Available online at www.sciencedirect.com
ScienceDirect
ELSEVIER Journal of Economic Dynamics & Control 31 (2007) 1753–1780
JOURNAL OF Economic Dynamics & Control
www.elsevier.com/locate/jedc

The competitive market paradox

Steven Gjerstad

Economics Department, Krannert School of Management, Purdue University, Lafayette, IN 47907, USA

Received 20 November 2004; accepted 14 June 2006
Available online 18 September 2006

Abstract

The competitive market model is paradoxical. In perfect competition, agents cannot influence price: they only select an output quantity. Such passive behavior does not conform to the intuitive notion of competition. This paper describes an experiment which demonstrates that near or even at a competitive equilibrium price, competition is undiminished. A substantial difference between the performance of sellers and buyers frequently results from this vigorous competition, even with low price variability and approximate efficiency. In double auction experiment sessions conducted with both automated and human agents, exogenous variation of the pace of asks and bids of automated agents demonstrates that the performance difference between sellers and buyers results primarily from a difference between the pace of asks and bids. If the buyers' pace is slower than sellers' pace, buyers make price concessions less frequently than sellers so that prices move below the equilibrium price. Then more buyers become active and fewer sellers remain active. Prices stabilize when changes to the numbers of active buyers and sellers offset the superior bargaining capability of one side or the other. In competitive equilibrium, to a first approximation agents *are* price takers, but that does not preclude vigorous competition: competitive behavior moves to the dimension of bargaining pace.

JEL classification: C78; C92; D41; D44

Keywords: Bargaining; Bounded rationality; Competitive equilibrium; Double auction; Experimental economics; Perfect competition

E-mail address: gjerstad@purdue.edu.

0165-1889/$ - see front matter
doi:10.1016/j.jedc.2006.07.001

1. Introduction

The competitive market model is paradoxical. In perfect competition, sellers and buyers are unable to influence price: their strategic behavior is reduced to selection of an output quantity. Such passive behavior does not conform to the intuitive notion of competition. This paper demonstrates that in the double auction – which is a robustly efficient market institution – the competitive market paradox can be resolved through a close examination of the bargaining process.

In the double auction, sellers and buyers can submit offer prices at any time during a trading period. When prices settle near the competitive equilibrium price, there is not much latitude for strategic behavior along the dimension of price alone. Strategic behavior then moves to the dimension of bargaining pace. When the pace of buyers' bids is slower than the pace of sellers' asks, then price concessions by sellers are more frequent than concessions by buyers. These price concessions by sellers lead to successively lower prices. Ultimately price concessions have a moderate effect on the average trade price, and a small effect on market efficiency, but they have a large effect on the split of surplus between buyers and sellers.

Fig. 1 illustrates the relationship between the efficiency loss and the ratio of buyers' to sellers' surplus when price differs from the equilibrium price. If the price and quantity exchanged are indicated by the dashed lines, then extracted surplus is 99% of the equilibrium surplus, but the ratio of buyers' to sellers' surplus is thirteen ninths, whereas at the competitive equilibrium price $p^* = 60$ that ratio is one.[1] Clearly, if one side of the market produces such a price movement away from equilibrium, agents on that side of the market benefit.

This paper reports a double auction experiment which demonstrates that price variability is low and allocations are almost fully efficient in all experiment sessions. Yet moderate differences between the average price and the equilibrium price and large differences between the performance of sellers and buyers are typical. The pace of sellers' asks and buyers' bids explains a large portion of these differences. These results are demonstrated with a novel experiment technique that combines automated sellers and buyers with human sellers and buyers in the double auction. This technique is a crucial element of the experiment design: the pace of asks by automated sellers and the pace of bids by automated buyers are varied by the experimenter, which produces observed effects on the average price and on the relative performance of sellers and buyers.

Among models of bargaining in the double auction – which include Wilson (1987), Friedman (1991), Gode and Sunder (1993),[2] Easley and Ledyard (1993), and Gjerstad and Dickhaut (1998) – the models by Wilson, by Friedman, and by Gjerstad and Dickhaut all include an explicit formulation of agents' timing decisions. The model by Gjerstad and Dickhaut is the only one of these three models

[1]The efficiency level shown in the figure is slightly below efficiency levels from the experiment sessions reported in this paper. In the 58 sessions reported in Section 4, the average efficiency level is 99.5%.

[2]Cason and Friedman (1996) thoroughly assess the models by Wilson, by Friedman, and by Gode and Sunder.

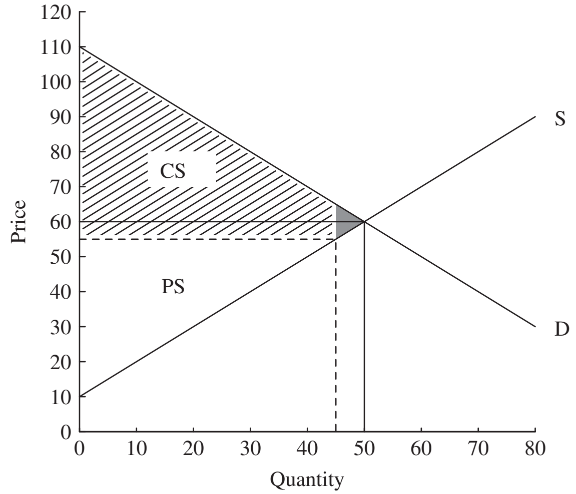

Fig. 1. Consumers' surplus (CS) and producers' surplus (PS) out of equilibrium.

though that has been simulated in continuous time, so their model is best suited for an examination of the impact of pace.

In the model of automated sellers and buyers by Gjerstad and Dickhaut, a seller forms a belief that his ask will be accepted by some buyer. He bases his belief on observed market data, including frequencies of asks, bids, accepted asks, and accepted bids. Then he chooses an ask that maximizes his expected surplus. A seller submits his ask at a random time, with an exponential distribution that depends on both his expected surplus and on the fraction of time remaining in the period. The belief of a buyer, the buyer's bid choice, and the timing of her bid are similar.[3] Simulations of this model demonstrate that sellers' asks and buyers' bids lead to efficient outcomes and stable prices, as in experiment sessions with human subjects. Buyers and sellers in the model base their decisions on the same information as human buyers and sellers, so automated agents can interact with human subjects.

The experiment includes five sessions with only human subjects (*baseline* sessions), 23 sessions with direct interaction between instances of the model and human subjects (*hybrid* sessions), and 30 simulations. Data from these sessions clearly show that prices converge and that allocations are almost fully efficient. Yet buyers frequently outperform sellers even when in equilibrium they earn equal surplus. The cause of this difference in relative performance is evaluated with a statistical model that treats the relative performance of sellers and buyers in a session as a function of the combination of types present in the session. Exogenous variation of the timing

[3]Beliefs and decision rules of sellers and buyers in the model by Gjerstad and Dickhaut are symmetric. Frequently, throughout this paper the description of a seller's or a buyer's behavior is detailed while the analogous statement for the opposite side of the market is less detailed when the symmetric description is obvious.

distribution parameter for the asks and bids by automated agents identifies the impact of bargaining pace on performance. Automated agent types are patient automated sellers, impatient automated sellers, patient automated buyers, and impatient automated buyers.[4] The analysis shows clearly that the relative performance of sellers and buyers is significantly affected by a difference between the pace of asks and the pace of bids, with an advantage to the more patient side of the market.[5]

Persistent vigorous negotiation, even at or near a competitive equilibrium, differs substantially from the conception of price adjustment as a restorative force that only affects the price in a market out of equilibrium. If pace and other aspects of bargaining skill pull prices below equilibrium, the number of active sellers in the market declines, and the number of active buyers increases. As more buyers become active, they outbid one another more frequently, which creates an upward pressure on prices. Similarly, a decrease in the number of active sellers reduces competition among sellers. When the superior bargaining skill of buyers is balanced by the increased number of active buyers and decreased number of active sellers, prices stabilize.[6] Changes to the numbers of active sellers and buyers balance the market price, while the competitive forces from bargaining are always present. This form of strategic behavior and price adjustment substantially resolves the competitive market paradox: once prices have collapsed to two adjacent prices in the bid–ask spread, timing of asks and bids becomes the dominant expression of competitive behavior. Competitive forces balance when the number of active buyers and sellers adjust to offset these differences in pace.

This paper is organized as follows. Section 2 describes the economic environment used in the experiment. Section 2 also describes the double auction mechanism and provides a summary of the heuristic belief learning (HBL) model of seller and buyer behavior used in hybrid sessions and simulations. The experiment design is described in Section 3. Experiment results are reported in Section 4. Conclusions are drawn in the final section.

2. The microeconomic system

Smith (1982) conceptually separates double auction experiments into (1) the induced economic environment, (2) the double auction mechanism, and (3) the

[4]The designations 'patient' and 'impatient' are relative, although the analysis in Section 4.5 benchmarks performance of the four types of automated sellers and buyers relative to human sellers and buyers.

