept the theory (a Popperian impossibility) nor to reject it (unlikely given the absence of sophisticated and explicit alternatives). Rather, our main goal is to assess the strengths and weaknesses of the current best theory. Failure to reject a prediction signals a strength of the theory in organizing behavioral data. Rejection of a BNE prediction for our laboratory SCM data provides a clue on where the predictive powers of the theory might be improved.

3. EXPERIMENTAL DESIGN

The experiment consists of 19 separate laboratory sessions conducted at UCSC, each with 30 or 40 trading periods. Subjects were recruited from large lower division classes in economics and biology. No inexperienced subject had ever participated in a previous SCM session. Subjects were randomly assigned a computer and trader position, and instructions were read orally while subjects followed along on their own copies. Four practice periods preceded the 30 or 40 trading periods. Including instructions, sessions lasted a little less than two hours. Total earnings ranged between \$5 and \$30 per subject with an average of about \$18.

Each session employed 8 traders (with one exception noted in Table I below); the 4 buyers and 4 sellers could enter offers in each period. Values and costs were induced in the standard fashion: each period t each buyer i received a specified "resale value" v_{it} for a single indivisible unit, and similarly each seller

j received a specified cost c_{jt}. If these traders transact at price p, then the exchange surplus $v_{it} - c_{jt}$ is composed of the buyer's profit $v_{it} - p$ and the seller's profit $p - c_{jt}$. These profits accumulated in a computer account throughout the session, and were paid in cash after the last trading period.

At the beginning of the session all traders were informed that all values and costs would be drawn independently each period from the discrete (rounded to the nearest penny) uniform distribution over the range [\$0.00, \$4.99]. Before the start of each trading period, each buyer saw her own value for that period and each seller saw his own cost, but no trader saw others' realized values or costs. At the conclusion of the trading period each trader observed the bids, asks, values, costs, and profits of all traders. We provided this complete ex post information to increase traders' opportunity to learn about their rivals' strategies. The instructions (available on request) provide additional details of the procedures and the operation of the market program.

The experiment employs two treatment variables—the pricing rule and trader experience. Recall that in the SCM each buyer submits a bid knowing only his own value, and each seller submits an ask knowing only her own cost. The computer aggregates (or "crosses") the revealed supply (asks) and demand (bids) and calculates the interval of market-clearing prices. The pricing rule k determines which point in the interval is used. For $k = 0$, 1, and 0.5 the market clearing price is respectively the upper endpoint, lower endpoint, and midpoint of the interval.

Trader experience is a composite treatment variable with five levels, including the usual two levels of inexperienced human opponents (8 inexperienced human traders) and experienced human opponents (8 human traders who previously had participated in an inexperienced human trader session). Following standard convention, throughout the data analysis we refer to these two levels as inexperienced and experienced. The next two levels involve robot traders (i.e., computer algorithms) programmed to use the BNE bid or ask functions graphed in Figure 2; instructions include the relevant graph. In the Nash robots treatment each human is inexperienced in any SCM market and faces 7 such robots. Strictly speaking, Nash robots is an individual choice task intended to overcome the learning difficulties cited above. The next treatment, which we shall refer to as Nash experienced, brings together (in a true market environment) 8 humans who had previously participated in Nash robots sessions. For reasons discussed below in Section 4.3, we included a fifth treatment called revealing robots that was identical to the Nash robots treatment except that the robot buyers entered bids equal to value, and robot sellers entered asks equal to cost.

A final aspect of the experience treatment deserves mention. All inexperienced sessions (robot and human) used the same set of value and cost draws, and all experienced (robot and human) sessions used a second (different) set of value and cost draws. This design feature eliminates a possible source of variability across sessions and enables us to make pairwise comparisons across sessions with differing price rules (k).

Table I summarizes the 19 sessions. The 13 inexperienced sessions had 30 periods and the 6 experienced sessions had 40 periods. The eight traders switched trading roles twice within a session, so that each could obtain experience facing the incentives of both sides of the market. In the 30-period inexperienced sessions, roles were switched before period 9 and before period 25. In the 40-period experienced sessions, roles were switched before period 11 and before period 31. The pricing rule k, the exact number of buyers and sellers, their human or robot status, and the possibility of role switches all were announced publicly each session before trade began.

4. RESULTS

We report the results in six subsections. Section 4.1 presents market level data such as overall trading efficiency and prices, comparing them to the BNE predictions and to the competitive equilibrium. Section 4.2 describes several features of the bid and ask data that are employed in the subsequent analyses. Section 4.3 presents tests of the comparative static predictions of the BNE for the pricing rule k (Hypotheses 1–2). Section 4.4 assesses the precise bid and ask predictions of the risk neutral BNE, and briefly relates the results to previous laboratory studies of one-sided auctions. Section 4.5 evaluates the symmetry properties of the bid and ask behavior implied by the BNE (Hypotheses 3–5)

TABLE I

SUMMARY OF LABORATORY SESSIONS

Session Name	k Treatment	Experience Treatment: Opponents	Experience	Label	Number of Periods
k0-hum-1	$k = 0$	Humans	None	Inexperienced	30
k0-hum-2	$k = 0$	Humans	None	Inexperienced	30
k0-hum-3x	$k = 0$	Humans	vs. Humans	Experienced	40
k0-rob-4	$k = 0$	Nash Robots	None	Nash Robots	30
k0-hum-5rx	$k = 0$	Humans	vs. Robots	Nash Experienced	40
k5-hum-6	$k = 0.5$	Humans	None	Inexperienced	30
k5-hum-7	$k = 0.5$	Humans	None	Inexperienced	30
k5-hum-8x	$k = 0.5$	Humans	vs. Humans	Experienced	40
k5-rob-9	$k = 0.5$	Nash Robots	None	Nash Robots	30
k5-hum-10rx	$k = 0.5$	Humans	vs. Robots	Nash Experienced	40
k1-hum-11	$k = 1.0$	Humans	None	Inexperienced	30
k1-hum-12	$k = 1.0$	Humans	None	Inexperienced	30
k1-hum-13x	$k = 1.0$	Humans	vs. Humans	Experienced	40
k1-rob-14[a]	$k = 1.0$	Nash Robots	None	Nash Robots	30
k1-rob-15	$k = 1.0$	Nash Robots	None	Nash Robots	30
k1-rob-16	$k = 1.0$	Nash Robots	None	Nash Robots	30
k1-hum-17rx	$k = 1.0$	Humans	vs. Robots	Nash Experienced	40
k0-rob-18	$k = 0$	Revealing Robots	None	Revealing Robots	30
k1-rob-19	$k = 1.0$	Revealing Robots	None	Revealing Robots	30

Notes: All markets involved 4 buyers and 4 sellers each period, whose values and costs were drawn independently from the uniform distribution over [0, \$4.99].

[a]Session k1-rob-14 employed only 5 subjects, each competing against 7 robot opponents. All other sessions employed 8 human subjects.

TABLE II

TRADING EFFICIENCY, BY TREATMENT CONDITION

Session or Benchmark	$k = 0$	$k = 0.5$	$k = 1.0$
Competitive Equilibrium	100	100	100
Risk Neutral BNE	98.0	98.9	98.6
Zero Intelligence Traders [Mean]	33.1	33.0	33.6
[5th percentile, 95th percentile]	[24.5, 42.5]	[23.7, 42.6]	[25.1, 43.0]
Inexperienced Mean	83.3	87.3	83.9
(Std. Error across periods)	(2.5)	(2.7)	(3.8)
Experienced Mean	92.4	91.4	85.0
(Std. Error across periods)	(1.8)	(11.0)	(3.4)
Nash Experienced Mean	96.6	92.2	91.9
(Std. Error across periods)	(1.4)	(2.6)	(3.4)

Notes: All efficiency values are the percentage of the maximum gains from exchange realized by traders. Zero intelligence estimates are based on a simulation of 200 sessions for each k treatment.

and the overall level of value and cost revelation in traders' offers. Section 4.6 concludes with a brief exploration of trader learning.

4.1. *Market Performance*

Trading efficiency is defined as the percentage of the potential gains from exchange actually extracted by traders. Table II presents the realized trading efficiency along with the risk neutral BNE efficiency prediction for each k treatment and for each relevant experience treatment.[5] The BNE sets a very high efficiency benchmark, between 98 and 99 percent. In every session, the traders fall short of that level. The inexperienced sessions' trading efficiency ranges between 83 and 88 percent, and with one exception the experienced sessions' efficiency ranges between 91 and 97 percent. Efficiency does not vary systematically with the k treatment.

Table II also presents an efficiency comparison based on a simulation of completely nonstrategic "zero intelligence" trading behavior. Simulated buyers' bids are randomly and uniformly distributed between 0 and their resale value; analogously, simulated sellers' asks are uniformly distributed between their cost and the highest possible cost draw ($4.99). This simulation exercise is not intended to be a serious representation of trader strategies, but rather as a second benchmark, polar to the BNE. Gode and Sunder (1993) have shown that such simple strategies lead to nearly fully efficient outcomes in the continuous double auction, often exceeding the efficiency of markets with human traders. In contrast, Table II demonstrates that zero intelligence trading behavior leads to average efficiency of about 33 percent in the SCM. Realized efficiency in each session exceeds the 95th percentile of a sample of 200 zero intelligence simula-

[5] Tables II and III do not report the Nash robots or the revealing robots data because the 7 robots in these sessions strongly bias overall performance toward the BNE or toward the competitive equilibrium.

tions, so we strongly reject the zero intelligence benchmark for market efficiency.

