



Learning in a Laboratory Market with Random Supply and Demand

TIMOTHY N. CASON

Department of Economics, Krannert School of Management, Purdue University, West Lafayette, IN 47907-1310
email: cason@mgmt.purdue.edu

DANIEL FRIEDMAN

Department of Economics, University of California at Santa Cruz, Santa Cruz, CA 95064

Abstract

We propose a simple adaptive learning model to study behavior in the call market. The laboratory environment features buyers and sellers who receive a new random value or cost in each period, so they must learn a strategy that maps these random draws into bids or asks. We focus on buyers' adjustment of the "mark-down" ratio of bids relative to private value and sellers' adjustment of the corresponding "mark-up" ratio of asks relative to private cost. The learning model involves partial adjustment of these ratios towards the ex post optimum each period. The model explains a substantial proportion of the variation in traders' strategies. Parameter estimates indicate strong recency effects and negligible autonomous trend, but strongly asymmetric response to different kinds of ex post error. The asymmetry is only slightly attenuated in "observational learning" from other traders' ex post errors. Simulations show that the model can account for the main systematic deviations from equilibrium predictions observed in this market institution and environment.

Keywords: experiment, call market, auction, bidding

JEL Classification: D44, D83, C92

1. Introduction

Market institutions are the heart of modern economies. They bring together numerous self-interested, privately informed people, and somehow produce mutually beneficial trades. A wide variety of market institutions, ranging from the continuous double auction to brokered search, operate in a wide variety of environments. Economists would like to know which market institutions perform best in particular environments. More fundamentally, economists need to understand more precisely how market institutions work.

The standard theoretical approach, called mechanism design, studies the equilibrium properties of institutions. But for a market institution to work well in practice it is not sufficient (and perhaps not even necessary)¹ that its equilibria are efficient. It is at least as important that the institution promotes rapid learning of behavior that supports efficient market outcomes. Unfortunately there is as yet no standard theory of how institutions (market or nonmarket) mediate learning. Our goal is to produce empirical tools and evidence that will guide the construction of such a theory.

We pursue our goal by fitting a simple error-driven learning model to behavior observed in the simplest viable market institution, the call market, also known as the Clearinghouse or the two-sided sealed auction. The call market 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 call market 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.² Contingent on his or her own value, each trader must choose an action (e.g., a bid) each period and these actions jointly determine the outcome. Thus each period we have a new episode of price formation based on new private information, and we can observe how traders adjust their contingent actions in response to accumulated experience.

Learning models try to capture the adjustment process and predict the extent to which traders' contingent actions and beliefs eventually reach equilibrium. Several strands of literature, old and new, inform the learning model we develop here. Psychologists following Bush and Mosteller (1955) model subjects' adjustment to experience in simple individual choice tasks. We will be especially interested in the "error-driven" learning processes following Rescorla and Wagner (1972) that deal with complex contingent feedback, and with the recent machine learning models following e.g., McClelland and Rumelhart (1988). The mushrooming literature on learning in games, as summarized in recent texts by Fudenberg and Levine (1998) and Camerer (1999), would also seem relevant. Unfortunately that literature focuses mainly on choices among a few discrete action alternatives rather than among the continuous families of contingent pricing strategies faced by market participants.³ We draw insights from the learning in games literature, but draw elements of our learning model mainly from the psychology and computer science literature.

Our work is also informed by earlier laboratory studies. Kagel (1994) and Cason and Friedman (1997) test the equilibrium theory of the call market summarized in Satterthwaite and Williams (1993). They find that the theory is superior to simple alternatives in predicting the general level of prices and efficiency, but does a poor job of predicting individual bid and ask functions. Traders in the laboratory tend to reveal value and cost more fully than the equilibrium theory predicts. Based on some ad hoc regressions, Cason and Friedman (1997) conjecture that the anomalous behavior arises from a biased learning process. The structural learning model presented below uses the same data and confirms the conjecture.

Daniel et al. (1998) report laboratory tests of equilibrium theory in the extreme case of a call market with only one buyer and one seller, i.e., bilateral monopoly. Using very asymmetric distributions for buyer values and seller costs, they find that sellers are less aggressive and buyers more aggressive than the equilibrium benchmarks. They show that a learning model can account for many of the regularities. Their learning model was developed independently from ours and differs in many ways, but also models the adjustment of an aggressiveness parameter in response to different sorts of experience.

Friedman and Ostroy (1995) investigate the call market and related trading institutions in a stationary environment with divisible goods. They argue that the data are best explained by an unmodelled learning process that eventually induces traders to offer at the competitive equilibrium price (a form of price underrevelation) but to fully reveal quantity. They demonstrate that such strategies are Nash equilibria that support competitive equilibrium outcomes. Cason (1993) presents a simple adaptive learning algorithm for the call market in a stationary environment that is based on Gode and Sunder's (1993) Zero Intelligence (ZI) algorithm. Simulations of the algorithm produce data that closely resembles laboratory markets populated with human subjects. Selten and Buchta (1998) introduce an informal "learning direction theory" that explains changes in bid functions in the first-price private-values auction, a trading institution with incentives similar to the call market. Garvin and Kagel (1994) find that subjects learn to avoid overbidding in first-price common-values auctions from their personal experience of the "winner's curse" and also from observing the fate of other players who overbid.

The remainder of this paper is organized as follows. Section 2 summarizes the rules of the call market institution, traders' equilibrium strategies, and the experiment design. Section 3 presents the learning model, which involves the adjustment of a single variable for each buyer (and each seller) indicating the degree to which bids reveal true value (and asks reveal true cost). Section 4 collects the results. The learning model, though parsimonious, explains a substantial proportion of the variation in traders' strategies. Parameter estimates suggest little or no autonomous time trend, and confirm that traders respond asymmetrically to the qualitatively different forms ex post error. The parameter estimates also indicate strong recency effects, as in Cournot belief learning, as opposed to fictitious play or simple Bayesian learning. Like Garvin and Kagel (1994), we find significant evidence that traders learn from other traders' ex post errors as well as from their own errors. A simulation of the estimated model confirms that it reproduces the main departures from equilibrium behavior. Section 5 summarizes the results and offers some concluding remarks.

2. The call market institution, equilibrium theory and experimental design

2.1. Institution pricing rules and equilibrium

The call market 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. With n buyers and n sellers, it turns out that p_l and p_u are respectively the n th and $(n+1)$ th lowest offers, counting both bids and asks. The experiment uses $n = 4$, so the price-setting offers are the fourth and fifth lowest, denoted s_4 and s_5 .

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 $[p_l, p_u]$. The experiment implements three pricing rules in different sessions, denoted $p_{k=1.0}$, $p_{k=0.5}$ and $p_{k=0}$ in figure 1.

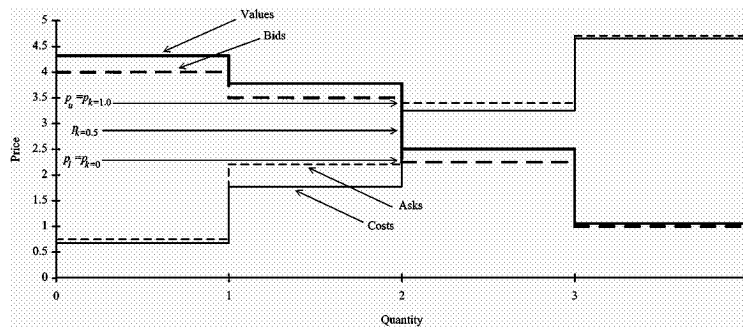


Figure 1. Example price determination for three price rules.