[5]Differences across agents between the frequency and magnitude of price concessions in this model result from differences in agents' strategies, rather than from a difference in preferences. Rubinstein (1982) examines bilateral bargaining between agents with different discount factors, and finds that the agent who discounts the future less obtains a larger fraction of the surplus. In his model, agents employ identical subgame perfect equilibrium strategies, but have heterogeneous preferences. In the model studied in this paper, agents have symmetric preferences and symmetrically defined beliefs, but the strategies of patient and impatient agents differ.

[6]An analogous argument applies if sellers pull the price above the equilibrium price, but the number of active buyers then decreases while the number of active sellers increases.

Table 1
Sellers' unit cost and buyers' unit value schedules for eight sessions

	1	2	3	4	5	6	7	8	9	10	11	12
Sellers 1 & 4	**54**	**66**	**78**	**90**	102	114	126	138	150	162	174	
Sellers 2 & 5	**65**	**70**	**75**	**80**	**85**	**90**	**95**	100	105	110		
Sellers 3 & 6	**61**	**69**	**77**	**85**	**93**	101	109	117	125	133	141	149
Buyers 1 & 4	**138**	**126**	**114**	**102**	90	78	66	54	42	30	18	
Buyers 2 & 5	**127**	**122**	**117**	**112**	**107**	**102**	**97**	92	87	82		
Buyers 3 & 6	**131**	**123**	**115**	**107**	**99**	91	83	75	67	59	51	43

behavior of sellers and buyers. This section describes these three elements, except human behavior, which is evaluated in Section 4 with the analysis of experimental data.

2.1. The economic environment

Sessions include a set of agents $\mathscr{I}$ partitioned into sets $\mathscr{I}_{\mathscr{S}}$ of six sellers and $\mathscr{I}_{\mathscr{B}}$ of six buyers. Sellers have unit costs and buyers have unit values for an abstract 'commodity,' as in Table 1. Units that can trade profitably (for both the seller and buyer) at the midpoint of the equilibrium price set $\{95, 96, 97\}$ are in bold type in the table. All sessions have similar symmetric supply and demand conditions, although six variant conditions differ by constants added to all unit values and costs.[7]

A seller incurs the cost for a unit each time that he completes a trade with a buyer; he receives units of experiment 'currency' from the buyer equal to the negotiated price. His profit is the price received from the buyer minus the unit cost, summed over all units sold. Each buyer has a currency endowment in each period that is no less than the sum of her unit values, in order that she can purchase all units for which she has a positive unit value, if she chooses to do so. This currency endowment is subtracted from the buyer's payoff at the end of the period, so that earnings result from the difference between value and price, summed over each unit purchased.

A seller sees his costs at the beginning of the period, but knows nothing about the costs of other sellers or the buyers' values. The cost schedule of a seller remains fixed from period to period, but he only learns his schedule in a new period when the period begins. A buyer's value schedule is also stationary, and her information is analogous to a seller's.

2.2. The double auction mechanism

The double auction is operationally simple, robustly efficient, and leads to stable transaction prices. In this mechanism, a seller may submit an ask at any time during a trading period in the area labeled 'Enter Ask' (shown in Fig. 2). A seller's ask is his

[7]These shifts reduce the possibility that subjects obtain information about the typical market price from participants in earlier sessions.

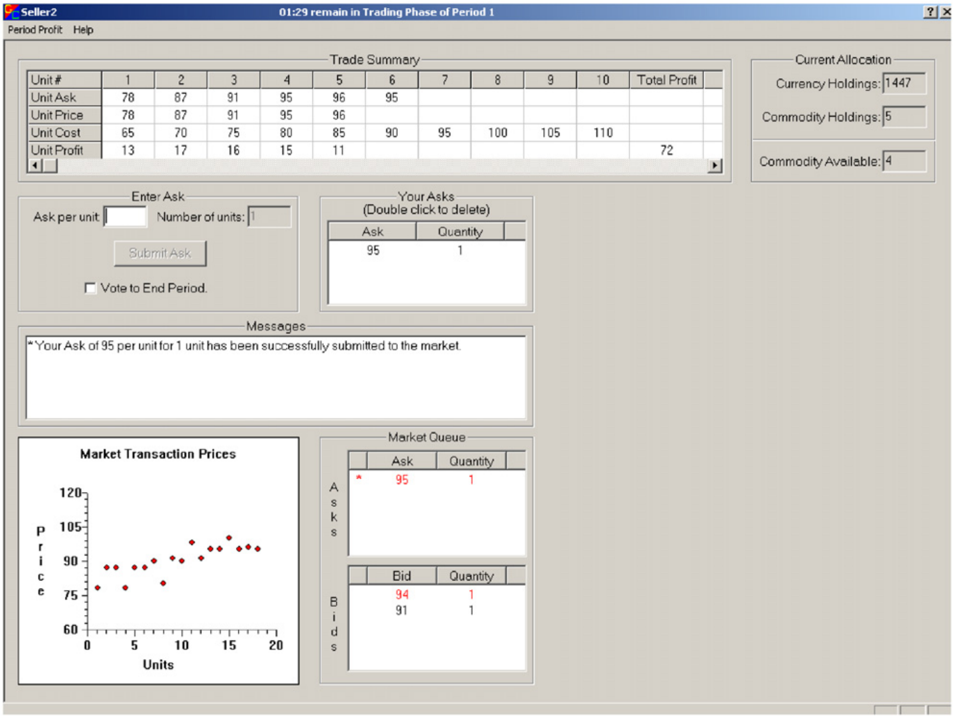

Fig. 2. A seller's screen with elements of market the institution.

current report of the lowest price that he will accept for a unit of the commodity. Similarly, at any time a buyer may submit a bid, which is her current report of the highest price that she will pay for a unit. An ask placed at or below the current high bid results in a trade at the bid price. A bid that meets or exceeds the current low ask results in a trade at the ask price. A seller may make any number of asks, and may trade any number of units that is consistent with his commodity endowment. A buyer may make any number of bids, and may trade any number of units that is consistent with her currency endowment.

Several specific rules are implemented in this version of the double auction. The most important is the 'spread reduction rule,' which requires that a new ask is below the current low ask and a new bid exceeds the current high bid. A seller has the option to remove his current ask. Each seller is permitted a single ask in the market queue at any time and any new ask by a seller replaces his current ask if he has one in the queue. Each ask is the unit price offered by the seller for a single unit: multiple unit trades are not permitted. Similar rules apply to a buyer's bids.

The 'Market Queue' displays current asks and bids. When a seller enters an ask into the queue, he receives a confirmation in the 'Messages' area. This ask also appears in the 'Unit Ask' row of the 'Trade Summary' table, in the column that corresponds to the unit the seller has offered for sale. Similarly, a buyer receives a confirmation message when she enters a bid into the bid queue, and the appropriate cell in her Trade Summary table is updated. When a seller and buyer complete a trade, they both receive a confirmation message, the 'Unit Price' and 'Unit Profit' are recorded in the Trade Summary tables for seller and buyer, and the price is displayed in the 'Market Transaction Prices' graph. The length of each trading period is known to each seller and to each buyer: a clock on the screen of each seller and buyer shows the time remaining in the period. Subjects' instructions explain the double auction mechanism. Appendix C outlines the sequence of steps through a seller's instructions. A buyer's instructions are analogous to a seller's.

2.3. Heuristic belief learning

The HBL model has three elements: a heuristic belief function, maximization of expected surplus relative to this belief, and the timing of an ask or a bid. Empirical acceptance frequencies motivate the definition of heuristic belief functions. Descriptions of expected surplus maximization and the timing of asks and bids follow the discussion of heuristic belief functions.

2.3.1. Empirical acceptance frequencies

For each ask a in the history maintained by seller $i \in \mathscr{I}_{\mathscr{S}}$, the empirical acceptance frequency for a is $\check{p}_i(a) = TA(a)/A(a)$, where $A(a)$ is the number of asks at a and $TA(a)$ is the number of asks at a accepted (or taken) by some buyer. The empirical acceptance frequency for buyer $i \in \mathscr{I}_{\mathscr{B}}$ is $\check{q}_i(b) = TB(b)/B(b)$, where $B(b)$ is the number of bids at b and $TB(b)$ is the number of those accepted (or taken) by a seller. Fig. 3(a) shows the empirical frequencies of asks and accepted asks (as the shaded area) during period 1 of baseline session 1. Fig. 3(b) shows the frequencies of bids and accepted bids from the same period.