Table III compares realized trading prices to two theoretical predictions: the competitive equilibrium (CE) price interval in Panel A (which is based on full revelation of values and costs), and the risk neutral BNE point prediction in Panel B. First note in Panel A that the theoretical BNE price lies within the CE interval in about 80 to 85 percent of the periods. The lower half of the table indicates that observed prices are within the CE price interval less frequently than this BNE benchmark, but much more frequently than the zero intelligence benchmark. In the inexperienced sessions, prices are within the CE price interval in about 40 percent of the periods. Prices are within the CE interval in more than one-half of the periods in the experienced sessions. Overall it appears that CE is not a particularly good predictor of prices in this random values

TABLE III

PRICE COMPARISONS RELATIVE TO COMPETITIVE EQUILIBRIUM INTERVAL AND RISK NEUTRAL BNE

Panel A: Percentage of Prices within the Competitive Equilibrium Price Interval

Session or Benchmark	$k = 0$	$k = 0.5$	$k = 1.0$
Competitive Equilibrium	100	100	100
Risk Neutral BNE	85.3	79.4	86.8
Zero Intelligence Traders, Mean	17.3	19.1	16.7
[5th percentile, 95th percentile]	[10.0, 26.7]	[11.3, 27.5]	[8.8, 25.0]
Inexperienced	46.7	45.0	36.7
[Two Individual Sessions]	[46.7, 46.7]	[33.3, 56.7]	[33.3, 40.0]
Experienced	52.6	57.9	50.0
Nash Experienced	50.0	71.1	52.6

Panel B: Mean Difference [Observed-BNE Price] and Mean Absolute Deviations from the Risk Neutral BNE Price

	$k = 0$		$k = 0.5$		$k = 1.0$	
Session or Benchmark	Mean Diff.	Absolute Dev.	Mean Diff.	Absolute Dev.	Mean Diff.	Absolute Dev.
Risk Neutral BNE	0	0	0	0	0	0
Zero Intelligence Traders, Mean	−0.26	0.63	0.01	0.51	0.31	0.63
[5th percentile, 95th percentile]	[−0.57, 0.08]	[0.44, 0.82]	[−0.24, 0.32]	[0.35, 0.67]	[0.01, 0.63]	[0.43, 0.86]
Inexperienced Mean	−0.17	0.43	0.11	0.32	0.30	0.48
(Std. Error across periods)	(0.07)	(0.04)	(0.05)	(0.03)	(0.07)	(0.05)
Experienced Mean	−0.22	0.30	0.01	0.15	0.30	0.33
(Std. Error across periods)	(0.05)	(0.03)	(0.03)	(0.02)	(0.05)	(0.04)
Nash Experienced Mean	−0.34	0.34	−0.06	0.18	0.25	0.32
(Std. Error across periods)	(0.05)	(0.05)	(0.05)	(0.04)	(0.05)	(0.04)

Notes: Zero intelligence estimates are based on a simulation of 200 sessions for each k treatment.

environment, and its performance improves only slightly as subjects gain experience.

Panel B of Table III indicates that the mean absolute price deviation from the BNE point prediction ranges between 15 to 48 cents, falling somewhat with experience. These deviations are generally significantly smaller than the absolute deviations in the zero intelligence simulations. Average prices are always greater than the BNE prediction in the $k = 1$ sessions, and less than the BNE prediction in the $k = 0$ sessions. The same is true for prices in the zero intelligence simulations. A possible explanation is that human traders do not fully respond to the price rule treatment (certainly zero intelligence traders do not). Recall from Figure 2 that BNE bid and ask functions decrease in k. Hence expected BNE prices also decrease in k, and that could account for the observed pattern of price deviations from BNE. We now examine the bid and ask data more directly to deepen our understanding of market performance.

4.2. *Bid and Ask Behavior*: *Some Preliminaries*

Figure 3 illustrates the nature of the bid and ask data by showing a scatterplot for session k5-hum-8x, a fairly typical $k = 0.5$ session with experienced traders. The solid line in each panel represents the full revelation benchmark bid = value or ask = cost, and the dotted line indicates the (approximate) risk neutral BNE bid or ask function. The open circles represent the 320 actual bids and asks by 8 traders over 40 periods.

In Figure 3 (and in all other sessions) there is a strong positive correlation between values and bids (and between costs and asks), and few traders overreveal by bidding above value or asking below cost. Nevertheless, there exists substantial variation in the bids and asks relative to the simple BNE bid and ask functions.

Variation seems perhaps more pronounced for very high and very low values and costs. In such cases traders have little incentive to use (and little opportu-

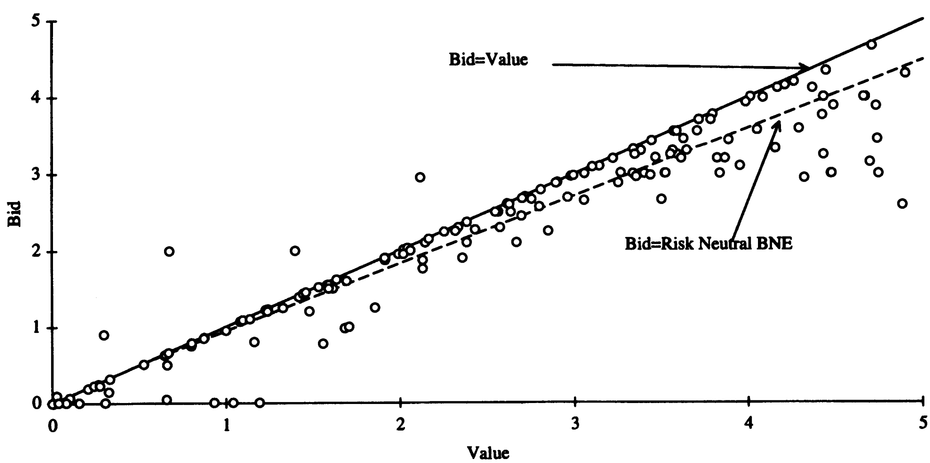

FIGURE 3A.—Bids in session k5-hum-8x ($k = 0.5$, Experienced Subjects).

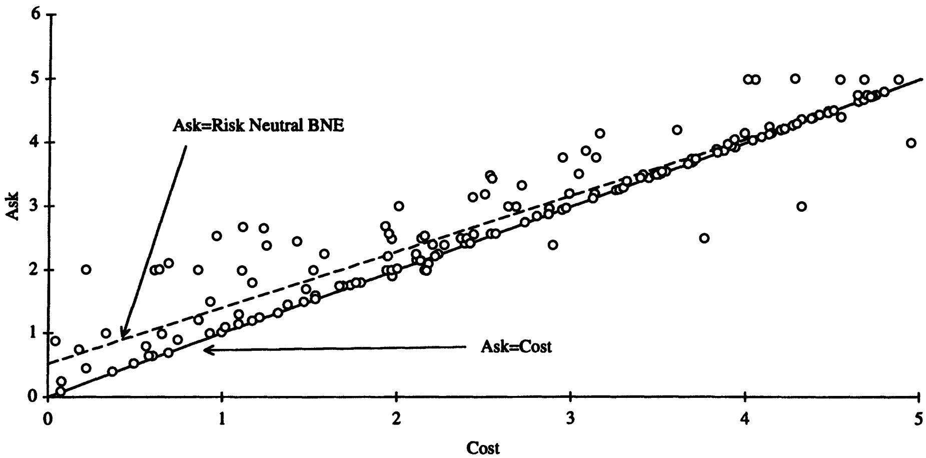

FIGURE 3B.—Asks in session k5-hum-8x ($k = 0.5$, Experienced Subjects).

nity to learn) the BNE bid and ask functions since the expected profit impact is so slight. For example, consider a buyer with value \$4.50 in a $k = 0.5$ session. She is very likely to transact at a price set by other traders whether she bids 80 percent, 95 percent, or even 105 percent of value instead of 89.7 percent of value as called for by her BNE bid function. For the same reason, low cost sellers' expected profits are insensitive to their asks. At the other extreme, high cost sellers and low value buyers are unlikely to transact even with a substantial deviation from the BNE ask or bid, so their expected profit is at best near 0. (Of course, they can expect to earn large negative profits if they bid far above value or ask far below cost.) The choices with the greatest payoff consequences are the bids and asks for values and costs near likely CE prices, because here the probability of transacting and the expected transaction price depend sensitively on the degree of revelation. Of course, BNE theory treats these mid-range and extreme draws alike, in that the bid and ask functions everywhere balance infinitesimal differences in expected payoffs.

Several authors recently have argued that it is useful to use a "payoff space metric" in assessing behavioral regularities, i.e., to weight observations by their expected payoff consequences—see Smith and Walker (1993); Harrison (1989); Friedman (1992); and Cox and Oaxaca (1995). Figure 4 illustrates the weights we have calculated for SCM buyers and sellers. We estimate the relative importance of each value or cost draw in terms of the expected loss from deviating unilaterally from BNE. Note that the most important values, representing about two-thirds of total expected payoff consequences, lie in the interval [\$2.00, 3.50] while the analogous interval of important costs is [\$1.50, 3.00]. In some of the analyses below we weight each observation by its importance, and in others we exclude the data outside the intervals of important values and costs.[6]

[6] We also conducted all the analyses with unweighted observations. The inferences drawn from the weighted data are often sharper but are never inconsistent with those drawn from unweighted data. Details of the weight calculations are available upon request.

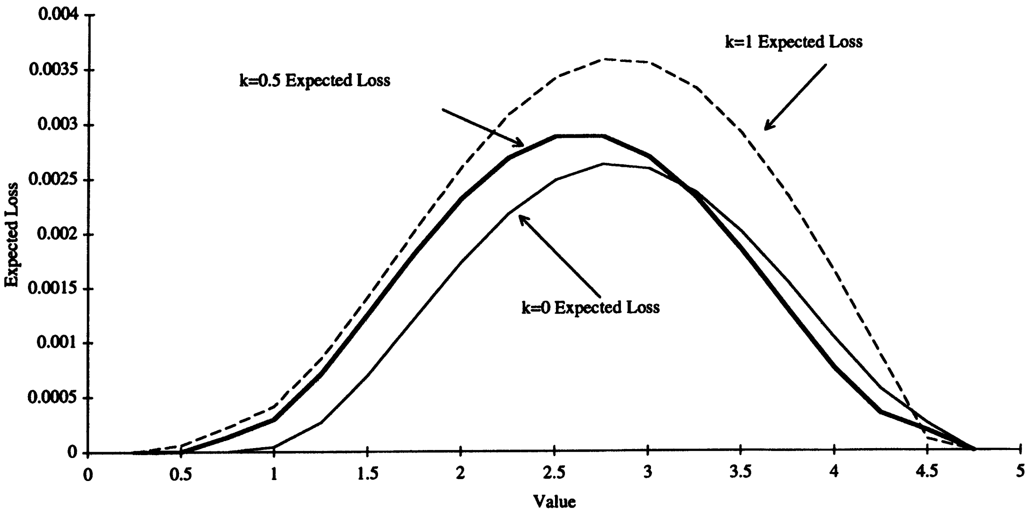

FIGURE 4.—Expected payoff importance weights based on a ten-cent bid deviation.