In Bayesian Nash Equilibrium, 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 call market for sellers when $k = 1$, because then p^* is always set by a buyer or by a non-transacting seller. Similarly 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. Rustichini et al. (1994) show that in Bayesian Nash Equilibrium both buyers and sellers reduce their bids and asks as k increases, as illustrated in figure 2. Cason and Friedman (1997) show that this bid and ask ordering holds for more general trader beliefs than in the Bayesian Nash Equilibrium. Although in theory the traders' strategies are sensitive to the pricing rule, Cason and Friedman (1997) find that laboratory subjects' offers do not respond to the differing incentives provided by the pricing rule k (see also Table 2 below). Our empirical work therefore will pool the data from the three pricing rules.

2.2. Experimental design

As reported in Cason and Friedman (1997), the experiment consists of 17 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 call market session. Subjects were randomly assigned a computer and trader position, and instructions were read orally while subjects followed along on their own written copy. 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 4 buyers and 4 sellers who 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

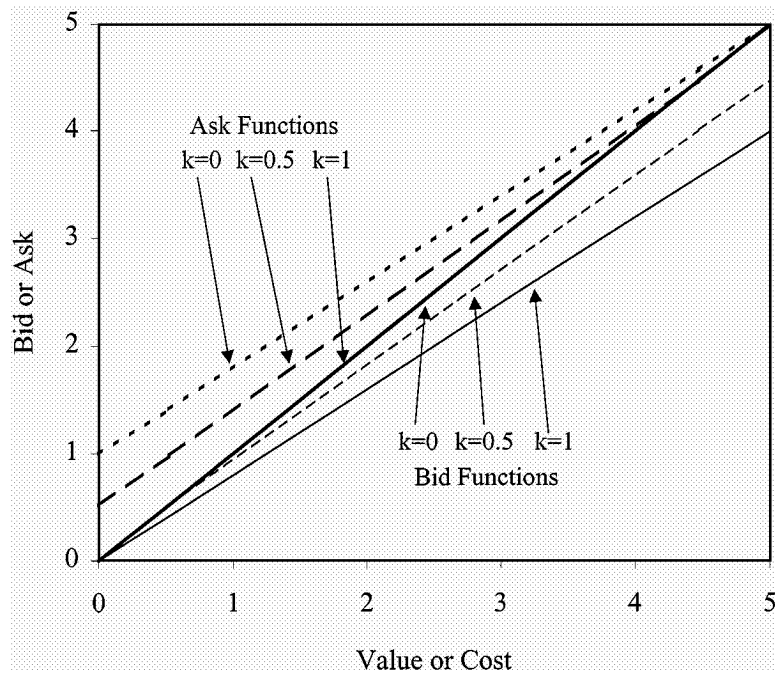


Figure 2. Approximate risk neutral BNE bid and ask functions for the three k treatments (4 buyers, 4 sellers).

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 all subjects received a summary of every participant's actions and outcomes, using tables that presented the bids, asks, values, costs and profits of all traders. We provided this complete information to increase traders' opportunity to learn about their rivals' strategies and to provide them with sufficient information to calculate optimal offers ex post. The instructions in the Appendix provide additional details of the procedures and the operation of the market program.

Because we are studying learning we differentiate the data based on trader experience. Trader experience is a composite treatment with four levels, including the usual two levels referred to below as Inexperienced (all 8 traders are inexperienced humans) and Experienced (all 8 human traders who previously had participated in an inexperienced session). The other two levels involve robot traders (i.e., computer algorithms) programmed to use the Bayesian Nash Equilibrium bid or ask functions graphed in figure 2; instructions include the relevant graph. In the Nash Robots treatment each human is inexperienced in any call market and

Table 1. Summary of laboratory sessions.

Session name	k -Treatment	Experience treatment:			Number of periods
		Opponents	Experience	Label	
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

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].

^aSession k1-rob-14 employed only five subjects, each competing against 7 robot opponents. All other sessions employed 8 human subjects.

faces 7 such robots. The last 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.⁴

Table 1 summarizes the 17 sessions. The 11 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.

3. Learning models

Traders in the call market have a formidable learning task. They implicitly must estimate the bid and ask functions used by other traders and must seek a best response, a bid function

$b = B(v)$ for a buyer or an ask function $a = A(c)$ for a seller. Arbitrary functions are quite difficult to learn, but there are good reasons to restrict attention to more tractable families of functions. First, it is straightforward to show that non-monotonic bid and ask functions are always dominated by increasing functions in this setting. Furthermore, as shown in figure 2, the equilibrium bid and ask functions are always approximately (and often precisely) linear in our environment with values v and cost c uniformly distributed on the interval $[0, M]$. Indeed, Appendix A in Cason and Friedman (1997) shows that precisely linear functions are the best response to a variety of other, nonequilibrium behavior by other traders. For example, the bid function $b = \alpha v$ is the unique best response to truth-telling by other traders (i.e., to $b = v$ and $a = c$) in a k -call market with $m > 1$ sellers and $m - 1$ other buyers, where $\alpha = m/(m + k)$. Linear functions are also natural rules of thumb (e.g., mark up cost by $x\%$), and they are relatively tractable. For all these reasons—tractability, intuitive appeal, and theoretical salience—we will henceforth assume that traders use linear bid and ask functions.

The trader's decision problem now can be posed as how much to reveal of her true willingness to transact. For a buyer, write $B(v) = \alpha v$, and by strict analogy write $A(c) = \alpha c + (1 - \alpha)M$ for sellers. Full revelation is represented by $\alpha = 1$, and the usual partial revelation by $0 < \alpha < 1$. Of course, complete unwillingness to transact ($\alpha = 0$) and overrevelation ($\alpha > 1$) are logical possibilities but they are dominated by partial or full revelation. The call market learning problem now comes down to finding how the revelation parameter α responds to experience.

Psychologists' error-driven learning approach (e.g., Rescorla and Wagner, 1972; Gluck and Bower, 1988; McClelland and Rumelhart, 1988; Friedman et al., 1995) seems especially relevant since it describes adjustment of a continuous choice variable given continuous feedback. At the end of each trial (trading period) t , each trader has feedback on the "correct" value α_t^o , i.e., each has the information required to see ex post what the most profitable value of α would have been. The trader uses this ex post optimal α_t^o to update her belief regarding the correct revelation parameter $\hat{\alpha}_t$:

$$\hat{\alpha}_t = [1 - \delta(t)]\hat{\alpha}_{t-1} + \delta(t)\alpha_t^o, \quad (1)$$

where $\delta(t)$ is a function describing the learning rate. We make the decision rule transparent, in that subjects adjust fully to the current belief $\hat{\alpha}_t$ when selecting the next period revelation parameter:

$$\alpha_{t+1} = \hat{\alpha}_t. \quad (2)$$

Combining Eqs. (1) and (2) and rearranging, we have the following error-driven learning equation:

$$\alpha_{t+1} = \alpha_t + \delta(t)e_t, \quad (3)$$

where the error $e_t = \alpha_t^o - \alpha_t$ is the difference between the ex post optimal α_t^o and the α_t actually used.