The empirical acceptance frequencies $\check{p}_i(a)$ and $\check{q}_i(b)$ are intuitively appealing as beliefs, but they have limitations. Empirical acceptance frequencies are irregular

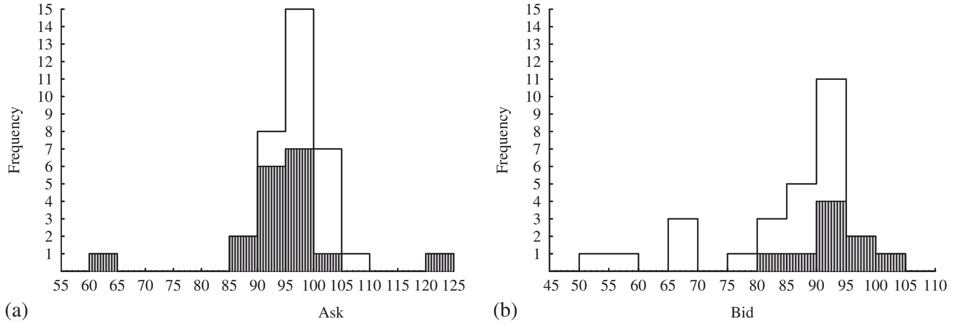

Fig. 3. Panel (a) shows all asks (as the outlined area) and all accepted asks (as the shaded area) from period 1 of baseline session 1. Panel (b) shows bids (outlined) and accepted bids (shaded) for the same period.

after a small number of trades. In the case of Fig. 3(a), after 35 asks there are intervals on which $\check{p}_i(a)$ is undefined. Also, the frequency is not necessarily monotonic, which is undesirable because a buyer is more likely to accept a lower ask. Similar considerations apply to the acceptance frequency of each buyer.

2.3.2. Heuristic belief functions

In order to address these issues, data that generate sellers' and buyers' empirical acceptance frequencies are transformed into heuristic belief functions $p_i(a)$ and $q_i(b)$ which are monotonic, and are defined for any number of asks, bids, and trades. This construction is described in Appendix A.

Gjerstad and Dickhaut show that a seller's belief function $p_i(a)$ is monotonically decreasing, and a buyer's belief function $q_i(b)$ is monotonically increasing. Figs. 4(a) and (c) show examples of a seller's and a buyer's belief. In simulations and hybrid sessions, beliefs are formed from asks, bids, and acceptances that lead up to the last 12 trades ($m = 12$ in the notation of Appendix A).[8] The heuristic belief function $p_i(a)$ in Fig. 4(a) is constructed from data in Figs. 4(b) and (d) using Eq. (A.1). Comparison of Figs. 4(a) and (c) indicates that the transformation of the data generates a monotonic belief function that is similar to the empirical acceptance frequency.

2.3.3. Expected surplus maximization

The expected surplus of seller i for ask a is $\mathrm{E}[S_i(a)] = (a - c_i^k)p_i(a)$, where c_i^k is the cost of his kth unit. The ask $a_i^* = \arg\max \mathrm{E}[S_i(a)]$ maximizes his expected surplus. The maximum expected surplus is $S_i^* \equiv \mathrm{E}[S_i(a_i^*)] = (a_i^* - c_i^k)p_i(a_i^*)$. For each buyer $i \in \mathscr{I}_{\mathscr{B}}$ the expected surplus function, the expected surplus maximizing bid, and the maximum expected surplus are defined similarly. In the original HBL model, the seller is too aggressive, because the ask a_i^* is treated as his last trade opportunity. Appendix B.1 develops a modification to the belief that addresses this issue.

2.3.4. Timing of asks and bids

Agents' asks and bids occur randomly according to exponential distributions. After the κth ask or bid seller $i \in \mathscr{I}_{\mathscr{S}}$ recalculates his expected surplus maximizing ask, as well as the length of time that he will wait before he makes the ask. His wait time is $t_{\kappa_i+1} - t_\kappa$. Similarly, each buyer calculates a new bid and wait time. For each agent $i \in \mathscr{I}$, the cumulative distribution function for the wait time $t_{\kappa_i+1} - t_\kappa$ is

$$\Pr[t_{\kappa_i+1} - t_\kappa < \tau] = 1 - \mathrm{e}^{-\tau/\lambda_i}. \tag{1}$$

The parameter λ_i depends on the seller's (or buyer's) current maximum expected surplus S_i^*, on the time t_κ that has elapsed in the trading period, and on the total time T in the trading period. With this specification, the wait time until the next ask or bid

[8] At the beginning of a new period (other than the first period), the history of asks, bids, and trades used to form beliefs includes some data from the previous period until the 12th trade has occured in the new period.

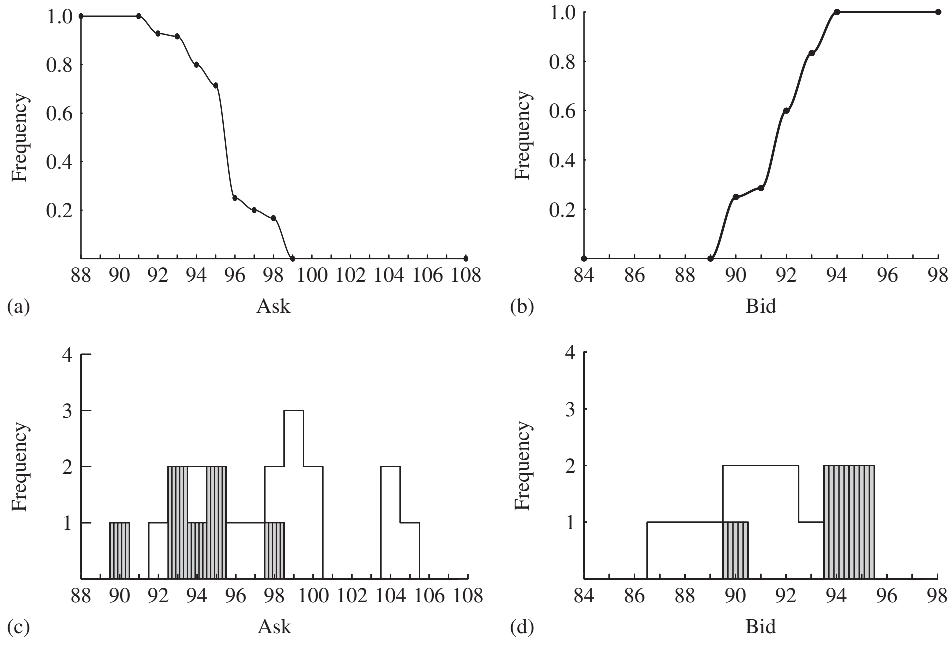

Fig. 4. Beliefs and empirical acceptance frequencies at end of period 1 in baseline session 1.

by some seller or buyer is the minimum of these wait times:

$$t_{\kappa+1} - t_{\kappa} = \min\{t_{\kappa_i+1} - t_{\kappa}\}_{i \in \mathscr{I}}.$$

The minimum of n independent exponential random variables with parameters $\{\lambda_i\}_{i\in\mathscr{I}}$ is an exponential random variable with parameter $\lambda^{-1} = \sum_{i\in\mathscr{I}} \lambda_i^{-1}$.

The specification of the mean wait time λ_i in the original HBL model is $\lambda_i(S_i^*, t_\kappa, T) = \beta_i(1 - \alpha_i t_\kappa/T)/S_i^*$. For this specification of $\lambda_i(\cdot)$, the mean time until agent i with the CDF in Eq. (1) places his ask or bid is λ_i. The mean time decreases in S_i^*: an agent with greater maximum expected surplus is more anxious to trade. The length of each period is $T = 180$ s. The factor α_i, which is $\alpha_i = 0.95$ for all sessions, increases the pace of a seller's ask or a buyer's bid as the period progresses.[9] The scale factor β_i is fixed for each session, but it is varied across sessions in order to determine the effect of timing decisions on performance.

In the original HBL model, the timing specification is too aggressive once the standing bid has reached the level of recent trade prices. Appendix B.2 describes a modification to the timing that reduces the pace of bids once the standing bid reaches the mean of recent trade prices.

[9]When $t = 0$, $1 - \alpha_i t/T = 1$, and as $t \to T$, $1 - \alpha_i t/T \to 1 - \alpha_i$. Surplus on units sold late in the trading period is typically much lower than it is early in the trading period. The factor $1 - \alpha_i t/T$ balances the effect of this difference by decreasing the wait time as the period progresses.

3. Experiment design

The experiment design includes two factors: (1) the pace of asks and bids by automated sellers buyers (described in the previous section), and (2) the fractions of human sellers and of human buyers. Variation of the fractions of automated sellers and buyers produces two conditions. In the balanced condition, half of the sellers and half of the buyers were instances of the modified HBL model, and the other half on each side of the market were human subjects. Performance of automated agents and human agents are directly comparable because there are three schedules of unit costs for sellers, with schedule 1 assigned to sellers 1 and 4, schedule 2 assigned to sellers 2 and 5, and schedule 3 assigned to sellers 3 and 6. In the balanced treatment sellers 1–3 are human agents and sellers 4–6 are automated agents. Buyers also have three schedules, assigned similarly.