Our data strongly support the conclusion that bidding (and asking) behavior is noisier when less expected profit is at stake. Define "behavioral noise" as the absolute value of the deviation from the estimated linear bid or ask function, where this estimate is based on all the data (unweighted) and is calculated separately for each experience and k treatment. This noise measure is negatively correlated with the importance weights in Figure 4 for all 8 data sets (4 experience treatments $\times$ 2 trader roles). Six of the 8 estimated correlation coefficients are significantly different from zero at the 5 percent level (two-tailed tests). The same conclusion also holds for other definitions of behavioral noise based on different summary statistics (such as the underrevelation ratios *VUR* and *CUR*) and other measures of choice variance.[7]

One other general feature of the data analysis requires mention. Because the same subject made a number of bids and asks, the independence assumption required for many standard tests is untenable. Therefore, all significance levels and estimates in the parametric analyses such as Tables IV and VII are based on a random effects error specification to capture systematic differences across subjects. In this specification, the error term associated with each observation has the form $e_{it} = u_i + \varepsilon_{it}$, where u_i is a random subject effect and ε_{it} is a standard i.i.d. error term. In our data the coefficient estimates change little when using this error structure compared to the usual error structure $e_{it} = \varepsilon_{it}$; however, standard errors often increase substantially and an LM test typically rejects homogeneity across subjects. At times we also employ *within-subject* tests that compare the behavior of each subject in the buyer role with the same

[7]In contrast, there is no evidence that behavioral noise decreases for later periods within a session. Kagel and Roth (1992) report a similar result for one-sided first-price private value auctions.

TABLE IV

PAIRED DIFFERENCES IN OFFERS FOR k TREATMENTS: OBSERVATIONS WEIGHTED BY PAYOFF IMPORTANCE

	Session Group	Mean Bid Difference			Mean Ask Difference		
		$k=0$ bid− $k=0.5$ bid	$k=0.5$ bid− $k=1$ bid	$k=0$ bid− $k=1$ bid	$k=0$ ask− $k=0.5$ ask	$k=0.5$ ask− $k=1$ ask	$k=0$ ask− $k=1$ ask
(1)	Risk Neutral BNE	0.26	0.29	0.55	0.29	0.26	0.55
(2)	Inexperienced	−0.12	−0.06	−0.19††	−0.12	0.12*	−0.01
		(0.034)	(0.045)	(0.044)	(0.041)	(0.050)	(0.042)
(3)	Experienced	−0.06	0.04	−0.02	0.02	−0.11†	−0.09
		(0.033)	(0.035)	(0.031)	(0.044)	(0.055)	(0.036)
(4)	Nash Robots	−0.01	0.19*	0.18*	0.11	0.01	0.13
		(0.047)	(0.042)	(0.041)	(0.036)	(0.038)	(0.045)
(5)	Nash Experienced	−0.05†	0.09	0.04	−0.09†	0.003	−0.08
		(0.031)	(0.037)	(0.031)	(0.035)	(0.032)	(0.040)

Notes: Means (and standard deviations in parentheses) are calculated using the expected loss weights in Figure 4. Significance levels are based on the (individual subjects) random effects error specification. Hypotheses 1 and 2 imply positive mean differences, and *, **, and *** respectively denote significance at the one-tailed 10, 5, and 1 percent levels for positive mean differences. The same significance levels for two-tailed tests are denoted †, ††, and †††, and are applied to negative mean differences.

subject's behavior in the seller role. Like the random effects error structure, this within-subject approach controls for systematic differences across subjects due to personal characteristics such as risk preferences.

The random effects specification combines nicely with the experimental design feature that the value and cost sequences are constant across k treatments, resulting in fairly powerful tests. The within-subjects tests generally are less powerful because of the small number of observations for any individual subject, so the analysis focuses on pooled data with the random-effects error correction.

4.3. *Bid and Ask Behavior*: *Hypotheses* 1 *and* 2

Table IV reports a direct test of Hypotheses 1 and 2. It pairs the identical value or cost draws from the differing k treatments, and presents the mean difference in offers for each pair. The first row shows that in BNE the mean shift in bids and asks is 26 or 29 cents in moving from $k=0$ to $k=0.5$ or from $k=0.5$ to $k=1$, and is 55 cents in the overall move from $k=0$ to $k=1$. Traders' actual mean bids and asks never shift this much. Indeed, in the inexperienced and experienced treatments (lines 2 and 3 of the table) the shifts often are in the wrong direction and occasionally are statistically significant.[8] The Nash robots treatment produces shifts in closest conformity to the BNE

[8] Our research hypothesis is that the differences are positive, so the significance levels reported for positive mean differences are based on one-sided tests. For negative mean differences we report two-sided significance levels, testing the null hypothesis that the difference is zero against the alternative that the difference is not equal to zero.

predictions. Even here the mean shifts are at best 19 cents or 11 cents instead of 29 cents, and at best the shift is significant at 10 percent level. Overall, the tests show that human traders, even when experienced, are surprisingly insensitive to the trading rule k.[9]

We have confirmed these results using underrevelation ratios in within subject tests for the $k = 1$ and $k = 0$ sessions. For example, in a $k = 1$ session a subject's *VUR* in the buyer role should exceed her *CUR* in the seller role. Paralleling the results of Table IV, the number of subjects with an ordering of median *VUR* and *CUR* that is consistent with the BNE is highest in the Nash robots treatment (19 of the 29 subjects). The ordering of the two medians is consistent with the BNE for about one-half of the subjects in the other experience treatments.[10,11]

Table IV indicates that the predicted k-treatment effect receives the greatest support in the Nash robots condition. One explanation is that this treatment

[9] To complement the matched-pair tests shown in Table IV we conducted Kolmogorov-Smirnov (K-S) distribution tests that the offer distributions were different in the different k treatments. These tests appear to have low power in this setting so they rarely reject the null hypothesis that the bids or asks in the different treatments are equal. This highlights the advantage of the matched-pair tests. For the few cases that the null is rejected in the K-S distribution tests, it is also rejected by the matched-pair tests. We also performed these tests using the subset of data from the final one-half of each session. If subjects require extensive experience to learn the market incentives, then support for Hypotheses 1 and 2 may only be present for the later periods of the experienced sessions. Results are similar in all cases, however, so these other data sets are not reported here.

[10] Traders participate in two roles within each session, so the null hypothesis of random behavior (or insensitivity to the pricing rule) would predict that one-half of the subjects' medians would be ordered according to the BNE prediction. One possible artifact of this trader role switchover within a session is a hysteresis effect if subjects try to apply their experience in one role improperly to the new strategic conditions of their new role. We examined the periods after the buyer and seller switch for "anchoring" of strategies to look for evidence supporting this conjecture. Define an "early" period as the first 4 periods of the session, as well as the first 4 periods after a buyer and seller switch. There is some weak evidence that subjects' underrevelation ratios move toward the BNE prediction in the later periods compared to these early periods. In the inexperienced and experienced conditions, the shift is in the direction of the BNE in 7 of the 8 cases (2 k treatments $\times$ 2 trader roles $\times$ 2 experience conditions). However, in the Nash robots and Nash experienced conditions the underrevelation ratios shift toward the BNE in only 4 of the 8 cases.

[11] A possible objection to the direct tests in Table IV is that they do not account for the overall shape of the bid and ask functions. Traders might, for example, use approximately linear functions whose slope but not intercept (or intercept but not slope) responds appropriately to k. To investigate this possibility we also fit the following regression equation for buyers (and an analogous equation for sellers):

$$b_{it} = \alpha_0 + \alpha_1 I_{0.5} + \alpha_2 I_1 + \beta_0 v_{it} + \beta_1 I_{0.5} v_{it} + \beta_2 I_1 v_{it} + e_{it},$$

where b_{it} is bid entered by buyer i in period t, v_{it} denotes the value draw, and the indicator variables I take on a value of 1 for the k treatment indicated by the subscript and 0 otherwise. The β_1 and β_2 coefficients of the value and treatment interactions pick up changes in the slope associated with changes in the k treatment, while the α_1 and α_2 coefficients pick up shifts in the intercept. The regression results are available upon request. As in the matched-pair tests of Table IV, the best support for the theory seems again to come from buyers facing Nash robots, with additional support in the Nash experienced treatment. The fitted intercept coefficients are too large and the slope coefficients too low in the other two treatments, which suggests that traders in these treatments tend to underrespond to their value and cost.

provides a stationary environment that promotes learning, because unlike the human opponent conditions the robots play known, fixed strategies. Another interpretation of this result has more to do with strategic uncertainty than with stationary learning environments. As emphasized by an anonymous referee, merely knowing that all other traders are "rational" greatly simplifies the trader's decision problem. This interpretation receives some support from the observation that the k-treatment effect becomes insignificant as subjects move from the Nash robots to the Nash experienced condition facing other human opponents (compare rows (4) and (5) of Table IV).

Yet another possible interpretation is that subjects tend only to *mimic* the BNE strategies of their robot opponents, rather than learning best responses to these Nash robots. We conducted the two revealing robots sessions (indicated at the bottom of Table I) to test this explanation. In these sessions all 7 robot opponents entered offers equal to their values or costs. As shown in Appendix A the unique best response for risk neutral buyers (sellers) is to underreveal values (costs) by 20 percent when $k=1$ ($k=0$). The point is that best response behavior is exactly the same with revealing robots as with Nash robots, but mimicking behavior is quite different.

The results provide no support for the mimicking hypothesis. Behavior in the revealing robots and Nash robots sessions is statistically indistinguishable. When pairing offers for the identical value or cost draw, the (payoff importance weighted) mean difference is about ten cents in the $k=1$ bids and the $k=0$ asks conditions; these are not statistically different from zero based on the same random effects error specification employed in Table IV. Similar to the Nash robots results of Table IV, bids decrease as predicted when moving from $k=0$ to $k=1$ in the revealing robots sessions. Asks, however, increase when moving from $k=0$ to $k=1$ in the revealing robots sessions.[12] Bid function estimates similar to those discussed in footnote 11 provide very good support for the theoretical best response predictions in the revealing robots sessions.

4.4. *Deviations from the Risk Neutral BNE*

There are more direct tests for the risk neutral BNE model. Table VA presents the difference in the offers from their risk neutral BNE predictions separately for each k treatment and experience condition. This difference is distributed asymmetrically and includes a number of outliers arising from "throwaway" offers (i.e., very low (high) bids (asks) when the value (cost) draw is relatively low (high)). Therefore, we employ the median as a measure of central tendency and the distance between the first and third quartiles as the dispersion measure.

[12]Although it does not bear directly on the mimicking hypothesis, we should mention that the $k=1$ asks in the two different robot opponent treatments are significantly different. Recall that full revelation is the dominant strategy for these sellers. Yet in the revealing robots $k=1$ session asks are on average 28 cents higher than in the Nash robots $k=1$ session.