The additive form of Eq. (3) is traditional in the psychology literature, but there are two reasons to consider an alternative multiplicative form. First, our revelation parameter α is by definition a nonnegative ratio. Multiplicative adjustment automatically preserves the nonnegativity constraint and generally seems better suited for working with a ratio. Second, recent computer science studies of learning, e.g., Kivinen and Warmuth (1995), demonstrate the objective superiority of a multiplicative form (“exponential gradient updating” in their terminology) over the additive form (“ordinary gradient descent”) in complex environments. To check for specification error, we include a residual time trend or drift parameter η and obtain the multiplicative learning equation

$$\frac{\alpha_{t+1}}{\alpha_t} = e^{\eta} \frac{[\alpha_t^o]^\delta}{[\alpha_t]}, \quad (4)$$

or, taking logs,

$$[\ln(\alpha_{t+1}) - \ln(\alpha_t)] = \eta + [\ln(\alpha_t^o) - \ln(\alpha_t)]\delta(t). \quad (5)$$

As a robustness check, we also estimate the additive form

$$[\alpha_{t+1} - \alpha_t] = \eta + [\alpha_t^o - \alpha_t]\delta(t). \quad (6)$$

Equation (6) can be used to model adaptive learning of any continuous action α , market or non-market, and (5) can be used whenever we can safely assume $\alpha > 0$. We do not impose the additional constraint that α does not exceed 1.

Several issues arise in applying (5) or (6) to our market data. First, based on previous results we must allow for the possibility that some types of errors have greater impact than others. The fundamental tradeoff in increasing α is between a higher probability of transacting and lower profit conditional on transacting. Sometimes the ex post error is from missing out on a profitable transaction opportunity (which we shall call error type m) and sometimes it is from adversely affecting the price of a realized transaction (error type p). Using a simple regression suggested by Selten and Buchta’s (1998) directional learning theory, Cason and Friedman (1997) find that traders seem to respond much more strongly to type m than type p errors. Therefore we want to estimate separate structural parameters δ_m and δ_p to assess the two types of ex post errors.

Second, we want to account for recency effects. Some popular learning models, e.g., back propagation as in McClelland and Rumelhart (1988) or the Cournot learning model mentioned below, assume that errors late in a session have the same impact on α as errors early in a session. It would be dangerous to impose such an assumption in our error-driven learning model since errors of a particular sort may be infrequent. Psychologists’ traditional power law of practice assumes that the impact diminishes with experience as measured by the number N of relevant previous trials. Indeed, models such as the fictitious play model mentioned below or simple Bayesian updating assume that the impact is proportional to N^{-1} . We include a parameter β to detect recency effects.

Thus our specification of the learning rate is

$$\delta(t) = \delta_m D_{mt} N_{mt}^{-\beta} + \delta_p D_{pt} N_{pt}^{-\beta}, \quad (7)$$

where D_{mt} and D_{pt} are indicator variables equal to 1 for error type m and p , respectively; and N_{mt} and N_{pt} are, respectively, the number of times an error of type m and p has been made by the subject, including the current period. The estimated η term in Eqs. (5) and (6) captures the fixed trend in the α , independent of the error feedback; it will be zero if the model is correctly specified.

In an early version of this paper (Cason and Friedman, 1995), we considered several additional complications: (a) weighting more recent evidence more strongly in the $\hat{\alpha}_t$ update; (b) permitting previous value or cost draws nearer the current value or cost draw to have more impact on the $\hat{\alpha}_t$ update; and (c) updating $\hat{\alpha}_t$ more strongly for errors with higher expected payoff consequences. The current specification drops (b) but captures (a) via the β parameter and captures (c) via weighted least squares estimation, as noted in the next section. The current specification has about the same explanatory power as the earlier version but is much more parsimonious.

The estimated parameters δ_m , δ_p and β have a natural interpretation in terms of classic models of expectation formation; see Cheung and Friedman (1997) for a parallel interpretation in the context of learning in games. According to the Cournot model of expectation formation, α adjusts completely to the previous period's ex post optimum: $\alpha_{t+1} = \alpha_t^o$. This implies the restrictions $\delta_m = \delta_p = 1$ and $\beta = 0$. The fictitious play model of expectation formation sets α_{t+1} equal to the average ex post optimum α of all previous periods. This implies the restrictions $\delta_m = \delta_p = 1$ and $\beta = 1$. Fictitious play implicitly assumes that other players keep their behavior constant, which is precisely correct in the Robot Opponents treatment, and may be a reasonable approximation in other treatments. In general, the β coefficient can be thought of as memory length, and $0 < \beta < 1$ indicates a stronger adaptive response to more recent observations. Recall that the parameters δ_m and δ_p are intended to capture possibly asymmetric error impacts, and that the Cason and Friedman (1997) finding of stronger response to missed transactions suggests that $\delta_m > \delta_p$.

A final implementation issues arises from the definition of the ex post optimum. Recall that the price is a k -weighted combination of the fourth and fifth highest offers. A buyer optimizes ex post by bidding one penny above s_4 , the fourth highest offer of the other 7 traders, whenever that bid would not exceed her current value. Tying s_4 keeps the price as low as possible and the penny breaks ties so the bidder is sure to transact. Similarly, the ex post optimal ask is one penny below s_4 , when s_4 is above the seller's cost. Ignoring the penny, we conclude that a current period ex post optimum revelation ratio is $\alpha_t^o = s_4/v_t$ for buyers and $\alpha_t^o = (M - s_4)/(M - c_t)$ for sellers, whenever the trader could have done better.

Of course, sometimes the trader could not have done better that period given the actions of the other 7 traders. The ex post optimum is not to transact when a buyer's value is too low or a seller's cost is too high to transact profitably. Also, many buyers and sellers transact at prices that they could not have profitably altered. In such cases there is no ex post error and, consistent with Eq. (3), we exclude these cases when estimating the model.⁵ Note also that in the dominant full revelation strategy cases of buyers with $k = 0$ and sellers with $k = 1$, the ex post optimum is not uniquely defined by the α_t^o expressions given above. In these

cases we use the midpoint of the range of optima, i.e., we set α_t^o midway between s_4/v_t and 1.0 for buyers, and midway between $(M - s_4)/(M - c_t)$ and 1.0 for sellers.

4. Results

To avoid unnecessary overlap with Cason and Friedman (1997), we present only a brief overview of the raw bid and ask data in the first subsection, together with a summary of the revelation ratios. The next two subsections present the main results from fitting the learning model, and the last subsection uses a simulation to check whether the learning model can explain the overall change observed in the aggregate revelation ratios.

4.1. Overview

Figure 3 illustrates the nature of the bid and ask data by showing a scatterplot for session k5-hum-8x, a fairly typical 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 Bayesian Nash Equilibrium 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 equilibrium bid and ask functions.

The learning analysis focuses on the revelation ratio α , so before proceeding it is instructive to summarize how this ratio varies across datasets and over time. Table 2 presents the median separately for the k -treatments, for buyers and sellers, and for each experience condition. The top row shows α in risk neutral Bayesian Nash Equilibrium. Consistent

Table 2. Median revelation (α) by experience, k -treatment and trader role.

Experience condition:	Buyers			Sellers		
	$k = 0$	$k = 0.5$	$k = 1.0$	$k = 0$	$k = 0.5$	$k = 1.0$
Risk Neutral Bayesian Nash Equilibrium	1.0	0.908	0.8	0.8	0.908	1.0
Inexperienced	0.872 [0.238]	0.929 [0.172]	0.934 [0.186]	0.907 [0.251]	0.903 [0.257]	0.952 [0.217]
Experienced	0.928 [0.152]	0.947 [0.152]	0.955 [0.137]	0.932 [0.197]	0.964 [0.135]	0.959 [0.253]
Nash Robots	0.951 [0.118]	0.930 [0.177]	0.876 [0.230]	0.915 [0.242]	0.935 [0.230]	0.947 [0.239]
Nash Experienced	0.993 [0.054]	0.996 [0.033]	0.973 [0.085]	0.993 [0.023]	0.997 [0.019]	0.991 [0.065]

Inter-quartile range shown in brackets to represent dispersion.