In the 'unbalanced' condition, either all of the sellers or all of the buyers were HBL model instances: human subjects filled all roles on the side of the market opposite the HBL model. This condition provides a robust test of the strategic capability of automated agents.

3.1. Experiment sessions

The experiment consists of five baseline sessions, 23 hybrid sessions, and 30 simulations. Table 2 lists the baseline and hybrid sessions chronologically, and also lists the 30 simulations. For each session, the table includes the midpoint of the equilibrium price range, the numbers of human buyer (HB) groups, impatient automated buyer (IAB) groups, patient automated buyer (PAB) groups, human seller (HS) groups, impatient automated seller (IAS) groups, and patient automated seller (PAS) groups. The site is reported for sessions with human subjects.

3.2. Experiment subjects and payments

One baseline session was conducted in June 2001 at the University of York in the UK. Two hybrid sessions were conducted in August 2001 at IBM Research in Yorktown Heights, New York with 12 graduate students in engineering and computer science who were at IBM Research for a summer internship program. Four baseline sessions and 21 hybrid sessions were conducted with undergraduate students at the University of Arizona in November 2001 and March 2002.

Equilibrium earnings for all subjects were between 1425 and 1680 units. Subjects at the University of York received an exchange rate of £0.015 per unit of experiment currency, and a payment of £3 for appearing on time for the session. At the Watson Research Center the exchange rate was \$0.012 per unit of experiment currency and there was no show-up payment. At the University of Arizona the exchange rate was \$0.015 and the show-up payment was \$5. Payments in York ranged from £22.60 to £31.60 and averaged £25.28. Payments at Watson ranged from \$5.00 to \$21.00 with a mean of \$16.08. At the University of Arizona payments ranged from \$11.50 to \$37.75 with a mean of \$26.18. Sessions lasted between 90 and 110 min, with 52.5 min

Table 2
Summary of baseline, hybrid, and simulation sessions

Session	p_e	HB	IAB	PAB	HS	IAS	PAS	Site
1	96	2			2			UK
2	83	1	1		1	1		NY
3	105	2				2		NY
4	96	2				2		AZ
5	96		2		2			AZ
6	108	2			2			AZ
7	108	1	1		1	1		AZ
8	108	1	1		1	1		AZ
9	120	1	1		1	1		AZ
10	108		2		2			AZ
11	108	2				2		AZ
12	120	2			2			AZ
13	96	1	1		1	1		AZ
14	96	1	1		1	1		AZ
15	120	1		1	1		1	AZ
16	120	1		1	1		1	AZ
17	136	2			2			AZ
18	108	1		1	1		1	AZ
19	96	2			2			AZ
20	120	2					2	AZ
21	120			2	2			AZ
22	136			2	2			AZ
23	136	2					2	AZ
24	108	1		1	1		1	AZ
25	108	1		1	1		1	AZ
26	108	1	1		1	1		AZ
27	96	1	1		1	1		AZ
28	96	1	1		1	1		AZ
29–38	96			2			2	
39–48	96		2				2	
49–58	96			2		2		

for the 15 periods, each with a 180 s trading phase and a 30 s review phase. Seating and instructions lasted between 30 and 50 min and payment required about 10 min per session.

4. Data analysis

As in previous double auction experiments and HBL model simulations, approximate price convergence and convergence of allocations to efficient outcomes typically occur in all sessions. In view of this convergence, it is surprising though that the ratio of sellers' realized surplus earnings to the buyers' realized surplus earnings (as fractions of their equilibrium earnings) differs significantly from one. Moreover, the ratio of an individual seller's earnings to his equilibrium earnings is highly correlated with the average ratio of the earnings of other sellers to their equilibrium earnings. This correlation justifies a statistical model that treats the relative performance of sellers and buyers as a function of the seller and buyer types present in each session. This performance analysis establishes the main result of the paper by regressing the ratio of sellers' to buyers' performance in each session on the composition of types present in the session. Types are human sellers, patient automated sellers ($\beta_i = 400$), impatient automated sellers ($\beta_i = 250$), human buyers, patient automated buyers ($\beta_i = 400$), and impatient automated buyers ($\beta_i = 250$). This analysis establishes that (1) patient automated sellers and buyers outperform impatient automated sellers and buyers, (2) impatient automated buyers and human buyers have similar performance, and (3) human buyers outperform human sellers.[10] Since impatient automated buyers and human buyers have similar performance, and since patient automated buyers outperform impatient automated buyers, it follows that buyer performance is enhanced by a reduction to the pace of bids, at least locally in the vicinity of the pace of human buyers.

4.1. Price convergence

Price convergence is estimated with the model $s(n) = an^{-b}\eta(n)$, where $s(n)$ is the period n standard deviation of prices and the random variables $\eta(n)$ are independent and lognormally distributed. This can be expressed as the linear model $\ln s(n) = \ln a - b \ln n + \varepsilon(n)$. Since $\mathrm{E}[s(1)] = a$, $\hat{a}$ estimates the period 1 price standard deviation; $\hat{b}$ estimates the rate at which prices converge.

Table 3 reports parameter estimates and tests of the hypotheses that $\hat{\alpha} > 0$ and $\hat{\beta} > 0$ for each session type. Bold type in Table 3 (and in all subsequent tables in Section 4) indicates coefficients that differ from zero at levels $p = 0.05$ or below. Reported p-values are for the estimates $\hat{b}$; all p-values for estimates $\hat{a}$ are less than 10^{-10}.

Prices initially vary less in the HBL model than in baseline sessions: much of the convergence in simulations occurs in the first period. Although estimates $\hat{b}$ show a convergence trend for all regressions at the level $p = 0.05$ or less, price variability declines less across periods in simulations than in baseline sessions. Each hybrid treatment produces estimates of $\hat{a}$ and $\hat{b}$ between those from simulations and those from baseline sessions. This is natural, since hybrid sessions include six human subjects and six instances of the HBL model. Most importantly though, prices stabilize by period 15 for all session types. Estimates of final period

[10] Fifteen price errors have been corrected in the 23 hybrid sessions, and three price errors have been corrected in the five baseline sessions. Most errors are due to an extra digit in the bid (e.g., a bid of 1,115 rather than 115), or a missing digit in an ask (e.g., an ask of 9 rather than an ask of 94). In 28 sessions, there were 12,962 trades, so these 18 price errors represent less that one error per 720 trades. Four errors were between a human seller and a human buyer, and 14 errors were made by a human seller or buyer with an automated buyer or seller, so the argument that automated sellers and buyers outperform human sellers and buyers would also hold, with greater estimates of the advantage to the automated sellers and buyers, if the price errors were not changed.

Table 3
Price convergence in simulations, hybrid sessions, and baseline sessions

Treatment	$\hat{a}$	$\hat{b}$	p-Value ($\hat{b}$)	R^2	F	$\hat{s}(15)$
Simulation	**2.195**	**0.181**	0.0002	0.092	14.9	1.34
Balanced (patient)	**4.215**	**0.508**	0.0000	0.420	107.3	1.06
Balanced (impatient)	**4.995**	**0.525**	0.0000	0.523	63.7	1.21
Unbalanced (automated buyers)	**5.357**	**0.489**	0.0000	0.250	19.3	1.42
Unbalanced (automated sellers)	**4.670**	**0.581**	0.0000	0.420	52.8	0.97
Baseline (all human)	**13.019**	**0.928**	0.0000	0.638	128.5	1.05

Table 4
Efficiency in simulations, hybrid sessions, and baseline sessions

Treatment	$\bar{a}$	$\hat{b}$	p-Value ($\hat{b}$)	R^2	F	$\hat{e}(15)$
Simulation	**0.0057**	0.0201	0.7855	0.0005	0.074	0.995
Balanced (patient)	**0.0377**	**0.7575**	0.0000	0.2594	41.3	0.995
Balanced (impatient)	**0.0847**	**1.2624**	0.0000	0.5499	70.8	0.997
Unbalanced (automated buyers)	**0.0165**	**0.6325**	0.0005	0.1921	13.7	0.997
Unbalanced (automated sellers)	**0.0877**	**1.0552**	0.0000	0.3569	40.5	0.995
Baseline (all human)	**0.0686**	**0.7806**	0.0000	0.3251	35.2	0.992

price variability, $\hat{s}(15)$, are all near one, whereas the range from the highest value to the lowest cost is 84.

4.2. Efficiency

Efficiency is also modeled with exponential convergence. If efficiency in period n is $e(n)$, and if foregone surplus declines exponentially across periods, then $1 - e(n) = an^{-b}\eta(n)$. This is equivalent to the linear model $\ln(1 - e(n)) = \ln a - bn + \varepsilon(n)$. Table 4 summarizes regression results for all experiment session types.