TABLE VA

MEDIAN DEVIATIONS FROM RISK NEUTRAL BNE BIDS AND ASKS: HIGH EXPECTED LOSSES SUBSAMPLE

Experience Condition	Median (Bid-BNE Bid)			Median (Ask-BNE Ask)		
	$k=0$	$k=0.5$	$k=1$	$k=0$	$k=0.5$	$k=1$
Inexperienced	−0.33***	0.09**	0.38***	−0.36***	−0.01	0.13***
	(0.59)	(0.27)	(0.22)	(0.39)	(0.58)	(0.49)
Experienced	−0.10***	0.17***	0.42***	−0.41***	−0.16**	0.13***
	(0.21)	(0.26)	(0.23)	(0.29)	(0.37)	(0.48)
Nash Robots	−0.14***	0.07	0.29***	−0.40***	−0.18	0.12***
	(0.24)	(0.31)	(0.38)	(0.54)	(0.42)	(0.42)
Nash Experienced	−0.01***	0.24***	0.43***	−0.53***	−0.23***	0.01***
	(0.09)	(0.13)	(0.13)	(0.11)	(0.10)	(0.04)

Notes: The interquartile range (Q3−Q1) is shown in parentheses to measure dispersion. Data are reported for the value and cost draws with the greatest expected losses from suboptimal offers: values in [\$2.00, \$3.50], and costs in [\$1.50, \$3.00]. Two-tailed median tests of the null hypothesis that the median difference equals zero: *** denotes significantly different from zero at 1 percent; ** denotes significantly different from zero at 5 percent; * denotes significantly different from zero at 10 percent.

Table VA sharpens some inferences regarding deviations from the BNE bid and ask functions. The most striking finding here is that in every case, the median deviation is significantly negative for $k=0$ and significantly positive for $k=1$. That is, buyers and sellers offer above the low predictions of the $k=1$ BNE and offer below the high predictions of the $k=0$ BNE. Median deviations for $k=0.5$ are sometimes significant and are positive for bids and negative for asks. The median deviation moves towards 0 with increasing experience in the dominant strategy full revelation cases ($k=0$ bids and $k=1$ asks), but elsewhere shows no clear trend.[13]

This movement toward the dominant strategy prediction with increased experience is encouraging, because this prediction is more basic than the nondominant BNE predictions. The rate at which offers conform to the dominant strategy is also similar to results from previous one-sided dominant strategy second-price auctions with independent private values. For example, inexperienced subjects in the second-price auctions reported in Kagel and Levin (1993) bid within 5 cents of value about 30 percent of the time. Our inexperienced subjects offer within 5 cents of the dominant strategy prediction about 33 percent of the time. For our experienced subjects this frequency increases to 54 percent.

In some other respects these bid and ask results are reminiscent of results in the extensive experimental literature on one-sided sealed bid auctions; see

[13] Kolmogorov-Smirnov distribution tests of the hypothesis that the observed and BNE offer distributions are equal generally agree with the median tests of the table: they sometimes do not reject the BNE when $k=0.5$, but they do reject the BNE when $k=1$ or $k=0$. We would prefer to conduct these tests using a random effects error structure as employed above, but this parametric error structure cannot be readily incorporated into nonparametric Kolmogorov-Smirnov and median tests.

Kagel (1995) for a survey. Kagel and Levin (1993) report first-price and second-price independent private value auction results that are similar to many other studies; their results are particularly useful for the present comparison because in their parameterization the risk neutral BNE predictions are almost identical to ours. Specifically, the predicted $VUR = 0.20$ (i.e., the BNE bid is 80 percent of value) for their first price $N = 5$ buyer treatment and for our $k = 1$ buyer (and $k = 0$ seller) treatment; and the predicted $VUR = 0.10$ for their first price $N = 10$ buyer treatment, which is essentially the same for our $k = 0.5$ buyer and seller treatments. Of course the dominant truth-telling strategy $VUR = 0$ applies to their second price treatments as well as to our $k = 0$ buyer (and $k = 1$ seller) treatments.

Similar to our results but contrary to theoretical predictions, Kagel and Levin (1993) find positive and significant intercepts for their bid function estimates in both first-price and second-price auctions (see their Table 2).[14] However, their bid function slope estimates are (a) consistent with theory in the second-price auction and (b) greater than the theoretical prediction in the first-price auction. Our corresponding slope point estimates are too low in 17 of the 24 cases (3 k-treatments $\times$ 2 trader roles $\times$ 4 experience conditions), and are almost uniformly too low in the inexperienced and experienced conditions. Kagel and Levin (1993) also find a significant shift in the estimated first-price bid function slope when increasing N from 5 to 10 that is consistent with the BNE comparative static prediction. Our results shown in Table IV and offer function estimates (see footnote 11) provide less support for BNE bid and ask function shifts in response to changes in the price rule k.

4.5. *Tests of Revelation Hypotheses* 3–5

The remaining hypotheses concern the predicted symmetry in value and cost revelation for the different k treatments. The summary statistics we employ for the amount of underrevelation in the different treatments are the underrevelation ratios *VUR* and *CUR* defined in Section 2. The mean is a poor indicator of central tendency for the ratios because of a substantial number of outliers, particularly when the denominator of these ratios is small. Hence we report the median as the measure of central tendency, and use the interquartile range to measure dispersion. We employ two-sample tests of the null hypothesis that the medians are equal, as well as Kolmogorov-Smirnov (K-S) tests of the null hypothesis that the distributions are equal, and restrict the sample to the "important" observations with the greatest expected payoff consequences.

The results reported in Table VI point to several conclusions. First, Hypotheses 3, 4, and 5 all are generally consistent with the data. Hypothesis 3 is that

[14] Other studies have also found positive bid function intercepts in first-price auctions. Cox et al. (1988) provide some theoretical justification for nonzero bid function intercepts in this context based on buyers' (nonmonetary) utility of winning the auction and a minimum potential income threshold for nonzero bids.

TABLE VB

MEDIAN DEVIATIONS FROM EMPIRICAL BEST RESPONSE (EBR) BIDS AND ASKS: HIGH EXPECTED LOSSES SUBSAMPLE

Experience Condition	Median (Bid-EBR Bid) $k = 0.5$	Median (Bid-EBR Bid) $k = 1$	Median (Ask-EBR Ask) $k = 0$	Median (Ask-EBR Ask) $k = 0.5$
Inexperienced	0.02	0.19***	−0.11***	0.03
	(0.25)	(0.23)	(0.35)	(0.61)
Experienced	0.15*	0.19***	−0.22***	−0.12**
	(0.26)	(0.20)	(0.29)	(0.39)
Nash Experienced	0.24***	0.35***	−0.46***	−0.24***
	(0.11)	(0.16)	(0.11)	(0.13)

Notes: The interquartile range (Q3−Q1) is shown in parentheses to measure dispersion. Data are reported for the value and cost draws with the greatest expected losses from suboptimal offers: values in [\$2.00, \$3.50], and costs in [\$1.50, \$3.00]. Two-tailed median tests of the null hypothesis that the median difference equals zero: *** denotes significantly different from zero at 1 percent; ** denotes significantly different from zero at 5 percent; * denotes significantly different from zero at 10 percent.

buyers and sellers underreveal to the same degree when $k = 0.5$. Although the median underrevelation is always below the predicted value of 9.4 percent, the p-values in the third column show that we never come close to rejecting the null hypothesis that the medians are the same for buyers as for sellers.[15] The K-S test rejects the equal distribution null in the inexperienced treatment, but at only the $p = 10$ percent significance level. Hypothesis 4 is that $k = 0$ sellers and $k = 1$ buyers underreveal to the same (theoretically maximal) degree. We reject the hypothesis only in the Nash experienced treatment, and even here the economic difference is rather slight (2.1 percent versus 0.5 percent underrevelation, both far below the predicted 20 percent). Hypothesis 5 is that the $k = 0$ buyers and $k = 1$ sellers underreveal to the same degree (theoretically 0). The hypothesis is rejected at the $p = 10$ percent level in the inexperienced treatment, and the K-S test rejects the equal distribution null in the experienced treatment at the $p = 5$ percent level. However, the economic difference in the experienced treatment (4.4 percent versus 6.1 percent) is also quite small.[16]

Perhaps the support in Table VI for Hypotheses 3–5 arises mainly from the low power of the tests. If so, the tests should detect no impact for other treatments such as experience. The middle and bottom rows of the table show otherwise. Of the 12 cases (3 k-treatments × 2 pairs of experience conditions × 2 trader types) the median underrevelation declines with experience in 11 cases. Nine of the median tests and nine of the K-S tests are significant at the 10 percent level. The majority of cases are significant at the 5 percent level.

[15] A within-subject test for the $k = 0.5$ treatment that buyer and seller strategies were symmetric leads to a similar conclusion. An F-test that the bid function and ask function are symmetric (where asks and costs are appropriately transformed to (4.99-ask) and (4.99-cost) to imply symmetry with respect to the bid function) fails to reject at the 5 percent level the symmetry null hypothesis for 29 of the 40 subjects.

[16] The median underrevelation ratios range between 0.035 and 0.064 in the revealing robots sessions, and none are significantly different from the comparable Nash robots sessions.

TABLE VI

MEDIAN VALUE AND COST UNDERREVELATION IN SYMMETRIC k TREATMENTS: HIGH EXPECTED LOSSES SUBSAMPLE

	Value Underrevelation Ratio (VUR) = (value − bid)/value Cost Underrevelation Ratio (CUR) = (ask − cost)/(4.99 − cost) Hypothesis 3:			Hypothesis 4:			Hypothesis 5:		
	(1) $k = 0.5$ VUR	(2) $k = 0.5$ CUR	p-values comparing col. 1 and 2	(3) $k = 1.0$ VUR	(4) $k = 0$ CUR	p-values comparing col. 3 and 4	(5) $k = 0$ VUR	(6) $k = 1.0$ CUR	p-values comparing col. 5 and 6
Risk Neutral BNE	0.094	0.094		0.20	0.20		0.0	0.0	
	Inexperienced:								
(1)	0.065 (0.109)	0.091 (0.224)	*0.419* **0.072**	0.061 (0.099)	0.077 (0.143)	*0.292* **0.466**	0.116 (0.177)	0.048 (0.166)	*0.079* **0.059**
	Experienced:								
(2)	0.025 (0.098)	0.024 (0.132)	*0.923* **0.319**	0.045 (0.078)	0.031 (0.129)	*0.383* **0.913**	0.044 (0.074)	0.061 (0.159)	*0.558* **0.031**
p-values comparing row 1 and 2	*0.069* **0.029**	*0.008* **0.018**		*0.128* **0.459**	*0.057* **0.108**		*0.017* **< 0.001**	*0.537* **0.812**	

TABLE VI—*Continued*

Value Underrevelation Ratio (VUR) = (value − bid)/value
Cost Underrevelation Ratio (CUR) = (ask − cost)/(4.99 − cost)

	Hypothesis 3:			Hypothesis 4:			Hypothesis 5:		
	(1) $k = 0.5$ VUR	(2) $k = 0.5$ CUR	p-values comparing col. 1 and 2	(3) $k = 1.0$ VUR	(4) $k = 0$ CUR	p-values comparing col. 3 and 4	(5) $k = 0$ VUR	(6) $k = 1.0$ CUR	p-values comparing col. 5 and 6
	Nash Robots:								
(3)	0.055	0.020	*0.674*	0.079	0.057	*0.711*	0.065	0.049	*0.349*
	(0.142)	(0.127)	**0.840**	(0.152)	(0.152)	**0.187**	(0.079)	(0.139)	**0.319**
	Nash Experienced:								
(4)	0.001	0.002	*0.491*	0.021	0.005	*< 0.001*	0.005	0.007	*0.496*
	(0.012)	(0.007)	**0.284**	(0.067)	(0.012)	**< 0.001**	(0.032)	(0.018)	**0.841**
p-values comparing row 3 and 4	*0.002*	*0.123*		*0.005*	*0.042*		*0.026*	*0.002*	
	< 0.001	**0.055**		**< 0.001**	**0.001**		**0.005**	**< 0.001**	

Notes: The interquartile range (Q3−Q1) is shown in parentheses to measure dispersion. Data are reported for the value and cost draws with the greatest expected losses from suboptimal offers: values in [\$2.00, \$3.50], and costs in [\$1.50, \$3.00]. Median test p-values are shown in *italics*, Kolmogorov-Smirnov test p-values are shown in **bold**.