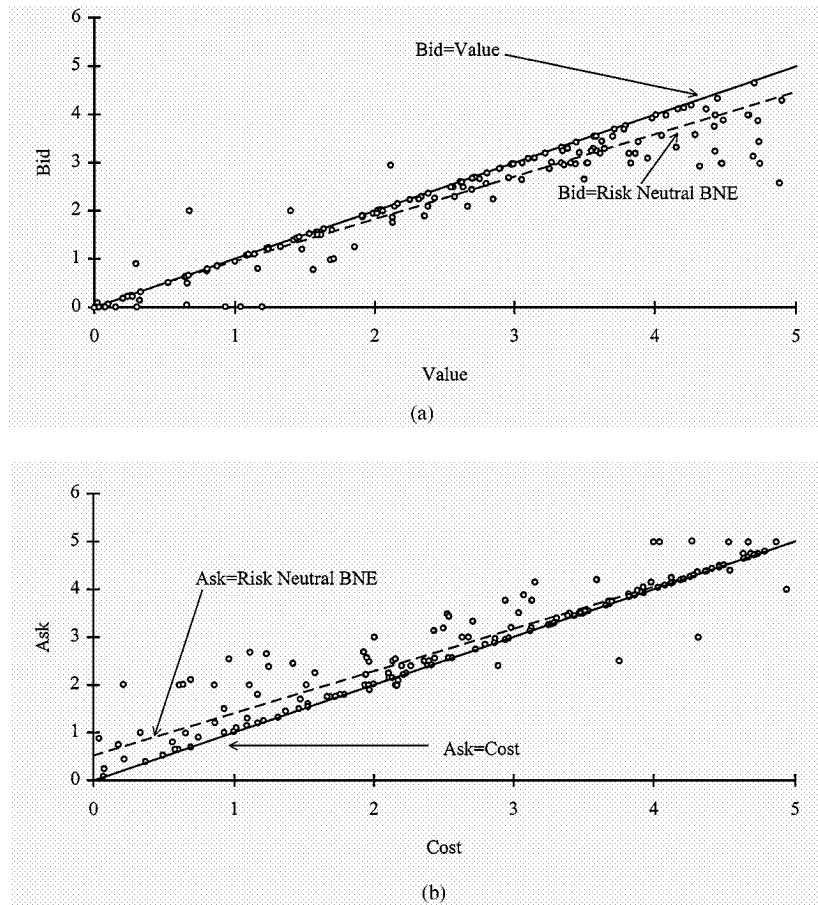


Figure 3. (a) Bids in session k5-hum-8x ($k = 0.5$, experienced subjects); (b) Asks in session k5-hum-8x ($k = 0.5$, experienced subjects).

with a major finding of Cason and Friedman (1997), the table shows that traders' median α generally does not track the equilibrium prediction as the pricing rule k changes. The median α shifts in the predicted direction only in the Nash Robots experience condition, and in all conditions the shifts in α are roughly an order of magnitude smaller than predicted.

Table 2 also highlights the Cason and Friedman (1997) finding that revelation increases significantly with experience; the entries in the experienced rows consistently exceed the corresponding entries in the other two rows. This experience effect can be seen more clearly in figure 4. Combining revelation ratios α across buyers and sellers and across k values, we plot the empirical cumulative distribution of α in periods 1 and 2 of the inexperienced treatments and the empirical cumulative distribution of α in the last two

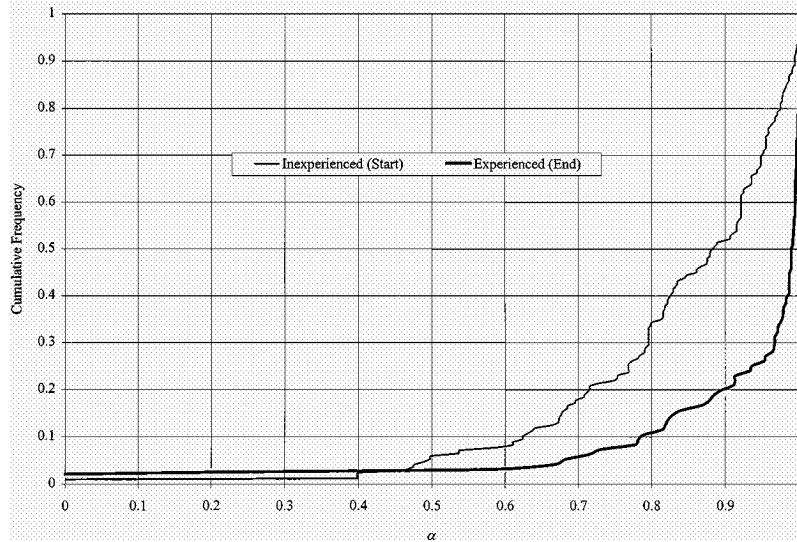


Figure 4. Beginning and ending α distributions.

periods of the experienced treatments. A primary task for the learning model analyzed in the next subsection is to explain the dramatic shift of the distribution towards $\alpha = 1.0$, or full revelation.

4.2. Model estimates

The adaptive, error-driven learning model presented in Section 3 is summarized by substituting Eq. (7) into Eq. (5):

$$[\ln(\alpha_{t+1}) - \ln(\alpha_t)] = \eta + [\ln(\alpha_t^o) - \ln(\alpha_t)] \{ \delta_m D_{mt} N_{mt}^{-\beta} + \delta_p D_{pt} N_{pt}^{-\beta} \}. \quad (8)$$

Equation (8) is estimated using weighted nonlinear least squares.⁶ States with no error constitute 3966 of the total 4470 offer observations and are omitted from the sample. Because these error observations are excluded and because there is no constant term in the expression in braces $\{ \}$, the δ_m and δ_p estimates capture directly the impact of missing a trade because of excessive underrevelation ($D_{mt} = 1$) and the impact of adversely affecting price ($D_{pt} = 1$). In both cases, recall that positive estimates indicate adjustment *towards* the ex post optimal α , an increase in α in event m and a decrease in α in event p . Estimates $\delta_m = 1$ or $\delta_p = 1$ indicate immediate adjustment to the ex post optimum, and estimates between 0 and 1 indicate partial adjustment. A positive η indicates a positive trend in the revelation parameter α independent of the feedback.

Table 3 presents the coefficient estimates for buyers (Panel A), sellers (Panel B) and pooled (Panel C).⁷ In some Inexperienced datasets the η estimates are significantly positive

Table 3. Multiplicative specification, nonlinear weighted least squares estimates.