All estimates $\hat{a}$ of the initial level of inefficiency $1 - \hat{e}(n)$ differ from zero with a p-value less than 10^{-10}. Significant convergence trends, measured by estimates $\hat{b}$, occur in baseline sessions and in all hybrid sessions. Simulations do not demonstrate a trend of convergence to efficiency across periods, yet the estimate of the efficiency level in the final period of the simulations, 99.5%, is similar to the estimated efficiency attained in hybrid sessions and baseline sessions by period 15.

4.3. Surplus split

In the benchmark case in Fig. 1, all trades occur at one price and efficiency is high, yet the surplus split is highly uneven. When all trades occur at a single price $p_\alpha = \alpha p_1 + (1 - \alpha)p_e$, where p_1 is the highest price at which supply is zero and p_e is the equilibrium price, the fraction of surplus lost is only α^2, but the ratio of surplus earned by buyers to surplus earned by sellers is $R(\alpha) = 1 + 4\alpha/(1 - \alpha)$. For example, if the fractional difference of the price from equilibrium is $\alpha = 0.07$, then the surplus loss is only 0.49%, and price variability is zero, but the ratio of buyers' surplus to sellers' surplus is 1.301. The value and cost schedules from Table 1 – which were the basis for all value and cost schedules used in the experiment – are approximately linear, so the relationship between price deviations and surplus split is similar to the benchmark case, as Table 5 demonstrates.

Fig. 5 shows sellers' surplus relative to buyers' surplus as a function of the difference between the market price (p_α) and the equilibrium price (p_e). Relative performance r in this context is defined as $r \equiv e_{\mathscr{S}}/e_{\mathscr{B}}$, where $e_{\mathscr{S}}$ is the surplus attained by sellers divided by the equilibrium surplus of sellers and $e_{\mathscr{B}}$ is defined similarly. The figure shows this relationship for three types of simulations.

The triangles on the upper right show observations from 10 simulations with patient sellers and impatient buyers. In each of these simulations, $p_\alpha - p_e > 0$ and $\ln(e_{\mathscr{S}}/e_{\mathscr{B}}) > 0$: the mean price is above the equilibrium price and patient automated sellers outperform impatient automated buyers. The 10 squares in the center of the

Table 5
Surplus and efficiency as a function of price deviation from equilibrium

$p_\alpha - p_e$	Buyer surplus	Seller surplus	Total surplus	Surplus ratio ($e_{\mathscr{S}}/e_{\mathscr{B}}$)	$\ln(e_{\mathscr{S}}/e_{\mathscr{B}})$	Efficiency (%)
4	476	720	1196	1.513	0.414	98.7
3	504	704	1208	1.397	0.334	99.7
2	534	674	1208	1.262	0.233	99.7
1	574	638	1212	1.111	0.106	100.0
0	606	606	1212	1.000	0.000	100.0
−1	638	574	1212	0.900	−0.106	100.0
−2	674	534	1208	0.792	−0.232	99.7
−3	704	504	1208	0.716	−0.334	99.7
−4	720	476	1196	0.661	−0.414	98.7

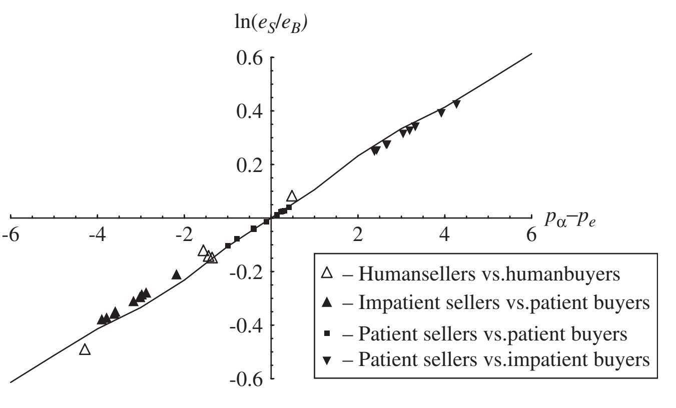

Fig. 5. Surplus ratio and price difference from equilibrium in simulations.

figure are from simulations with patient sellers and buyers. These are distributed fairly evenly around the equilibrium price and the equilibrium surplus split. The 10 observations on the lower left in the figure are from simulations with impatient sellers and patient buyers. The curve in Fig. 5 traces out the logarithm of the theoretical surplus split (from column 6 of Table 5) as a function of $p_\alpha - p_e$. A difference between the pace of asks and bids has a pronounced effect on the relative performance of sellers and buyers in these simulations. Fig. 5 also shows the ratio of sellers' to buyers' surplus as a function of price differences from equilibrium in the five baseline sessions. In four of five baseline sessions, human buyers outperform human sellers.

Performance differences between sellers and buyers are also typical in hybrid sessions, as shown in Fig. 6. In all five hybrid sessions with automated sellers and human buyers the average price is above the equilibrium price and automated sellers outperform human buyers. In four hybrid markets with human sellers and automated buyers the average price is at or below the equilibrium price and automated buyers outperform human sellers. These two observations suggest that automated sellers outperform human buyers and automated buyers outperform human sellers.

Balanced hybrid sessions are shown in Fig. 6 as squares. Thirteen of 14 balanced hybrid sessions resulted in an average price at or below the equilibrium price. Sellers outperformed buyers in only one of 14 balanced sessions. Since automated sellers and automated buyers are specified symmetrically, the performance advantage of buyers in balanced hybrid sessions most likely results from the superior performance of human buyers to human sellers in these sessions.[11]

Performance ratio statistics from baseline sessions and hybrid sessions do not have the degree of regularity observed in the simulations, yet a consistent pattern emerges between the composition of types in a session and the division of surplus between sellers and buyers. A statistical model, developed in Section 4.5, explains these patterns of relative performance in terms of the types present in each session. Coefficients for each type are interpreted as contributions to the performance of sellers relative to buyers. In order to carry out this analysis, it is important to demonstrate that the performance of sellers is positively correlated with one another, as is the performance of buyers. High correlation of performance within each side of the market justifies statistical analysis in which the performance of each side of the market is modeled as dependent on the profile of types included in a session. The next subsection demonstrates this performance correlation.

4.4. Performance correlation

Let e_i be the ratio of the surplus earned by seller i to his equilibrium surplus. The average of e_j for sellers other than i is $e_{-i} = \frac{1}{5}\sum_{j \neq i} e_j$. If the coefficient α_1 is close to

[11]Smith and Williams (1982, p. 20) and Williams and Smith (1984, p. 111) also observed a tendency for buyers to outperform sellers. They describe this tendency as the 'weak seller hypothesis.' There is no clear reason for this tendency, though a referee suggests that it may be due to the frequency with which most people engage in transactions as buyers and the relative infrequency with which most people act as sellers.

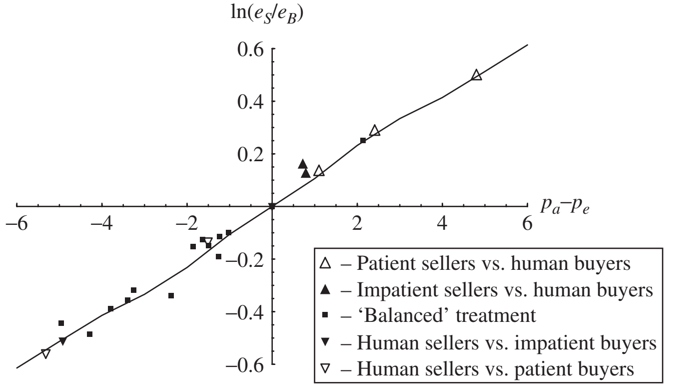

Fig. 6. Surplus ratio and price difference for each hybrid session.

one in the regression $e_i - 1 = \alpha_0 + \alpha_1(e_{-i} - 1) + \varepsilon_i$ then the average performance of all other sellers is a good predictor of the performance of an individual seller.

The null hypothesis is that the performance of each agent is independent. This implies that if the normalized performance of other sellers is $e_{-i} - 1$, the expected performance for seller i is $\mathrm{E}[e_i - 1] = -\frac{5}{7}(e_{-i} - 1)$. An intuitive argument for this is that, if the market is efficient, then an increase to the performance of the other five sellers must be compensated by a decrease in the performance of seller i and the six buyers.