The decline of underrevelation with experience is an important finding in its own right. For the cases covered in Hypothesis 5, the implication is that traders' behavior comes rather close to the BNE benchmark with experience. Indeed, in the "best" experience treatment in line (4), the table shows median underrevelation of less than 1 percent, very close to the predicted full revelation. For the other cases (covered by Hypotheses 3 and 4) the move towards full revelation is equally strong. Of course, in these cases the move is away from BNE, not towards it.[17]

For a while we entertained the following potential explanation of our results. A few traders irrationally overreveal their values or costs, and then most other traders rationally overreveal in response. We encountered both theoretical and empirical difficulties as we examined the explanation. The theoretical difficulty is that (contrary to our intuition) bids are not strategic complements in the SCM. Proposition 1 of Appendix A shows that with pricing rule k, the best response to truth-telling is $VUR = k/(4+k)$, which (as indicated in the first line of Table VI) is essentially the same as the best response to BNE behavior by other traders. Proposition 2 of Appendix A shows that in general the best response depends on something like the Mills ratio for the offer distribution, not on the degree of underrevelation per se.

The theoretical problem is not insurmountable because the best response to the empirical distribution of bids and asks might still turn out to involve about the same degree of underrevelation as we actually observed. For each session we calculated the degree of underrevelation in the empirical best response (EBR).[18] Table VB compares actual bids and asks to the EBR strategies using the same format and conventions as Table VA uses for BNE strategies. It omits the cases where EBR coincides with BNE because the strategy is dominant (first and last columns of Table VA) or because the empirical distribution is fixed (the Nash robots row of Table VA). The median deviations reported in Table VB generally have the same sign as those in Table VA and typically are smaller in absolute value, so the empirical best response explanation goes in the right direction. But quantitatively it is inadequate. Replacing BNE in Table VA by EBR in Table VB reduces the median deviation consistently only in the inexperienced sessions, and perhaps not at all in the Nash experienced sessions where the impact would be expected to be strongest. Moreover, the Tables VA and VB have

[17] Nash experienced traders are closest to full revelation, significantly closer than experienced traders for $k=0$ and $k=1$ asks. (This statement is based on a matched-pair analysis (for identical cost draws) using importance weights and a subject random effects error specification.) We are unable to explain this difference, although it suggests that the type of trader opponent experience influences learning and therefore subsequent behavior.

[18] This calculation involved three steps: (i) fit the empirical distributions of bids and asks with a fourth-order polynomial; (ii) draw 10,000 sets of seven "other" bids and asks from this fitted empirical distribution; and (iii) calculate the best response to this simulated distribution of rival offers. We carried out this approximation separately for each laboratory session. Details appear in Appendix B. While our procedure is based on a simulation and does not provide exact results, we believe that sensible alternative procedures would produce substantially the same results.

essentially the same pattern of significant deviations. We conclude that, although some traders might be making approximate best responses to the empirical distribution, most traders do not. In the most relevant conditions, the median trader reveals significantly more than prescribed by either her EBR or her BNE strategy.[19]

4.6. *Learning Direction Theory and the Movement Toward Increased Revelation*

It is worth noting that when pooling across subjects the underrevelation ratio distributions summarized above in Table VI are generally unimodal and skewed toward positive values. However, the distribution of median underrevelation ratios when calculated separately for each subject is often bimodal. A number of subjects cluster around full revelation, while another group of subjects' ratios fall in the range of 10 to 25 percent underrevelation. The following analysis demonstrates that subjects tend to migrate to the full revelation group as they obtain experience.

Define a "high revelation" subject as one who has a median underrevelation ratio less than 10 percent, and a "low revelation" subject as one who has a median underrevelation ratio greater than or equal to 10 percent. The number of low revelation subjects falls significantly with experience. For the $k = 0.5$ treatment, 46 percent of the 48 inexperienced subjects are classified as low revelation, while only 16 percent of the 32 experienced subjects are classified as low revelation. This difference is significant at one percent (chi-squared (1 d.f.) = 7.84). For the $k = 1$ buyers and $k = 0$ sellers (in which BNE predicts an underrevelation ratio of 20 percent), 48 percent of the 61 inexperienced subjects and 16 percent of the 32 experienced subjects are classified as low revelation. This difference is also significant at one percent (chi-squared (1 d.f.) = 9.22). Finally, for the $k = 0$ buyers and $k = 1$ sellers (with zero underrevelation as the dominant strategy), 44 percent of the 61 inexperienced subjects and 22 percent of the 32 experienced subjects are classified as low revelation. This difference is significant at five percent (chi-squared (1 d.f.) = 4.54). Therefore, we conclude that a substantial number of subjects increase their value and cost revelation with increased experience, which leads to the movement toward greater revelation for the sample overall.

The data also provide evidence that subjects obtain more profit in this setting when they reveal more of their values and costs. Define the profit performance

[19] Results on individual subjects presented in Appendix B also support this same conclusion. Most individual subjects overreveal compared to EBR for bids with $k = 1$ and asks with $k = 0$, and underreveal compared to the dominant strategy of full revelation for bids with $k = 0$ and asks with $k = 1$. Individual subjects' offers are more evenly distributed above and below the EBR in the intermediate case of $k = 0.5$, except in the Nash experienced condition where overrevelation relative to EBR is much more common than underrevelation.

of each subject as his actual profit divided by his profit at the competitive equilibrium benchmark. The correlation between profit and median underrevelation ratio is -0.41 for experienced subjects, which is significantly different from zero at better than one percent. This implies that those subjects who revealed more of their values and costs earned more profit on average, which could be one reason that value and cost revelation increased with experience. It also suggests that those traders who underreveal fail to identify the optimum and underrevealed too much.

More direct insight into the surprising result that increasing experience leads to greater revelation—even when this is contrary to the BNE predictions—is provided by a brief exploration of traders' adjustment process. We begin with Selten and Buchta's (1994) "directional learning model" for single-sided sealed bid auctions, which is easily adapted to the SCM. The basic idea is that traders who missed a profitable trading opportunity due to (ex post) excessive underrevelation will reveal more the next time, and traders who transacted at a price less favorable than the (ex post) optimum will reveal less next time. The model makes no predictions for the event that the market "outpriced" the buyer (seller) with a clearing price above value (below cost).

Table VII tests this directional learning idea after subsetting to the most important bid and ask data. In keeping with the naive nature of ex post reasoning, column (1) looks at the event that the trader missed a profitable

TABLE VII

IMMEDIATE OFFER RESPONSES TO PREVIOUS MARKET OUTCOMES

	Previous Market Outcome:		Dependent Variable: Change in Underrevelation Ratio				
	Missed Profitable Trade (1)	Affected Market Price (2)	Value of Lost Surplus when Missing Profitable Trade	Value of Lost Surplus when Affecting Market Price	Intercept	R^2	Observations
increase	*Inexperienced*						
revelation/	42/54	22/40	-0.164^{***}	0.091	-0.006	0.11	275
total cases	($p < 0.001$)	($p = 0.486$)	(0.031)	(0.065)	(0.007)		
increase	*Experienced*						
revelation/	23/31	19/52	-0.190^{***}	0.045	-0.002	0.10	263
total cases	($p = 0.003$)	($p = 0.027$)	(0.039)	(0.032)	(0.005)		
increase	*Nash Robots*						
revelation/	57/67	7/13	-0.338^{***}	0.897	0.002	0.10	173
total cases	($p < 0.001$)	($p = 0.343$)	(0.092)	(0.874)	(0.018)		
increase	*Nash Experienced*						
revelation/	14/14	28/52	-0.181^{***}	0.018	-0.004	0.05	262
total cases	($p < 0.001$)	($p = 0.789$)	(0.040)	(0.019)	(0.004)		

Notes: Data are reported for the value and cost draws with the greatest expected losses from suboptimal offers: values in [\$2.00, \$3.50], and costs in [\$1.50, \$3.00]. The p-values on the left side of the table are based on Fisher's Exact Test (Null Hypothesis: The change in value or cost revelation is independent of the indicated previous market outcome.) Regression results on the right side of the table are estimated using GLS with a random effects error specification $e_{it} = u_i + \varepsilon_{it}$, with subjects as the random effect, and are adjusted to correct for significant first-order serial correlation. Standard errors are in parentheses.

trading opportunity when price was below value (above cost) for a buyer (seller) who did not transact.[20] Column (2) examines the event that a trader adversely affected the price. For example, a transacting buyer bid too high (ex post) if she is on the margin (thus affecting the clearing price) and could have entered a lower successful bid to lower the price she paid. We see from column (1) that traders indeed are much more likely to increase than decrease (or leave constant) the degree of revelation after missing a profitable trade. On the other hand, column (2) shows that the evidence at best is mixed that they are likely to decrease revelation after adversely affecting price.

To estimate the strength of the effects, we move from Selten and Buchta's binary variables (the event occurred or not, and revelation increased or not) to continuous variables measuring the ex post lost profit in each event and the change in underrevelation. The right side of Table VII shows that missed trades have a very statistically (and economically) significant impact on underrevelation. The negative coefficients are appropriate; they imply that when traders miss a profitable trade they reveal more (significantly reduce the degree of underrevelation) on the next (important) opportunity. The coefficients for adversely affecting price also have the appropriate sign, but are statistically (and often economically) insignificant. These statements hold about equally well for each experience treatment. Thus we find evidence of an asymmetry in traders' learning process that pushes them toward truth-telling.