Experience	η	δ_m	δ_p	β	Adj. R^2	Miss trade obs.	Price impact obs.	$H_0: \delta_m = \delta_p$ [χ^2 (1 d.f.)]
<i>Panel A: buyers</i>								
Inexperienced	0.01 (0.06)	0.61 (0.40)	-0.18 (0.26)	0.01 (0.62)	0.07	27	48	3.02 ($p = 0.08$)
Inexperienced ^a	0.05 (0.04)	0.55 ^d (0.28)	-0.19 (0.20)	0.11 (0.51)	0.10	27	46	5.05 ($p < 0.025$)
Experienced	0.02 (0.02)	3.35 ^b (0.90)	0.39 ^d (0.21)	8.17 (135.56)	0.28	12	38	13.04 ($p < 0.01$)
Robot opponents	-0.02 (0.06)	0.80 ^b (0.19)	-0.54 (0.36)	-0.63 ^b (0.23)	0.51	32	34	11.26 ($p < 0.01$)
Robot experienced	-0.09 (0.08)	0.23 (0.82)	-0.002 (0.014)	-2.60 (3.56)	0.05	10	33	1.07 ($p = 0.30$)
<i>Panel B: sellers</i>								
Inexperienced	0.10 ^b (0.04)	0.78 ^b (0.16)	0.22 (0.24)	0.21 (0.29)	0.28	50	58	3.15 ($p = 0.08$)
Experienced	-0.07 (0.04)	1.24 ^b (0.26)	-0.38 ^c (0.16)	-0.01 (0.23)	0.29	22	45	18.05 ($p < 0.01$)
Robot opponents	0.10 ^b (0.03)	0.35 ^b (0.13)	-0.30 ^d (0.16)	-1.13 ^c (0.48)	0.42	19	24	12.83 ($p < 0.01$)
Robot experienced	-0.01 (0.02)	1.06 ^b (0.05)	-0.02 (0.06)	0.18 (0.15)	0.90	10	41	68.24 ($p < 0.01$)
<i>Panel C: pooled across buyers and sellers</i>								
Inexperienced ^a	0.07 ^c (0.03)	0.75 ^b (0.14)	-0.03 (0.15)	0.19 (0.25)	0.21	77	104	11.19 ($p < 0.01$)
Experienced	-0.03 (0.03)	1.20 ^b (0.23)	-0.24 ^c (0.12)	-0.01 (0.22)	0.24	34	83	25.24 ($p < 0.01$)
Robot opponents	0.01 (0.04)	0.71 ^b (0.14)	-0.46 ^c (0.23)	-0.72 ^b (0.19)	0.49	41	58	19.18 ($p < 0.01$)
Robot experienced	-0.06 (0.05)	1.09 ^b (0.19)	-0.13 (0.16)	0.00 (0.45)	0.29	20	94	16.15 ($p < 0.01$)

$[\ln(\alpha_{t+1}) - \ln(\alpha_t)] = \eta + [\ln(\alpha_t^o) - \ln(\alpha_t)]\{\delta_m D_{mt} N_{mt}^{-\beta} + \delta_p D_{pt} N_{pt}^{-\beta}\}$; standard errors are in parentheses.

Each observation is weighted by the expected payoff costs of an offer deviation from the optimal offer, as explained in text.

^aTwo outlier observations excluded, as explained in note 6.

^bDenote significantly different from zero at 1% (all 2-tailed tests).

^cDenote significantly different from zero at 5% (all 2-tailed tests).

^dDenote significantly different from zero at 10% (all 2-tailed tests).

but elsewhere they are insignificant.⁸ In every dataset the estimates indicate $\delta_m > \delta_p$, and in most datasets the χ^2 test reported in the rightmost column indicates that the difference is highly significant. Moreover, the δ_p estimates are either insignificant or *negative*, indicating adjustment *away from* the ex post optimum (and toward greater revelation). The pooled δ_m estimates center near the full adjustment towards the ex post optimum (again towards greater revelation).

Table 4. Additive specification, nonlinear weighted least squares estimates.

Experience	η	δ_m	δ_p	β	Adj. R^2	Miss trade obs.	Price impact obs.	$H_0: \delta_m = \delta_p$ [χ^2 (1 d.f.)]
<i>Panel A: buyers</i>								
Inexperienced	0.44 ^b (0.07)	-1.27 ^d (0.74)	5.55 ^b (0.42)	0.88 ^b (0.11)	0.58	27	48	30.19 ($p < 0.01$)
Inexperienced ^a	0.06 (0.04)	0.37 (0.40)	-0.06 (0.23)	-0.13 (0.99)	0.04	27	46	1.10 ($p = 0.29$)
Experienced	0.02 (0.03)	2.61 ^c (1.17)	0.12 (0.28)	1.22 (1.28)	0.13	12	38	3.93 ($p = 0.05$)
Robot opponents	-0.01 (0.04)	0.79 ^b (0.21)	-0.63 ^d (0.37)	-0.80 ^b (0.28)	0.28	32	34	9.78 ($p < 0.01$)
Robot experienced	-0.03 (0.03)	0.44 (0.86)	0.001 (0.014)	-1.74 (1.95)	0.18	10	33	7.14 ($p < 0.01$)
<i>Panel B: sellers</i>								
Inexperienced	0.12 ^b (0.04)	0.52 ^d (0.28)	0.41 (0.32)	-0.12 (0.53)	0.11	50	59	0.07 ($p = 0.79$)
Experienced	0.04 (0.07)	0.61 (0.48)	0.37 (0.32)	-0.42 (0.60)	0.13	22	46	0.19 ($p = 0.67$)
Robot opponents	0.09 ^b (0.03)	0.21 (0.14)	-0.26 (0.24)	-1.66 ^d (0.87)	0.24	20	24	6.71 ($p < 0.01$)
Robot experienced	0.01 (0.02)	1.03 ^b (0.10)	0.03 (0.07)	0.21 (0.19)	0.79	10	41	35.80 ($p < 0.01$)
<i>Panel C: pooled across buyers and sellers</i>								
Inexperienced ^a	0.09 ^b (0.03)	0.50 ^c (0.24)	0.17 (0.19)	-0.07 (0.50)	0.07	77	105	1.20 ($p = 0.27$)
Experienced	0.04 (0.04)	0.75 ^d (0.40)	0.24 (0.20)	-0.27 (0.47)	0.10	34	84	1.36 ($p = 0.24$)
Robot opponents	0.02 (0.03)	0.65 ^b (0.16)	-0.53 ^c (0.26)	-0.86 ^b (0.26)	0.25	52	58	15.26 ($p < 0.01$)
Robot experienced	-0.01 (0.02)	1.02 ^b (0.17)	0.00 (0.09)	-0.04 (0.32)	0.43	20	74	25.69 ($p < 0.01$)

$[\alpha_{t+1} - \alpha_t] = \eta + [\alpha_t^o - \alpha_t] \{ \delta_m D_{mt} N_{mt}^{-\beta} + \delta_p D_{pt} N_{pt}^{-\beta} \}$; standard errors in parentheses.

Each observation is weighted by the expected payoff costs of an offer deviation from the optimal offer, as explained in text.

^aTwo outlier observations excluded, as explained in note 6.

^bDenote significantly different from zero at 1% (all 2-tailed tests).

^cDenote significantly different from zero at 5% (all 2-tailed tests).

^dDenote significantly different from zero at 10% (all 2-tailed tests).

We also estimated Eq. (8) separately for each of the 17 sessions to provide an alternative formal test of the hypothesis that $\delta_m > \delta_p$ against the null $\delta_m = \delta_p$. Of the resulting 17 statistically independent estimates of the pair (δ_m, δ_p) , the ordering in 16 pairs is $\delta_m > \delta_p$. A binomial test strongly rejects ($p < 0.001$) the null hypothesis that the orderings are equally

likely in favor of the alternative hypothesis that traders respond more strongly to m -errors than to p -errors.

The results shown in Table 3 also indicate that in most cases we cannot reject the Cournot (short memory) hypothesis $\beta = 0$. A puzzling exception is the negative estimate for the Robot Opponents data. Our prior belief, noted above, is that larger β (indeed $\beta = 1$) is more plausible in this environment.