The argument can be made more formally as well. Assume that surplus for each agent in equilibrium is s^*. Assume also that the market is efficient. (The first of these assumptions holds approximately for the economic environment of the experiment; the second holds approximately for the experiment data.) Then average efficiency is $\frac{1}{12}\sum_{i=1}^{12} e_i = \frac{1}{12}\sum_{i=1}^{12} s_i/s^*$. Since equilibrium surplus for each agent is $s_i = s^*$, $\frac{1}{12}\sum_{i=1}^{12} e_i = 1$. So independence of agents' performance measures implies that $\sum_{i=1}^{12}(e_i - 1) = 0$. This in turn implies that they are independent and identically distributed on the set $\{e : \sum_{i=1}^{12}(e_i - 1) = 0\}$, so $\mathrm{E}[e_i - 1] = -\sum_{i' \in (\mathcal{I}_S \setminus \{i\}) \cup \mathcal{I}_B} \mathrm{E}[e_{i'} - 1]$. If the average of the normalized performance measures for all sellers other than seller i is $e_{-i} - 1$ then $\mathrm{E}[e_i - 1] = -5(e_{-i} - 1) - \sum_{i' \in \mathcal{I}_B} \mathrm{E}[e_{i'} - 1]$. Under the null hypothesis, $\mathrm{E}[e_i - 1] = \mathrm{E}[e_{i'} - 1]$ for all $i' \in \mathcal{I}_B$ so $\mathrm{E}[e_i - 1] = -\frac{5}{7}(e_{-i} - 1)$ is the expectation of the normalized performance measure of seller i.

Table 6 shows regression results for each design treatment. Since all estimated slope coefficients α_1 are closer to one than to $-\frac{5}{7}$, it is appropriate to model relative performance of sellers to that of buyers as a function of the distribution of types in a session.

To a large extent, performance correlation within each side of the market results from the bid improvement rule. In a market with several buyers and sellers, each of whom has multiple units to transact, a narrow bid–ask spread is typical. Prices are therefore bounded by the bid–ask spread. As a consequence, traders who are

Table 6
Performance correlation among sellers and among buyers

Treatment	Type	α_0	α_1	R^2	F
Baseline	Sellers	−0.0259	**1.1479**	0.3506	238.6
	Buyers	0.0005	**0.7812**	0.3066	195.4
Balanced	Sellers	**−0.0072**	**0.9230**	0.6578	2083
	Buyers	**0.0113**	**0.8267**	0.4072	744.6
Automated sellers	Sellers	0.0016	**0.9763**	0.8725	3066
	Buyers	**−0.0470**	**0.6186**	0.1547	81.9
Automated buyers	Sellers	**−0.0188**	**0.8685**	0.5194	386.9
	Buyers	0.0090	**0.9345**	0.6913	801.6

Table 7
A simplified hypothetical experiment design and the price for each session

HS	IAS	PAS	HB	IAB	PAB	$\bar{p}$
2			2			95
		2	2			97

strategically weak can make price concessions, but a weak buyer can pay no more than the standing ask and a weak seller can accept no less than the standing bid.

Although there is some variation in performance within each side of the market, the effect of the pace variables could be identified even if there were no variation within each side of the market. Suppose that every seller has the same performance, but the performance of the sellers is affected by the composition of types in the market, such as human sellers or impatient automated sellers. The simple hypothetical experiment in Table 7 would identify the pace of human buyers (HB) relative to human sellers (HS) and to patient automated sellers (PAS), since the average price varies with the composition of the agent types in the session, and the midpoint of the equilibrium price range is 96. The next section develops this idea to evaluate performance.

4.5. Relative performance model

This section develops and estimates a model of the relative performance of sellers and buyers. The model treats relative performance as a function of the profile of types present in each session, where relative performance r is $r \equiv e_{\mathscr{S}}/e_{\mathscr{B}}$ as in Section 4.3.

Assume that relative performance is a random function $r(X)$ from $\mathbb{R}_+$ to $\mathbb{R}_+$ where X is a random variable that measures the pace of asks relative to bids. Relative performance $r(X)$ should treat sellers and buyers symmetrically. If relative pace is inverted, relative performance should be inverted: $r(X^{-1}) = r(X)^{-1}$. Functions of the form $r(X) = X^{\gamma}$ have this property for any $\gamma \in \mathbb{R}$.

Assume that the relative pace variable X has a lognormal distribution where $X = \mathrm{e}^{Y}$ and $Y \sim \mathrm{N}(\mu, \sigma)$. Let $S = \{s_1, s_2\}$ and $B = \{b_1, b_2\}$ denote the two seller and two buyer groups. Assume that for each seller group there is a random variable Y_{s_j} that measures the pace of seller group j for $j = 1, 2$ and for each buyer group there is a random variable Y_{b_j}. Let $Y = \sum_{j=1}^{2} Y_{b_j} - \sum_{j=1}^{2} Y_{s_j}$. The underlying assumption in this model is that the random variable $X = \mathrm{e}^{Y}$ is a function of contributions from each of the two seller groups and each of the two buyer groups. Note that if Y_{s_j} and Y_{b_j} have a common expectation, then the expected value of Y is zero, so that $\mathrm{E}[X] = 1$, in which case the relative performance of sellers and buyers is equal in expectation.

The numbers in columns three through eight in Table 2, with the last three rows (for the three types of simulations) repeated ten times, represent the design matrix for the relative performance model. Since there are two groups of sellers with identical costs and two groups of buyers with identical values in each session, the numbers in columns three through five sum to two in each row, as do the entries in columns six through eight. The column rank of the design matrix is four, so the regressions estimate type contributions to relative performance with pairwise comparisons, first between human subjects and patient automated sellers and buyers, and then between human subjects and impatient automated sellers and buyers.

The regression model is

$$\ln r(X_k) = \alpha_0 + \alpha_{\mathrm{h}} s_{\mathrm{h},k} + \alpha_{\mathrm{p}} s_{\mathrm{p},k} - \beta_{\mathrm{h}} b_{\mathrm{h},k} - \beta_{\mathrm{p}} b_{\mathrm{p},k} + \varepsilon_k, \tag{2}$$

where X_k is the relative pace of sellers to buyers in session k. The variables $s_{\mathrm{h},k}$ and $s_{\mathrm{p},k}$ indicate the number of human seller groups and the number of patient automated seller groups included in session k, where each of these variables takes the value 0, 1, or 2 and $s_{\mathrm{h},k} + s_{\mathrm{p},k} \leqslant 2$. (The number of impatient automated seller groups in a session is $s_{\mathrm{i},k} = 2 - s_{\mathrm{h},k} - s_{\mathrm{p},k}$.) Similarly, $b_{\mathrm{h},k}$ and $b_{\mathrm{p},k}$ are the number of human buyer and patient automated buyer groups in session k. The negative signs on the buyer terms in regression equation (2) have been chosen so that the strength of seller and buyer coefficients can be compared directly, since the negative signs adjust for the fact that strong buyers have the opposite impact from strong sellers on the ratio of seller performance to buyer performance. With this specification of the model, a higher value for any of the estimates $(\hat{\alpha}_{\mathrm{h}}, \hat{\alpha}_{\mathrm{p}} \hat{\beta}_{\mathrm{h}}, \hat{\beta}_{\mathrm{p}})$ indicates better performance for the group associated with that parameter estimate.

Table 8 reports coefficient estimates, standard errors, and p-values for tests of the hypotheses that all coefficients differ from zero. The R^2 statistic for this model is 0.6979: the model explains a large fraction of the variability in relative performance. The F statistic is 30.6, which has significance at a level less than 10^{-12}.

A constant term that differs significantly from zero indicates an experiment design that favors either the sellers or the buyers, independently of the composition of seller and buyer types. This hypothesis is decisively rejected. Pace coefficients can be interpreted as contributions to the relative performance of sellers and buyers. The presence of human sellers in a session k ($s_{\mathrm{h},k} > 0$) has a detrimental effect on the ratio

Table 8
Estimates for humans vs. patient automated sellers and buyers

	$\hat{\alpha}_0$	$\hat{\alpha}_h$	$\hat{\alpha}_p$	$\hat{\beta}_h$	$\hat{\beta}_p$
Estimate	−0.003	**−0.116**	**0.142**	−0.001	**0.157**
SE	0.066	0.036	0.028	0.034	0.028
p-Value	0.962	0.002	0.000	0.986	0.000

of seller surplus to buyer surplus in that session, since the estimate $\hat{\alpha}_h$ is significantly negative. Human buyers have a neutral effect on the relative performance statistic (since $\hat{\beta}_h$ is not significantly different from zero); relatively then, human buyers are stronger than human sellers. Patient automated buyers and patient automated sellers have comparable effects on performance. Finally, both human sellers and human buyers are weaker than patient automated sellers or buyers, since both $\hat{\alpha}_p$ and $\hat{\beta}_p$ differ significantly from zero and are similar in magnitude. Patient automated sellers have a significant positive impact on the relative performance of sellers. Patient automated buyers contribute negatively to the relative performance of sellers. Consequently, patient automated buyers are stronger than human buyers.