Trading efficiency increases if traders' bids and asks reveal more of their underlying values and costs, and the results indicate the subjects tend on average to reveal more of their values and costs than predicted by the risk neutral BNE. It may therefore seem difficult to reconcile this "over-revelation" result with the earlier result shown in Table II that efficiency falls below the BNE benchmark. However, efficiency suffers primarily because of the subset of "low revelation" traders who tend to underreveal too much. Overall trading efficiency can easily fall below 90 percent if only one or two of the eight traders frequently underreveal their values or costs more than the BNE suggests, and consequently miss out on profitable trades. Most sessions included at least one such trader.

5. CONCLUSION

Our goal in this laboratory study has been to evaluate the Bayesian Nash equilibrium (BNE) as an explanation of price formation in the single call market (SCM). The experimental design includes a number of features intended to give the theory its "best shot" at organizing the data, and the results identify some strengths and some weaknesses of BNE. The SCM institution guarantees a unified price each period, and the prices and market efficiency we observe are

[20] A more sophisticated definition, which differs in a few marginal cases for $k < 1$ buyers and $k > 0$ sellers, is that (given the other 7 bids and asks) there is a bid (ask) the nontransacting buyer (seller) could have made that would have been profitable even allowing for its effect on price. Results are very similar when using this alternative definition.

closer to the BNE predictions than they are to the competitive equilibrium on the one hand, or to the "zero-intelligence" predictions on the other. But we observe systematic deviations of actual outcomes from BNE predictions. For example, BNE predicts greater responsiveness of transaction prices to the pricing rule and predicts higher efficiency than we observe in our laboratory markets.

To better understand the discrepancies in market outcomes, we focused our analysis on buyers' bids and sellers' asks. For the most part, buyers and sellers do not shade their offers away from actual values and costs as much as BNE predicts, and they tend to offer below BNE predictions for the low $(k = 0)$ pricing rule and above BNE for the high $(k = 1)$ pricing rule. Experience on the whole does not bring trader behavior closer to BNE predictions. Indeed, with increasing experience, traders tend to more fully reveal their true values and costs, whether or not greater revelation is consistent with BNE.

We do not regard our results as a rejection of the BNE theory. One rejects a theory only for some alternative and, despite its weaknesses, the BNE theory clearly does a better job of explaining our data than explicit alternative theories such as competitive equilibrium. However, the results do give us some important clues that may lead to a deeper understanding of price formation in the SCM. First, the BNE predictions are most accurate in sessions in which inexperienced traders face robots that use the BNE bid and ask functions. The revealing robots sessions allow us to reject the hypothesis that this arises simply from subjects mimicking the robots' BNE strategies. Second, behavioral noise is smallest at intermediate value and cost realizations, where the expected payoff consequences are largest. Third, the main systematic effect of trader experience seems to be a reduction in the extent to which traders shade offers away from true values or costs (or more precisely, a reduction in the number of traders who strongly shade their offers).

These clues suggest to us that learning models may provide a useful account of the price formation process. The Nash robots treatment was chosen precisely to aid learning, and indeed it produced behavior generally closest to BNE. It is not obvious whether this is due to the stationary learning environment or because this treatment eliminates strategic uncertainty, and we believe robot opponents are a useful tool to address this issue in future experiments. Although the learning analysis based on directional learning theory is tentative and preliminary, it suggests a learning bias which may account for the fact that experience did not quickly lead traders to BNE behavior, and instead led to excessive value and cost revelation. In the BNE, traders are fully rational and balance the marginal loss from the reduced probability of transacting against the marginal gain conditional on transacting. In contrast, rationality is quite bounded in the directional learning approach, and subjects seem to respond more to the strong negative reinforcement of missing obviously profitable (ex post) trading opportunities than to the more subtle signal of affecting the clearing price. Following a period of adverse price impact, some traders do enter more aggressive offers, but other traders offer *less* aggressively, perhaps to avoid "coming so close" to missing a profitable trade.

We see several useful directions for follow-up work. The current random values SCM environment can be used to look more deeply at the learning process. One could employ treatments that encourage subjects to focus more directly on the marginal impact of offers on clearing price in order to explore the robustness of the learning asymmetry we detected. In terms of data analysis, one could adapt formal learning models to the present environment. The models should allow for heterogeneous behavior across traders since several of our present results suggest a significant role for heterogeneity. No definitive learning model yet is on the horizon, but some insight may be gained from existing models in the cognitive psychology and recent economics literature.

There is also a lot of important work to be done with other market institutions. The directional learning approach can be applied to trader behavior in most market institutions, and may provide further insights beyond those provided initially by Selten and Buchta in one-sided auctions. More generally, one should compare price formation in the SCM to price formation in related market institutions such as the multiple call market and the uniform price double auction (McCabe et al. (1993), Friedman (1993)). The equilibrium theory (not to mention the learning theory) is less developed for these other market institutions, but empirical comparisons may yield stylized facts that will help spur theoretical progress.

Dept. of Economics, University of Southern California, Los Angeles, CA 90089-0253, U.S.A.

and

Dept. of Economics, University of California, Santa Cruz, CA 95064, U.S.A.

Manuscript received March, 1995; final revision received April, 1996.

APPENDIX A: MATHEMATICAL DETAILS

Rustichini, Satterthwaite, and Williams (RSW) (1994 and 1990) and Satterthwaite and Williams (1989 and 1993) include a thorough analysis of BNE for the SCM. However, some of our hypotheses involve behavior outside of BNE and some key BNE formulae, especially in RSW (1990), are not easily accessible. The mathematical details collected below extend the RSW results in some minor respects, but their main purpose is to make the present paper more self-contained. To reduce the notational burden we focus on the case of $m = 4$ buyers and $m = 4$ sellers whose values and costs are i.i.d. draws from a uniform distribution on the interval $[0, M]$. In the text $M = \$4.99$ but we renormalize below to $M = 1$.

Let $(\hat{s}_1, \dots, \hat{s}_8)$ be the offers (bids and asks) of all traders sorted from lowest $(\hat{s}_1)$ to highest $(\hat{s}_8)$. Recall that for $k \in [0,1]$ the price in the k-SCM then is

$$\text{(A1)} \qquad p = (1-k)\hat{s}_4 + k\hat{s}_5.$$

We begin by deriving a general expression for a buyer's expected profit $\Pi(b, v)$ given a value draw v and bid b; the formulas for sellers are analogous with $1-k$ replacing k. Let $f(s_4, s_5)$ be the density of the fourth and fifth lowest offers by the 4 sellers and 3 *other* buyers $\{s_1, \dots, s_7\}$; Lemmas 2 and 3 below characterize f. Let $\mathbf{1}_e$ denote the indicator function for event e, so $\mathbf{1}_e = 1$ if e and $\mathbf{1}_e = 0$ if not e. Recall that the buyer's profit is 0 unless her bid equals or exceeds the price, i.e., unless $\hat{e} = [b \geq p]$. Assuming that f is a proper density, $\hat{e}$ coincides with $e = [b > s_4]$ except on the measure 0 event $[b = s_4]$. The buyer's profit in event e is $v - p$. Hence we have established

$$\Pi(b,v) = E_f(v-p)\mathbf{1}_{[b>s_4]} \equiv \int_{s_4=0}^{b}\int_{s_5=s_4}^{1}(v-p)f(s_4,s_5)\,ds_5\,ds_4. \tag{A2}$$

In working with (A2) it will be useful to refer to several calculus facts.

LEMMA 1: (a) *Let $g(x,y)$ be a continuous real valued function with continuous second partial derivative g_2 on the set $[0 \leq y \leq x < \infty]$, and let $G(x) = \int_0^x g(x,y)\,dy$. Then $G'(x) = g(x,x) + \int_0^x g_2(x,y)\,dy$.*

(b) *Let $g(x)$ be continuous on (a,c) and let $G(x) = \int_a^c \min\{x,u\}g(u)\,du$. Then $G'(x) = \int_x^c g(u)\,du$.*

$$\int_a^1 (1-x)^n dx = (n+1)^{-1}(1-a)^{n+1}. \tag{c}$$

PROOF: Direct computations from the definition of the derivative as $\lim_{h\to 0} h^{-1}(G(x+h) - G(x))$ yield (a) and (b). Part (c) can be obtained by direct calculation or by using the change of variable $y = 1 - x$. *Q.E.D.*

Now we characterize the density function f of the key order statistics s_4, s_5 of bids and asks. We begin with the full-revelation (or truth-telling) case.

LEMMA 2: *Suppose each seller asks his true cost and each other buyer bids her true value. Then*

$$f(s_4,s_5) = cs_4^3(1-s_5)^2, \tag{A3}$$

where c is a normalizing constant. For $M = 1$ we have $c = 7!/(3!2!)$.

PROOF: Here there are 7! equally likely orderings of bids and asks, whose joint density is 1 on the 7-hypercube. Hence the joint density of order statistics $s_1, \ldots, s_7$ has density 7! on the set $[0 \leq s_1 \leq s_2 \leq \cdots \leq s_7 \leq 1]$. Integrating out everything but s_4 and s_5, and using Lemma 1(c) for s_6 and s_7, we get

$$f(s_4,s_5) = \int_{s_1=0}^{s_2}\int_{s_2=0}^{s_3}\int_{s_3=0}^{s_4}\int_{s_6=s_5}^{1}\int_{s_7=s_6}^{1} 7!\,ds_1 ds_2 ds_3 ds_6 ds_7 = 7!\frac{s_4^3}{3!}\frac{(1-s_5)^2}{2!}. \qquad Q.E.D.$$

The general case is considerably more complicated. Suppose each seller (each other buyer) uses the strictly increasing, continuous, and piecewise differentiable ask function $a = A(c)$ (resp., bid function $b = B(v)$). Let $\alpha = A^{-1}$ and $\beta = B^{-1}$ be defined on the ranges of possible bids and asks, and let $e_i = (k_i, \mathbf{1}_{4i}, \mathbf{1}_{5i})$ denote the event that exactly k_i of $\{s_1, s_2, s_3\}$ are bids and that s_4 is a bid ($\mathbf{1}_{4i} = 1$) or an ask ($\mathbf{1}_{4i} = 0$), and that s_5 is a bid ($\mathbf{1}_{5i} = 1$) or an ask ($\mathbf{1}_{5i} = 0$). Given the constraint that k_i is a nonnegative integer bounded above by $3 - \mathbf{1}_{4i} - \mathbf{1}_{5i}$ since there are only 3 other buyers, there are 11 such events $i = 1, \ldots, 11$. The probability p_i of each event e_i can be determined in principle from the A and B functions but we have no simple formula.