Table 4 presents corresponding results for the alternative additive (rather than multiplicative) specification

$$[\alpha_{t+1} - \alpha_t] = \eta + [\alpha_t^o - \alpha_t] \{ \delta_m D_{mt} N_{mt}^{-\beta} + \delta_p D_{pt} N_{pt}^{-\beta} \} \quad (9)$$

obtained by inserting Eq. (7) into Eq. (6). Table 4 includes one more data point in each of three datasets (in each case, a negative α arising from an ask above $M = 4.99$). The results are mostly similar to those of Table 3 but a bit more erratic; for example, the additive specification of Table 4 is more sensitive to the outlier observations included in the top row of estimates.⁹

In both Tables 3 and 4 we pool across individual traders. It is conceivable that the conclusions regarding the difference in the adjustment parameters δ_m and δ_p are caused by a subset of traders who react more strongly to missing a profitable trade, while other traders do not. Based in part on previous work (e.g., Cheung and Friedman, 1997; Friedman et al.; 1995; Friedman and Massaro, 1996), we had planned to estimate this learning model for individuals to test this conjecture and evaluate learning heterogeneity. However, each subject in each (buyer or seller) role makes only a few offers—roughly 15 when inexperienced and 20 when experienced—and errors are infrequent—only 14 of the 170 inexperienced subjects made at least 5 errors, and only 12 of the 96 experienced subjects made at least 5 errors. Hence reliable estimates for individual subjects are not feasible with our data.

4.3. Observational learning

Garvin and Kagel (1994) define a discount rate for bidding in common-value auctions that is in some ways analogous to our revelation ratio α . Their data analysis indicates that subjects adjust the discount rate in the appropriate direction when they personally lose money by overpaying for the auctioned good. Perhaps more surprisingly, subjects adjust almost as much when they see that they would have lost money had they applied their own discount rate to the auction winner's signal.

Our experiment also provides subjects with the opportunity to learn from observing the experience of other subjects. The complete ex post information allows each seller to see what the ex post optimum revelation ratio was for the other three sellers as well as her own ex post optimum, and similarly for buyers.¹⁰ Perhaps traders will adjust α towards the ex post optimum α_t^{oo} observed for another trader of her own type as well as towards her own personal ex post optimum. This possibility is also reminiscent of the Camerer and Ho (1999) parameter for discounting payoffs that can be inferred but that are not personally experienced.

Given the empirical results of the previous subsection, our specification of observational learning should allow for separate coefficients for other traders m -errors and p -errors, but can be streamlined by dropping the time trend and recency parameters. Therefore we estimate the equation

$$\begin{aligned} [\ln(\alpha_{t+1}) - \ln(\alpha_t)] = & [\ln(\alpha_t^o) - \ln(\alpha_t)]\{\delta_m D_{mt} + \delta_p D_{pt}\} \\ & + [\ln(\alpha_t^{oo}) - \ln(\alpha_t)]\{\gamma_m D_{mt}^o + \gamma_p D_{pt}^o\}. \end{aligned} \quad (10)$$

The number of observations here is four times as large as in the previous models because each error also leads to three “observational” errors that can be used by the other three subjects in the same (buyer or seller) role. We also estimate an additive specification analogous to Eq. (9). Our expectation was that the estimated observational learning coefficient would be smaller (perhaps much smaller) than the direct learning coefficient for missed trades, $0 \leq \gamma_m < \delta_m$, and that the observational learning coefficient γ_p for pricing errors would not be significant.

Table 5 reports our findings. The estimates are as expected for the robot experienced treatment. The estimates for the other three treatments are much better than we expected: Both observational learning coefficients are large and significant in all but the robot experienced treatment. The estimated missed trades coefficients are ordered $\gamma_m < \delta_m$ as expected, but particularly in the inexperienced data (both human and robot opponents) they are similar in size. Consistent with the direct learning results, we also find a stronger impact for observational m -errors than for observational p -errors ($\gamma_m > \gamma_p$), but unlike their direct learning counterparts, the observational p -errors have a significant positive effect. These results are robust to inclusion of a trend parameter η .

4.4. Model simulation

The model estimates in Tables 3 and 4 clearly point towards larger alphas and greater revelation in later periods. The positive δ_m estimates indicate that the typical subject increases revelation after missing out on a profitable trade. The negative (or roughly zero) estimates of δ_p also indicate an increase (or no change) in revelation following an adverse impact on price because the movement is away from the ex post optimum of less revelation. But are the effects quantitatively of the right magnitude to account for traders’ observed tendency towards full revelation?

To answer that question, we simulate subject behavior according to the learning model using parameter values fitted to (approximately) two-thirds of the sessions, randomly chosen. We then compare the final period simulation results to the actual final period results of the other (out-of-sample) one-third of the sessions. Following a standard convention, e.g., Roth and Erev (1995), the simulation begins with random draws from the empirical initial α distribution, viz., the Inexperienced (Start) cumulative distribution function shown in figure 4. The simulated subjects participated in 15 call auction periods and used the Inexperienced (or Robot Opponent) estimates to update α , and then participated in 20 call auction periods and used the Experienced (or Robot Experienced) estimates to update α .¹¹

Table 5. Observational learning.

Experience	δ_m	δ_p	γ_m	γ_p	Adj. R^2
<i>Panel A: multiplicative specification, weighted least squares estimates (Pooled across buyers and sellers)</i>					
	$[\ln(\alpha_{t+1}) - \ln(\alpha_t)] = [\ln(\alpha_t^o) - \ln(\alpha_t)]\{\delta_m D_{mt} + \delta_p D_{pt}\} + [\ln(\alpha_t^{oo}) - \ln(\alpha_t)]\{\gamma_m D_{mt}^o + \gamma_p D_{pt}^o\}$				
Inexperienced ^a	0.82 ^b (0.11)	-0.22 (0.13)	0.75 ^b (0.05)	0.17 ^b (0.03)	0.29
Experienced	1.09 ^b (0.17)	-0.16 (0.08)	0.53 ^b (0.07)	0.12 ^b (0.03)	0.21
Robot opponents	1.10 ^b (0.09)	-0.52 ^c (0.24)	0.94 ^b (0.08)	0.40 ^b (0.05)	0.42
Robot experienced	1.00 ^b (0.11)	0.01 (0.08)	0.08 (0.12)	0.05 (0.04)	0.18
<i>Panel B: additive specification, weighted least squares estimates (Pooled across buyers and sellers)</i>					
	$[\alpha_{t+1} - \alpha_t] = [\alpha_t^o - \alpha_t]\{\delta_m D_{mt} + \delta_p D_{pt}\} + [\alpha_t^{oo} - \alpha_t]\{\gamma_m D_{mt}^o + \gamma_p D_{pt}^o\}$				
Inexperienced ^a	0.83 ^b (0.21)	-0.18 (0.20)	0.79 ^b (0.06)	0.31 ^b (0.02)	0.40
Experienced	1.07 ^b (0.22)	0.14 (0.11)	0.76 ^b (0.06)	0.13 ^b (0.03)	0.29
Robot opponents	0.94 ^b (0.11)	-0.74 ^b (0.25)	0.81 ^b (0.10)	0.41 ^b (0.04)	0.35
Robot experienced	1.01 ^b (0.11)	0.03 (0.05)	0.05 (0.08)	0.04 (0.03)	0.17

Each observation is weighted by the expected payoff costs of an offer deviation from the optimal offer, as explained in text.