These relations can be summarized with an ordering $\prec$ over the six types in the experiment. These types are human sellers and buyers (HS and HB), impatient automated sellers and buyers (IAS and IAB), and patient automated sellers and buyers (PAS and PAB). Results of the regression in Table 8 can be summarized as

$$HS \prec HB \prec PAB \sim PAS. \tag{3}$$

In order to assess the impact of impatient automated sellers and buyers, the regression equation is presented in the alternative form

$$\ln r(X_k) = \alpha_0 + \alpha_h s_{h,k} + \alpha_i s_{i,k} - \beta_h b_{h,k} - \beta_i b_{i,k} + \varepsilon_k, \tag{4}$$

where $s_{i,k}$ and $b_{i,k}$ are the numbers of impatient automated seller and impatient automated buyer groups in session k. As a result of the collinearity in the design matrix, regression equation (4) has the same explanatory power as regression equation (2), so that the R^2 and F statistics are identical, but it produces different coefficient estimates. Coefficients for the alternative formulation in Eq. (4) are shown in Table 9.

The most important additional comparison in Table 9 is between human buyers and impatient automated buyers. These coefficients are similar, and both have small standard errors. Consequently, the timing specification for impatient automated buyers yields performance that is similar to human performance. These results can also be summarized with the order $\prec$ used in (3). The regression results in Table 9 can be summarized as

$$HS \prec HB \sim IAB \sim IAS. \tag{5}$$

Combining the results in (3) and (5) results in

$$HS \prec HB \sim IAB \sim IAS \prec PAB \sim PAS. \tag{6}$$

Table 9
Estimates for humans vs. impatient automated sellers and buyers

	$\hat{\alpha}_0$	$\hat{\alpha}_\mathrm{h}$	$\hat{\alpha}_i$	$\hat{\beta}_\mathrm{h}$	$\hat{\beta}_i$
Estimate	−0.028	**−0.258**	**−0.142**	**−0.158**	**−0.157**
SE	0.041	0.032	0.028	0.032	0.028
p-Value	0.497	0.000	0.000	0.000	0.000

These regressions demonstrate that, aggregated across sessions, a substantial proportion of the difference between the performance of sellers and buyers is explained by the types present in the market. An alternative regression, which utilizes only the relative performance from the final trading period, produces similar results. This is important, since one possible scenario is that strategic behavior is important along the path to a competitive equilibrium, but its importance diminishes once prices stabilize and the opportunity to influence price reduces. All of the qualitative results noted above for the full data set are obtained in the final period as well, and the coefficients are similar: pace remains an important factor even after prices and allocations converge to an approximate equilibrium.

5. Conclusions

This paper demonstrates that a substantial deviation from equilibrium earnings is consistent with the price stability and approximate efficiency that are typical in double auction experiments. Deviation from equilibrium earnings is driven to a large extent by a difference between the pace of sellers' asks and buyers' bids. Strategic interaction among sellers and buyers is a crucial element of the price discovery process, but it is surprising that strategic considerations figure prominently in the interactions of market participants after low levels of price variability are attained. Once prices stabilize, there is frequently a clearly established bid–ask spread. Sellers and buyers vie with each other to influence the market price, and consequently to improve relative performance. If, for example, sellers are more likely to yield and accept the standing bid rather than wait until a buyer accepts the standing ask, then the average price will decline. The cumulative effect of these concessions has a considerable effect on relative performance, as demonstrated by the statistical model in Section 4.5. Consequently, price taking behavior – which is virtually synonymous with competitive equilibrium – is not implied by price stability and approximate efficiency, that is, by approximate competitive equilibrium. There is a substantial role for strategic behavior even under competitive outcomes: the pace of sellers' asks and buyers' bids is a key element of such strategic behavior in the temporally unstructured bargaining of the double auction.

Acknowledgments

This paper has benefitted from suggestions by Jon Leland, Jason Shachat, two referees, the coeditor Cars Hommes, and from seminar participants at the University of Arizona in November 2002, the ESA conference in Pittsburgh in June 2003, the Stanford Institute for Theoretical Economics workshop on Price Formation in Financial Markets in July 2003, the Econometric Society meetings in January 2004, and the SAET conference in Vigo, Spain in July 2005. Experiment software was developed at the T.J. Watson Research Center at IBM by Amit J. Shah. Support from NSF Grant DUE-022634 is gratefully acknowledged. Jim Cox in his former capacity as director of the Economic Science Laboratory at the University of Arizona provided valuable support for this work.

Appendix A

A.1. Construction of belief functions

Step 1: The initial belief of agent $i \in \mathscr{I}$ is defined on the set $[0, M_i]$. Buyer $i \in \mathscr{I}_{\mathscr{B}}$ believes that a bid at zero will be accepted with probability zero and a bid at or above M_i will be accepted with probability one; seller $i \in \mathscr{I}_{\mathscr{S}}$ believes that an ask at zero will be accepted with probability one and an ask at M_i will be accepted with probability zero.

In simulations and hybrid sessions, for seller $i \in \mathscr{I}_{\mathscr{S}}$ with the vector of unit costs $\{c_i^k\}_{k=1}^{K_i}$, the bound M_i is $M_i \equiv 2^{2+[\log_2(\max_k\{c_i^k\})]}$. For buyer $i \in \mathscr{I}_{\mathscr{B}}$ with unit values $\{v_i^k\}_{k=1}^{K_i}$, the bound M_i is $M_i \equiv 2^{1+[\log_2(\max_k\{v_i^k\})]}$ where $[\log_2(\max_k\{v_i^k\})]$ is the greatest integer less than or equal to $\log_2(\max_k\{v_i^k\})$.

Step 2: Each agent $i \in \mathscr{I}$ has finite memory length m_i. Each new ask is adjoined to a set $\mathscr{A}$ of asks that have previously occurred; each new bid is adjoined to a set of bids $\mathscr{B}$. Define $\mathscr{A}_{m_i} \subseteq \mathscr{A}$ as the set of asks that have been made during the negotiations leading to the last m_i trades. Define $\mathscr{B}_{m_i} \subseteq \mathscr{B}$ similarly and let $\mathscr{D}_{m,i} \equiv \mathscr{A}_{m_i} \cup \mathscr{B}_{m_i} \cup \{0, M_i\}$. The belief of seller $i \in \mathscr{I}_{\mathscr{S}}$ is defined at each possible ask $a \in \mathscr{D}_{m,i}$ as

$$\widehat{p}_i(a) = \frac{\sum_{d \geqslant a} TA(d) + \sum_{d \geqslant a} B(d)}{\sum_{d \geqslant a} TA(d) + \sum_{d \geqslant a} B(d) + \sum_{d \leqslant a} RA(d)},$$

where $TA(\cdot)$ and $B(\cdot)$ are as defined in the description of the empirical acceptance frequencies and $RA(d) \equiv A(d) - TA(d)$ counts rejected asks at d. Similarly, the belief of buyer $i \in \mathscr{I}_{\mathscr{B}}$ is defined at each possible bid $b \in \mathscr{D}_{m,i}$ as

$$\widehat{q}_i(b) = \frac{\sum_{d \leqslant b} TB(d) + \sum_{d \leqslant b} A(d)}{\sum_{d \leqslant b} TB(d) + \sum_{d \leqslant b} A(d) + \sum_{d \geqslant b} RB(d)}.$$

The rationale for the definition of a seller's belief function $\widehat{p}_i(a)$ is that any ask at or above a that has been accepted reinforces the seller's belief that an ask at a will be accepted, as does any bid at or above a. Any ask at a value at or below a that has not been accepted decreases the seller's belief that a will be accepted. The rationale for the definition of a buyer's belief function $\widehat{q}_i(b)$ is analogous.

Step 3: In markets that include the spread reduction rule, which requires that any new ask is placed at a value below the current low ask (or 'standing ask') *sa* and any new bid is above the current high bid *sb*, beliefs are modified to account for this restriction. This is accomplished by setting $\widehat{p}_i(d) = 0$ for all $d \geqslant sa$ and setting $\widehat{q}_i(d) = 0$ for all $d \leqslant sb$.

Step 4: The beliefs $\widehat{p}_i(a)$ and $\widehat{q}_i(b)$ are extended from $\mathscr{D}_{m,i}$ to $[0, M_i]$ by cubic spline interpolation. For each $d_k \in \mathscr{D}_{m,i}$, define $p_i(a)$ on $[d_k, d_{k+1})$ by $p_i(a) = \beta_0 + \beta_1 a + \beta_2 a^2 + \beta_3 a^3$, where four conditions uniquely determine the coefficients $(\beta_0, \beta_1, \beta_2, \beta_3)$ on the interval $[d_k, d_{k+1})$. These conditions are $p_i(d_k) = \widehat{p}_i(d_k)$, $p_i(d_{k+1}) = \widehat{p}_i(d_{k+1})$, $p_i'(d_k) = 0$, and $p_i'(d_{k+1}) = 0$. The extension of beliefs $\widehat{q}_i(b)$ for buyer i to $[0, M_i]$ is identical.