LEMMA 3: *For the general symmetric case, let* $j_i = 3 - \mathbf{1}_{4i} - \mathbf{1}_{5i} - k_i$ *be the number of bids in* $\{s_6, s_7\}$. *Then*

$$f(s_4, s_5) = \sum_{i=1}^{11} \hat{p}_i \alpha(s_4)^{3-k_i} \beta(s_4)^{k_i} [1 - \alpha(s_5)]^{j_i} [1 - \beta(s_5)]^{2-j_i} \alpha'(s_4)^{1-\mathbf{1}_{4i}} \beta'(s_4)^{\mathbf{1}_{4i}} \times \alpha'(s_5)^{1-\mathbf{1}_{5i}} \beta'(s_5)^{\mathbf{1}_{5i}}, \tag{A4}$$

where $\hat{p}_i = 7! p_i / [(3-k_i)! k_i! (2-j_i)! j_i!]$.

PROOF: Suppose the ordering of asks A and bids B happens to be, say, $ABABABA$. Then the joint density of $(s_1, \dots, s_7)$ is $\alpha'(s_1)\beta'(s_2)\alpha'(s_3) \cdots \alpha'(s_7)$ times the density (7!) of the relevant values $v_i = \alpha(s_i)$ and costs $c_i = \beta(s_i)$. Integrating out this density as in Lemma 2 we get

$$f(s_4, s_5 | ABABABA) = 7! \alpha(s_4)^2 \beta(s_4) [1 - \alpha(s_5)][1 - \beta(s_5)] \beta'(s_4) \alpha'(s_5). \tag{A5}$$

The same expression holds for every other case in $e_i = (1,1,0)$, e.g. $AABBAAB$. Note that p_i times (A5) coincides with the integrand in (A4). But $f(s_4, s_5)$ is simply the sum over the events e_i of the joint density $p_i f(s_4, s_5 | e_i)$, which is precisely (A4). *Q.E.D.*

PROPOSITION 1: *A k-SCM buyer's best response to truth-telling opponents is* $B(v) = [4/(4+k)]v$.

PROOF: Plugging (A1) and (A3) into (A2) we get

$$\Pi(b, v) = c \int_{s_4=0}^{b} \left\{ \int_{s_5=s_4}^{1} [v - (1-k)s_4 - k \min(b, s_5)] s_4^3 (1 - s_5)^2 \, ds_5 \right\} ds_4. \tag{A6}$$

It is easy to see that Π is smooth, decreasing at $b = 1$ and increasing at $b = 0$ for $v \in (0,1)$ so its maximum in b for given v will occur at an interior solution of the first-order condition $0 = \partial \Pi / \partial b$. Differentiating (A6) and using Lemma 1(a), (b), and (c), and the identity $(1-k)b + kb = b$, we get

$$0 = c\{\ \}|_{s_4=b} + c\int_0^b \frac{\partial}{\partial b}\{\ \}\, ds_4 = c(v-b)b^3 \frac{(1-b)^3}{3} - ck\frac{b^4}{4}\frac{(1-b)^3}{3}.$$

Cancelling the positive common factor $cb^3(1-b)^3/3$, the first-order condition reduces to $0 = v - b - (kb/4)$ with solution $b = [4/(4+k)]v$. *Q.E.D.*

REMARK 1.1: Note that $b = 0.8v$ is optimal against truth-telling opponents as well as BNE opponents when $k = 1$. This strategic isomorphism may seem at first puzzling, but it is present in some other trading institutions. For example, Selten and Buchta (1994) show that the bid function $b = [(N-1)/N]v$ (where N is the number of buyers) is optimal for risk neutral buyers in one-sided, first-price independent private value auctions in the BNE as well as against opponents that use any linear bid function with slope greater than $(N-1)/N$, including of course truth-telling.

PROPOSITION 2: *A k-SCM buyer with value* v *facing the marginal bid/ask density* $f(s_4, s_5)$ *will optimally choose a bid* b *satisfying the first-order condition*

$$(v - b)\int_{s_5=b}^{1} f(b, s_5)\, ds_5 = k\int_{s_4=0}^{b}\int_{s_5=b}^{1} f(s_4, s_5)\, ds_4\, ds_5. \tag{A7}$$

PROOF: As in the previous proof, we find that the optimal bid will satisfy the first-order condition obtained by plugging (A1) (but now not (A3)) into (A2) and differentiating. We get

$$0 = \left\{ \int_{s_5 = s_4}^{1} [v - (1-k)s_4 - kb] f(s_4, s_5)\, ds_5 \right\} \Bigg|_{s_4 = b}$$
$$- \int_{s_4=0}^{b} \frac{\partial}{\partial b} \int_{s_5=b}^{1} [v - (1-k)s_4 - k \min(b, s_5)] f(s_4, s_5)\, ds_4\, ds_5.$$

Equation (A7) follows immediately on simplifying, using the identity $(1-k)b + kb = b$ and Lemma 1(a) and (b). *Q.E.D.*

COROLLARY 2: *Optimal bids are fully revealing when* $k = 0$.

PROOF: Obvious from (A7).

COROLLARY 3: *Suppose* $f(s_4, s_5) = g(s_4)h(s_5)$ *for some functions* f *and* g *on the triangle* $0 < s_4 < s_5$, *and let* G *be the cumulative distribution function associated with the marginal density* g. *Then the optimal bid satisfies*

(A8) $\quad (v - b)g(b) = kG(b).$

Consequently, the degree of underrevelation is increasing in k, *ceteris paribus in* g. *Indeed, it is proportional to* k *if the Mills ratio* g/G *is constant.*

PROOF: (A8) follows from plugging $f = gh$ into (A7) and canceling the common factor $\int_b^1 h(s_5)\, ds_5$. *Q.E.D.*

REMARK 2.1: An interpretation of Corollary 3 is that a bidder's subjective beliefs about others' behavior, summarized in $f(s_4, s_5)$, may be insensitive to k and may be approximately separable. In this case an expected-profit maximizing buyer will obey Hypothesis 1.

REMARK 2.2: Corollary 3, and more generally equation (A7), show that bids are not generally strategic complements (or substitutes). For fixed $k > 0$, the degree of underrevelation by other traders affects the best response only via the Mills ratio in the separable case, or more generally via the ratio of the integrals in (A7).

REMARK 2.3: The decision problem for sellers is symmetric with $1 - k$ replacing k and $M - x$ replacing x. Hence we obtain analogous results for sellers' profit functions and ask functions. The parameter $M > 0$ (the upper bound on values and costs) affects the constant c but it cancels out in the propositions and corollaries.

APPENDIX B: ESTIMATION OF THE EMPIRICAL BEST RESPONSES

Estimation of the Empirical Best Responses (EBR) proceeded in three steps: (1) fit the empirical distributions of bids and asks; (2) draw 10,000 sets of seven "other" bids and asks from this empirical distribution; and (3) calculate the best response to this simulated distribution of rival offers. We carried out this procedure separately for each laboratory session.

Step 1: *Fitting the Empirical Distribution of Bids and Asks*

Each inexperienced session generated approximately 120 bids and 120 asks (e.g., 4 buyers for 30 periods), and each experienced session generated approximately 160 bids and 160 asks. We first assembled the bids and asks for each session separately into empirical cumulative distribution functions. We then approximated these empirical distribution functions with a fourth-order polynomial to permit an efficient simulation of rival offers in Step 2. Only a minimal amount of information was lost in this approximation, since it accounts for over 99 percent of the variation in the exact empirical distribution for all sessions.

Step 2: *Drawing* 10,000 *Rival Offer Sets per Session*

Again separately for each session, we used the empirical distribution functions from Step 1 to draw 10,000 sets of offers for the seven other traders. For buyers we drew 3 other buyer bids independently from the empirical distribution of bids, and drew 4 seller asks independently from the empirical distribution of asks. We ranked these 7 offers and retained the fourth and fifth lowest (s_4 and s_5 in the notation of Appendix A) because only those two are necessary to calculate the profit from an arbitrary buyer bid. For sellers we independently drew 4 buyer bids and 3 other seller asks from the appropriate empirical distributions, and retained the third and fourth lowest (s_3 and s_4) for calculating the profit from any seller ask in Step 3.

Step 3: *Calculating Best Responses*

We calculated the profit for each of 31 possible offers (0 percent underrevelation, 1 percent underrevelation, $\dots$, 30 percent underrevelation) for a discrete set of value and cost draws (\$0.00, \$0.20, $\dots$, \$4.80) for each of the 10,000 sets of other offers that were calculated in Steps 1 and 2 based on the empirical offer distributions. We then used the mean profit for each underrevelation level (over the 10,000 draws) as an estimate of the expected profit of each underrevelation level, and the underrevelation level with the highest expected profit is our estimate of the empirical best response for that value or cost draw.

Results Summary

Our empirical comparison with the actual offers focuses on the range of values and costs with the highest expected losses from suboptimal bidding (i.e., values in [\$2.00, \$3.50] and costs in [\$1.50, \$3.00]; see Appendix A). Figure B-1 presents an example of the results for buyers in this range. Results are symmetric for sellers and support three main conclusions. First, the EBR requires more value and cost revelation than the BNE prediction, with the bias toward greater revelation highest when the BNE predicts the least revelation (i.e., $k = 1$ for buyers). Second, the BNE comparative static prediction for the pricing rule k (Hypotheses 1 and 2) holds for the EBR, but with a smaller predicted difference with the change in k. Third, the deviation of the EBR from the BNE is the lowest for the Nash experienced condition. This final result is expected because, as shown in Proposition 1 of Appendix A, best responses to truthtelling offers are almost identical to the BNE predictions in the various k treatments, and traders' offers were closest to truth-telling in the Nash experienced condition.

The comparison to the actual offers requires an estimate of the EBR for all values and costs, rather than only the 20-cent increments for which we explicitly calculated an EBR. To fill in the EBR estimate for all values and costs we estimated linear equations based on the points graphed in Figure B-1. These equations fit with an R-squared typically above 99 percent.