Standard errors in parentheses.

^aTwo outlier observations excluded, as explained in note 6.

^bDenote significantly different from zero at 1% (all 2-tailed tests).

^cDenote significantly different from zero at 5% (all 2-tailed tests).

Figure 5 presents the results based on 400 simulated subjects for both the multiplicative specification of Table 3 and the additive specification of Table 4. For comparison purposes, figure 5 also presents the out-of-sample inexperienced starting α distribution and the out-of-sample experienced ending distribution. The mean value of the out-of-sample experienced distribution is 0.943, with a 95% confidence interval of [0.903, 0.983]. The mean of both simulations fall within this confidence interval, with the additive model simulation mean = 0.945 almost dead center and multiplicative model simulation mean = 0.907 near the lower endpoint. Both simulations, however, exhibit too little dispersion compared to the actual distribution. This is to be expected because we fit a representative agent, and allowed only a very limited sort of heterogeneity in the simulation. We expect that simulations based on heterogeneous agents would improve the fit, but as discussed above individual subject estimates are not feasible in the current data. In particular, neither simulation captures the large number of fully revealing (i.e., $\alpha = 1$) final offers.

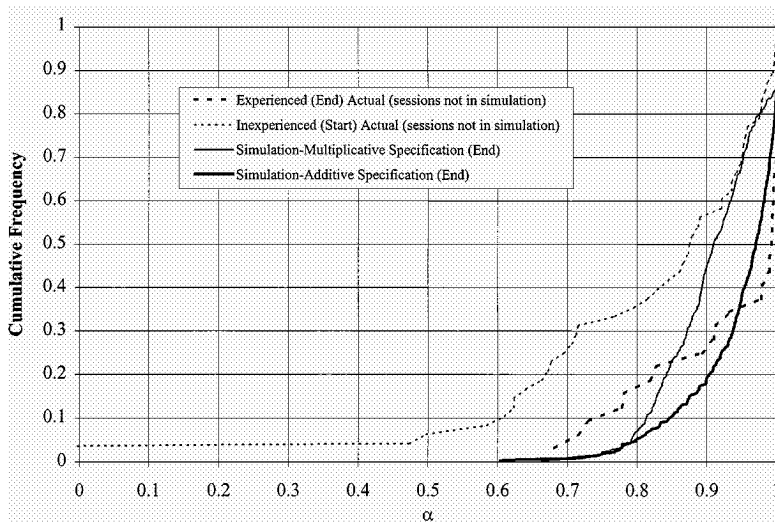


Figure 5. Simulated and actual α distributions for ending periods.

Simulations based on the observational learning model of Section 4.3 lead to similar conclusions, including a lower simulated distribution for the multiplicative version than for the additive version. The observational learning model simulations produce somewhat lower distributions than their direct learning counterparts, mainly because of the positive γ_p coefficients. In the direct model, as α rises there is little countervailing force to reduce it because δ_p is very low (and is indeed often negative, further increasing α). By contrast, in the observational learning model the positive γ_p estimates oppose the rise in α .

5. Discussion

By now there are numerous empirical learning models for studying individual behavior in laboratory games and individual choice tasks, but (with a few partial exceptions noted in the introduction) such models hardly exist for market behavior. In this paper we offer a parsimonious yet fully specified learning model intended to track buyer and seller behavior in a simple laboratory market institution known as the call market. The model is based on the idea that traders respond to ex post errors by adjusting the degree to which their bids or asks reveal true value or cost.

The model enjoys some empirical success. It accounts for a substantial portion of the variation in trader behavior, typically 20–50% but up to 90% in the case of Robot Experienced Sellers. More importantly, it accounts well for traders' response to ex post errors. Parameter estimates indicate that traders systematically react more strongly to ex post "underrevelation" that causes them to miss out on a profitable trade (m -errors) than to ex post "overrevelation" that causes them to have an adverse impact on the market price (p -errors). The parameter estimates indicate a strong recency effect in that accumulated experience

does not significantly dampen traders' reaction to the ex post errors. The estimates also indicate substantial "observational learning" from other traders' ex post errors, again with a greater response to m -errors than to p -errors.

Trader behavior changes over time in the call market institution. Experienced traders in later periods reveal true values and costs to a much greater extent than inexperienced traders do in early periods (or than do ideal risk neutral traders do in Bayesian Nash equilibrium). Two different kinds of evidence show that our learning model accounts well for this behavioral change. The parameter estimates indicate little or no scope for an autonomous time trend in revelation, given coefficients that pick up systematic reaction to ex post errors. Perhaps more importantly, simulation of the estimated model roughly reproduces the mean shift in revelation.

The estimation results include two puzzles. Although previous analysis of the data led us to expect a weaker reaction to the ex post error of adversely affecting price (p -errors) than to missing profitable transactions (m -errors), we did not expect to see negligible direct reactions to p -errors or reactions in the wrong direction. A possible (but ex post) explanation is that some traders react to adversely affecting price by thinking that if they had underrevealed to a much greater extent they would have missed out on the transaction.¹² The other puzzle is that the estimated recency effect is strongest in the treatment with pre-programmed robot opponents where a priori we expected it to be the weakest. Other than to mention that the subjects in this treatment were inexperienced, we have no explanation to offer. We believe that the puzzle should be taken seriously only if it is confirmed with new data and with new model specifications.

Important work lies ahead. Potentially, the most important application of the current line of research is to predict the performance of many different market institutions in a variety of environments. Learning models should be tested in random values environments with the continuous double auction, multiple call market and uniform price double auction institutions. The present study provides design lessons for such experiments. Perhaps the most important negative lesson is that more data are required to estimate heterogeneous models of individual learning, so subjects should remain in the same (buyer or seller) role in longer sessions.¹³ It also may be worthwhile investigating asymmetric distributions for values and costs to produce greater variability in the equilibrium predictions. Although a random values environment is appropriate for studying call market applications such as opening price procedures, other sorts of environments will be appropriate when studying other trading institutions. For example, a stationary values environment or an environment with values that follow a random walk might be appropriate for studying non-auction trading institutions that feature long-term attachments between buyers and sellers. We leave these avenues for future research.

Acknowledgments

Financial support was provided by the National Science Foundation (SBR-9223830 and SBR-9223461). We are grateful to Andrew Davis and Tai Farmer for programming support and to Carl Plat for research assistance. We received helpful comments from audiences at the 1995 Economic Science Conference, the 1996 conference in honor of Robert

Clower at the University of Trento, the 1997 Bonn conference on bounded rationality, and at Middlebury College. Editor Charles Holt and anonymous referees offered several useful suggestions, most notably that we consider observational learning. The usual caveat applies.