Table VB in the text indicates that the median deviations of actual offers from the EBR are typically smaller than the median deviations from the BNE, but also usually differ significantly from the EBR. Table B-I presents a classification of individual subjects' offers relative to the EBR to determine if the bias relative to the EBR observed in the aggregate is based on a small set of

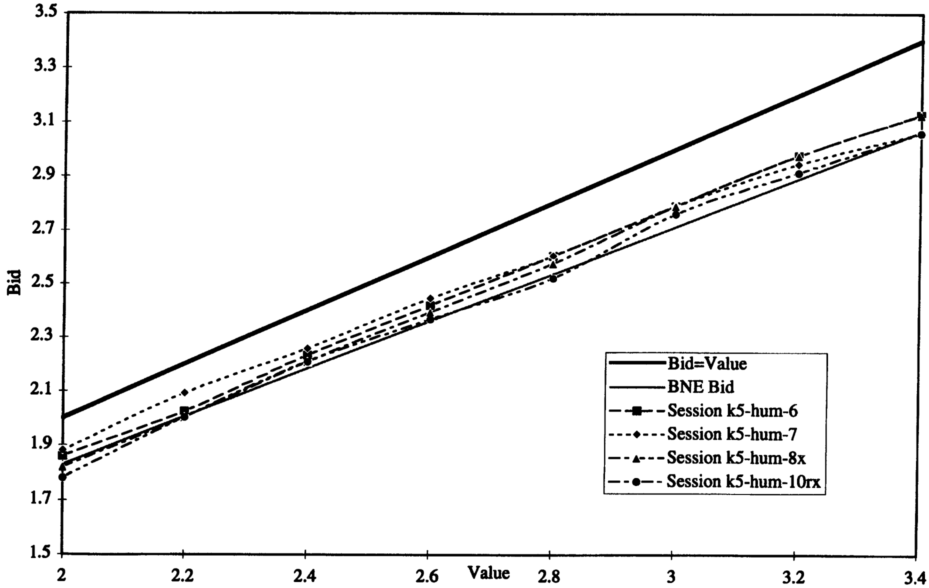

FIGURE B-1A.—Bid empirical best responses for $k = 0.5$ (important values).

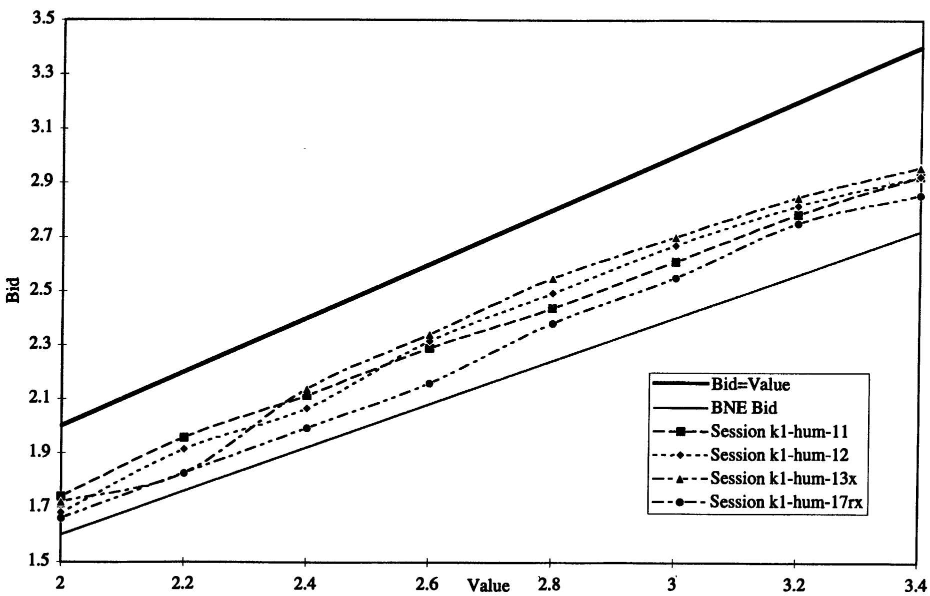

FIGURE B-1B.—Bid empirical best responses for $k = 1.0$ (important values).

TABLE B-I

Number of Subjects who Overrevealed and Underrevealed Relative to the Empirical Best Response (EBR)

		(1) $k=1$ bids & $k=0$ asks	(2) $k=0.5$ bids & $k=0.5$ asks	(3) $k=0$ bids & $k=1$ asks
BNE underrevelation ratio		0.2	0.094	0
EBR underrevelation ratio (average of separate estimates for each session)		0.134	0.082	0
Experience Condition	Median Difference Between the Actual Offer and EBR Offer Indicates:			
Inexperienced (32 subjects)	Overrevelation	22	17	3
	Underrevelation	10	15	29
Experienced (16 subjects)	Overrevelation	12	9	0
	Underrevelation	4	7	16
Nash Robots (16 or 29 subjects)	Overrevelation	25	10	4
	Underrevelation	4	6	25
Nash Experienced (16 subjects)	Overrevelation	16	14	1
	Underrevelation	0	2	15

Notes: Overrevelation and underrevelation classification is based on the median paired difference between the actual offers and the EBR offers for the identical value and cost draws. This calculation is based on the value and cost draws with the greatest expected losses from suboptimal offers: values in [\$2.00, \$3.50], and costs in [\$1.50, \$3.00].

irrational traders, or is a widespread phenomenon.[21] The pattern of deviations is relatively consistent across subjects, and is consistent with the conclusions drawn in the text. Overrevelation relative to the EBR was most common when the BNE implies the most underrevelation (column 1). A binomial test rejects the hypothesis that over and underrevelation are equally likely for each dataset in this column (although only at the ten percent significance level in the experienced treatment). Overrevelation relative to the EBR is less common when $k=0.5$ (column 2), but is still more common than underrevelation relative to the EBR in all experience conditions. However, a binomial test rejects the hypothesis that over and underrevelation are equally likely only in the Nash experienced condition shown in the lowest row of the table. Finally, subjects nearly always underreveal compared to the dominant strategy full revelation case (column 3), as they rarely submitted bids above value or asks below cost. Here the binomial test always rejects the hypothesis that over and underrevelation are equally likely.

REFERENCES

Chamberlin, E. (1948): "An Experimental Imperfect Market," *Journal of Political Economy*, 56, 95–108.

Cox, J., and R. Oaxaca (1995): "Is Bidding Behavior Consistent with Bidding Theory for Private Value Auctions?" in *Research in Experimental Economics*, Vol. 6, ed. by R. M. Issac. Greenwich, CT: JAI Press, forthcoming.

[21] It would be preferable to classify subjects into three categories: (i) those who significantly overreveal relative to the EBR; (ii) those who significantly underreveal relative to the EBR; and (iii) those whose offers are insignificantly different from the EBR. Unfortunately, we do not have a sufficient number of observations per subject to provide meaningful statistical tests. For these calculations we subset the data to the value and cost range with substantial payoff consequences of suboptimal offers, so in all cases we observe less than ten offers per subject in each trader role.

COX, J., V. SMITH, AND J. WALKER (1988): "Theory and Individual Behavior of First-Price Auctions," *Journal of Risk and Uncertainty*, 1, 61–99.

CHATTERJEE, K., AND W. SAMUELSON (1983): "Bargaining Under Incomplete Information," *Operations Research*, 31, 835–851.

FRIEDMAN, D. (1992): "Theory and Misbehavior of First Price Auctions: A Comment," *American Economic Review*, 82, 1374–1378.

——— (1993): "How Trading Institutions Affect Financial Market Performance: Some Laboratory Evidence," *Economic Inquiry*, 31, 410–435.

GODE, D., AND S. SUNDER (1993): "Allocative Efficiency of Markets with Zero Intelligence (ZI) Traders: Market as a Partial Substitute for Individual Rationality," *Journal of Political Economy*, 101, 119–137.

GRESIK, T., AND M. SATTERTHWAITE (1989): "The Rate at Which a Simple Market Converges to Efficiency as the Number of Traders Increases: An Asymptotic Result for Optimal Mechanisms," *Journal of Economic Theory*, 48, 304–332.

HARRISON, G. (1989): "Theory and Misbehavior of First-Price Auctions," *American Economic Review*, 79, 749–762.

KAGEL, J. (1994): "Double Auction Markets with Stochastic Supply and Demand Schedules: Clearing House and Continuous Double Auction Trading Mechanisms," unpublished manuscript, Department of Economics, University of Pittsburgh.

——— (1995): "Auctions: A Survey of Experimental Research," in *The Handbook of Experimental Economics*, ed. by J. Kagel and A. Roth. Princeton, NJ: Princeton University Press, pp. 501–585.

KAGEL, J., AND D. LEVIN (1993): "Independent Private Value Auctions: Bidder Behavior in First-, Second- and Third-Price Auctions with Varying Numbers of Bidders," *Economic Journal*, 103, 868–879.

KAGEL, J., AND A. ROTH (1992): "Theory and Misbehavior in First-Price Auctions: Comment," *American Economic Review*, 82, 1379–1391.

KAGEL, J., AND W. VOGT (1993): "The Buyers Bid Double Auction: Preliminary Experimental Results," in *The Double Auction Market: Institutions, Theories and Evidence*, ed. by D. Friedman and J. Rust. Redwood City, CA: Addison-Wesley, pp. 285–305.

MCCABE, K., S. RASSENTI, AND V. SMITH (1993): "Designing a Uniform Price Double Auction: An Experimental Evaluation," in *The Double Auction Market: Institutions, Theories and Evidence*, ed. by D. Friedman and J. Rust. Redwood City, CA: Addison-Wesley, pp. 307–332.

RUSTICHINI, A., M. SATTERTHWAITE, AND S. WILLIAMS (1990): "Convergence to Price-Taking Behavior in a Simple Market," Northwestern University Discussion Paper 914.

——— (1994): "Convergence to Efficiency in a Simple Market with Incomplete Information," *Econometrica*, 62, 1041–1063.

SATTERTHWAITE, M., AND S. WILLIAMS (1989): "The Rate of Convergence to Efficiency in the Buyer's Bid Double Auction as the Market Becomes Large," *Review of Economic Studies*, 56, 477–498.

——— (1993): "The Bayesian Theory of the k-Double Auction," in *The Double Auction Market: Institutions, Theories and Evidence*, ed. by D. Friedman and J. Rust. Redwood City, CA: Addison-Wesley, pp. 99–123.

SELTEN, R., AND J. BUCHTA (1994): "Experimental Sealed Bid First Price Auctions with Directly Observed Bid Functions," University of Bonn Discussion Paper B-270.

SMITH, V. (1962): "An Experimental Study of Competitive Market Behavior," *Journal of Political Economy*, 70, 111–137.

SMITH, V., AND J. WALKER (1993): "Rewards, Experience and Decision Costs in First-Price Auctions," *Economic Inquiry*, 31, 237–244.

SMITH, V., A. WILLIAMS, W. K. BRATTON, AND M. VANNONI (1982): "Competitive Market Institutions: Double Auctions vs. Sealed Bid-Offer Auctions," *American Economic Review*, 72, 58–77.

VICKREY, W. (1961): "Counterspeculation, Auctions, and Competitive Sealed Tenders," *Journal of Finance*, 16, 8–37.

WILLIAMS, S. (1991): "Existence and Convergence of Equilibria in the Buyer's Bid Double Auction," *Review of Economic Studies*, 58, 351–374.

WILSON, R. (1985): "Incentive Efficiency of Double Auctions," *Econometrica*, 53, 1101–1115.