Notes

1. We are thinking of Smith-type auctions of complex bundles like airport landing rights. We know of no efficiency proof for these auction mechanisms (indeed, there may be proofs of inefficiency) yet they are quite efficient in practice.
2. McCabe et al. (1992) is a notable exception from the stationary repetitive environment usually employed in market experiments. They add a random constant each period to all values and costs, which causes equilibrium prices to fluctuate randomly but keeps the equilibrium quantity constant.
3. An exception is Camerer and Ho's (1999) analysis of p -beauty contest games, where they discretize the continuous set of choices into intervals to estimate their Experience Weighted Attraction (EWA) model. A similar procedure could be applied to the continuous action space of these call markets, but the EWA model would require additional adjustments to capture the learning asymmetry that we find in the data. Another exception is Capra et al. (1999), who apply a learning model to a continuous-choice duopoly market. Unlike our approach, Capra et al. model learning and beliefs in terms of the probability distribution of others' price choices, and they apply a logit probabilistic choice model to generate choice probabilities.
4. Each inexperienced session employed a common set of random value and cost draws, and each experienced session employed another common set of random value and cost draws.
5. Including the cases with no error does not change the conclusions; cf. note 7.
6. As documented in Cason and Friedman (1997), the offers exhibit heteroscedasticity: Error variance is higher when the expected cost of errors is smaller; i.e., for low and high value draws and low and high cost draws (cf. Smith and Walker, 1993). We assume that the variance is proportional to the reciprocal of the expected cost of an offer error and weight observations by the expected error cost, as calculated in Cason and Friedman (1997).
7. We also present estimates with two outliers excluded in the inexperienced buyers estimate and pooled buyers estimate. In one case a bid of 2.40 was submitted when the value draw was 0.21 ($\alpha = 11.5$), and in the other case a bid of 2.35 was submitted when the value draw was 0.37 ($\alpha = 6.4$). These outliers have a much greater impact on the additive specification of Table 4.
8. We also reestimated the model shown in Table 3 after including the observations with no ex post error. Because $D_{mt} = D_{pt} = 0$ in these cases, for these observations the term in brackets in Eq. (8) is zero. The parameter estimates change very little and our conclusions are robust to including these no error observations, although the η estimates are frequently significantly different from zero due in part to the substantially increased sample size.
9. Standard goodness-of-fit comparisons of the two specifications (such as Adjusted R -squared or likelihood function comparisons) are inappropriate due to the different dependent variables of the multiplicative and additive specifications.
10. Conceivably a buyer might learn to adjust her revelation ratio from observing sellers' ex post errors and a seller might learn from observing buyers' errors. We a priori chose not to consider such cross-observational learning because (except for $k = 0.5$) the optimal revelation ratios of buyers differ from those of sellers. Of course, our learning model's definition of ex post error already incorporates learning from the other side of the market in the narrower sense of observing their actual behavior, but not trying to calculate their optimal behavior.
11. Recall that although each inexperienced (experienced) session had 30 (40) periods, subjects participated in each of the buyer and seller roles half of the time. Also, one half of the simulated subjects followed the Inexperienced-Experienced experience pattern, and one half of the simulated subjects followed the Robot Opponent-Robot Experienced experience pattern.

12. A referee suggests more generally that subjects regard positive trading profits as positive reinforcement that does not demand a change in behavior, while they regard missing a valuable trade and earning zero profit as negative reinforcement that does demand a change. This interpretation is consistent with our own intuition and can be linked to the psychology literature.
13. Indeed, a referee notes that players in the $k = 1$ and $k = 0$ sessions may have to unlearn as buyers the α they learn as sellers, and vice versa.

References

- Bush, R. and Mosteller, F. (1955). *Stochastic Models of Learning*. New York: John Wiley and Sons.
- Camerer, C. (1999). *Behavioral Game Theory*. Book Manuscript, California Institute of Technology.
- Camerer, C. and Ho, T.-H. (1999). "Experience-Weighted Attraction Learning in Normal-Form Games." *Econometrica*. 67, forthcoming.
- Capra, C.M., Goeree, J., Gomez, R., and Holt, C. (1999). "Learning and Noisy Equilibrium Behavior in an Experimental Study of Imperfect Price Competition." Manuscript, University of Virginia.
- Cason, T. (1993). "Call Market Efficiency with Simple Adaptive Learning." *Economics Letters*. 40, 27–32.
- Cason, T. and Friedman, D. (1995). "Learning in Laboratory Markets with Random Supply and Demand." Manuscript, University of Southern California.
- Cason, T. and Friedman, D. (1997). "Price Formation in Single Call Markets." *Econometrica*. 65, 311–345.
- Cheung, Y.-W. and Friedman, D. (1997). "Individual Learning in Games: Some Laboratory Results." *Games and Economic Behavior*. 19, 46–76.
- Daniel, T., Seale, D., and Rapoport, A. (1998). "Strategic Play and Adaptive Learning in the Sealed Bid Bargaining Mechanism." *Journal of Mathematical Psychology*. 42, 133–166.
- Friedman, D. and Ostroy, J. (1995). "Competitiveness in Auction Markets: An Experimental and Theoretical Investigation." *Economic Journal*. 105, 22–53.
- Friedman, D., Massaro, D., Kitsis, S., and Cohen, M. (1995). "A Comparison of Learning Models." *Journal of Mathematical Psychology*. 39, 164–178.
- Friedman, D. and Massaro, D. (1996). "Probability Matching and Underconfidence: An Exploration of the Decision Process." Manuscript, University of California at Santa Cruz.
- Fudenberg, D. and Levine, D. (1998). *Theory of Learning in Games*. Cambridge, MA: MIT Press.
- Garvin, S. and Kagel, J. (1994). "Learning in Common Value Auctions: Some Initial Observations." *Journal of Economic Behavior and Organization*. 25, 351–372.
- Gode, D. and Sunder, S. (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.
- Gluck, M. and Bower, G. (1988). "From Conditioning to Category Learning: An Adaptive Network Model." *Journal of Experimental Psychology: General*. 117, 225–244.
- Kagel, J. (1994). "Double Auction Markets with Stochastic Supply and Demand Schedules: Clearing House and Continuous Double Auction Trading Mechanisms." Manuscript, University of Pittsburgh.
- Kivinen, J. and Warmuth, M. (1995). "Additive versus Exponentiated Gradient Updates for Linear Prediction." Proceedings of the 27th Annual ACM Symposium on Theory of Computing. New York: ACM Press, pp. 209–218.
- McCabe, K., Rassenti, S., and Smith, V. (1992). "Designing Call Auction Institutions: Is Double Dutch the Best?" *Economic Journal*. 102, 9–23.
- McClelland, J. and Rumelhart, D. (1988). *Explorations in Parallel Distributed Processing: A Handbook of Models, Programs and Exercises*. Cambridge, MA: MIT Press.
- Rescorla, R. and Wagner, A. (1972). "A Theory of Pavlovian Conditioning: Variants in the Effectiveness of Reinforcement and Nonreinforcement." In A. Baleck and W. Prokasy (eds.), *Classical Conditioning II: Current Theory and Research*. New York: Appleton-Century Crafts, pp. 64–99.
- Roth, A. and Erev, I. (1995). "Learning in Extensive Form Games: Experimental Data and Simple Dynamic Models in the Intermediate Term." *Games and Economic Behavior*. 8, 164–212.
- Rustichini, A., Satterthwaite, M., and Williams, S. (1994). "Convergence to Efficiency in a Simple Market with Incomplete Information." *Econometrica*. 62, 1041–1063.

- Satterthwaite, M. and Williams, S. (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.
- Satterthwaite, M. and Williams, S. (1993). "The Bayesian Theory of the k-Double Auction." In D. Friedman and J. Rust (eds.), *The Double Auction Market: Institutions, Theories and Evidence*. Redwood City, CA: Addison-Wesley, pp. 99–123.
- Selten, R. and Buchta, J. (1998). "Experimental Sealed Bid First Price Auctions with Directly Observed Bid Functions." In D. Budescu, I. Erev, and R. Zwick (eds.), *Games and Human Behavior: Essays in Honor of Amnon Rapoport*. Mahwah, NJ: Erlbaum Ass.
- Smith, V. and Walker, J. (1993). "Rewards, Experience and Decision Costs in First-Price Auction." *Economic Inquiry*, 31, 237–244.