Appendix B

B.1. Modification to decision rule

Human buyers realize that a bid below recent trade prices can be made at low cost, since bids can be replaced. Consequently, human buyers often make bids below recent trade prices, and bids progress up toward recent trade prices. Fig. 3(b) in Section 2.3.1 shows all bids in period 1 of baseline session 1. Twenty-two of 28 bids were below the average price ($\bar{p} = 92.8$) during the period shown in the figure (and only five of those 22 bids were accepted). Five of the six bids placed at values above the mean price were accepted. The figure demonstrates that human subjects frequently place bids that are below prices where trades have occurred.

In the HBL model, a buyer's belief $q_i(b)$ represents the buyer's assessment of the probability that a bid at b will be accepted by some seller. Maximization of expected surplus relative to this belief is too aggressive, in the sense that it ignores the option to replace bids that are not accepted. Formulation of a continuation value for the decision problem would be desirable, but a number of obstacles present themselves. Among these obstacles are the non-stationarity of beliefs, the fact that buyers typically purchase multiple units, and the unknown number of bidding opportunities for a buyer in a trading period. If none of these problems were present, a buyer with n remaining opportunities to place a bid b would have the continuation value

$$\begin{aligned} \mathrm{E}[S_i(b)] &= (v-b)\{q(b) + (1-q(b))q(b) + \cdots + (1-q(b))^{n-1}q(b)\} \\ &= (v-b)(1-(1-q(b))^n). \end{aligned} \tag{B.1}$$

The purpose of the decision rule modification is to decrease the expected surplus maximizing bid. Proposition 1 shows that the bid that maximizes expected surplus is lower when the surplus function is $\tilde{S}_i(b) = (v-b)q(b)^{1/n}$. The model has been tested with this modification.[12]

[12] The function $\tilde{S}(b)$ has been used as the alternative objective function. This form for the objective was selected prior to formulation of the continuation value argument described above and characterized in Eq. (B.1).

Proposition 1. *Let* b_k^* *be the argmax of* $S(b) = (v - b)q(b)$ *on* $[d_k, d_{k+1})$, *and let* $\tilde{b}_k^*$ *be the argmax of* $\tilde{S}(b) = (v - b)q(b)^{1/n}$ *on* $[d_k, d_{k+1})$. *Then* (1) *on each interval* $[d_k, d_{k+1})$ *the argmax* $\tilde{b}_k^*$ *of* $\tilde{S}(b)$ *is less than the argmax* b_k^* *of* $S(b)$, *and* (2) *if* $\tilde{S}(\tilde{b}_k^*) > \tilde{S}(\tilde{b}_{k-1}^*)$ *then* $S(b_k^*) > S(b_{k-1}^*)$.

Proof. The idea of the proof of Proposition 1 is straightforward: assertion (1) states that the local argmax of $\tilde{S}(b)$ on each interval $[d_k, d_{k+1})$ is less than the local argmax of $S(b)$. Assertion (2) guarantees that the global argmax of $\tilde{S}(b)$ lies in either the same interval $[d_k, d_{k+1})$ as the global argmax of $S(b)$, or it lies in an interval to the left of the one where the global argmax of $S(b)$ occurs. Consequently, $\tilde{b}^* < b^*$.

If the slope of $\tilde{S}(b)$ is negative at b_k^* then $\tilde{S}(b)$ is increasing to the left of b_k^*. Differentiate $\tilde{S}(b)$ to get $\tilde{S}'(b) = (v - b)(1/n)q(b)^{1/n-1}q'(b) - q(b)^{1/n}$. From the definition of b_k^*, $S'(b_k^*) = 0$. Since $S'(b) = (v - b)q'(b) - q(b)$ it follows that $(v - b_k^*)q'(b_k^*) = q(b_k^*)$. Substitute this expression for $(v - b_k^*)q'(b_k^*)$ into $\tilde{S}'(b)$ evaluated at b_k^* to get

$$\begin{aligned}\tilde{S}'(b_k^*) &= (v - b_k^*)\frac{1}{n}q(b_k^*)^{1/n-1}q'(b_k^*) - q(b_k^*)^{1/n} \\ &= \left(\frac{1}{n} - 1\right)q(b_k^*)^{1/n}.\end{aligned}$$

For all $n > 1$, this expression is negative, so assertion (1) holds.

Next, consider assertion (2). The problem is to show that when $\tilde{S}(\tilde{b}_k^*) > \tilde{S}(\tilde{b}_{k-1}^*)$ then $S(b_k^*) > S(b_{k-1}^*)$. From the definition of $\tilde{S}(\tilde{b}_k^*)$, the first condition is $(v - \tilde{b}_k^*)q(\tilde{b}_k^*)^{1/n} > (v - \tilde{b}_{k-1}^*)q(\tilde{b}_{k-1}^*)^{1/n}$. Since $\tilde{b}_{k-1}^*$ maximizes $\tilde{S}(b)$ on $[d_{k-1}, d_k)$, it follows that

$$(v - \tilde{b}_k^*)q(\tilde{b}_k^*)^{1/n} > (v - \tilde{b}_{k-1}^*)q(\tilde{b}_{k-1}^*)^{1/n} > (v - b_{k-1}^*)q(b_{k-1}^*)^{1/n}.$$

If the first and last terms in this inequality are raised to the power n the result is

$$(v - \tilde{b}_k^*)^n q(\tilde{b}_k^*) > (v - b_{k-1}^*)^n q(b_{k-1}^*),$$

which is equivalent to $(v - \tilde{b}_k^*)^{n-1}(v - \tilde{b}_k^*)q(\tilde{b}_k^*) > (v - b_{k-1}^*)^{n-1}(v - b_{k-1}^*)q(b_{k-1}^*)$. The product of the second and third terms on both sides of this inequality are expressions for expected surplus under the original belief function, so the last inequality can be written as $(v - \tilde{b}_k^*)^{n-1}S(\tilde{b}_k^*) > (v - b_{k-1}^*)^{n-1}S(b_{k-1}^*)$. Two short steps complete the proof. Write this last inequality as

$$\left(\frac{v - \tilde{b}_k^*}{v - b_{k-1}^*}\right)^{n-1} S(\tilde{b}_k^*) > S(b_{k-1}^*).$$

Finally, multiply the left side of the inequality by $S(b_k^*)/S(b_k^*)$ to get

$$\left(\frac{v - \tilde{b}_k^*}{v - b_{k-1}^*}\right)^{n-1} \frac{S(\tilde{b}_k^*)}{S(b_k^*)} S(b_k^*) > S(b_{k-1}^*).$$

Since $v - \tilde{b}_k^* < v - b_{k-1}^*$ and $S(\tilde{b}_k^*) < S(b_k^*)$, it follows that $S(b_k^*) > S(b_{k-1}^*)$. □

B.2. Modification to timing

The second problem with the original HBL model is that bids frequently exceed recent prices. This is also addressed with a simple modification: the expected time until a buyer submits her bid is decreased when current maximum expected surplus for the buyer exceeds the expected surplus that the buyer would attain at the average of recent prices. In the modified definition of the random variable that governs the timing of bids, let $\bar{p}_m$ be the average price over the last m trades, and let $\bar{S}_i \equiv \max\{0, \mathrm{E}[\tilde{S}_i(\bar{p}_m)]\} = \max\{0, (v_i - \bar{p}_m)q(\bar{p}_m)^{1/n}\}$. The timing modification is implemented by taking one plus the expected surplus at the average price, $1 + \bar{S}_i$ divided by one plus the current maximum expected surplus $1 + \tilde{S}_i^*$ for the buyer, and raising this ratio to a positive power c. The modified specification of $\lambda_i(\tilde{S}_i^*, t_{\kappa-1}, T)$ is

$$\tilde{\lambda}_i(\tilde{S}_i^*, \bar{S}_i, t_\kappa, T) = \begin{cases} \beta_i(T - \alpha_i t_\kappa)/(\tilde{S}_i^* T) \times ((\bar{S}_i + 1)/(\tilde{S}_i^* + 1))^{c} & \text{if } \bar{S}_i > 0, \\ \beta_i(T - \alpha_i t_\kappa)/(\tilde{S}_i^* T) & \text{if } \bar{S}_i \leqslant 0, \end{cases} \tag{B.2}$$

where $\alpha_i \in [0, 1)$, $\beta_i > 0$, and $\tilde{S}_i^*$ is the maximum expected surplus of agent i.

If the current standing bid sb is greater than or equal to $\bar{p}_m$ then $\bar{S}_i > \tilde{S}_i^*$, so $((\bar{S}_i + 1)/(\tilde{S}_i^* + 1))^{c} > 1$. Consequently, $\tilde{\lambda}_i > \lambda_i$ in this case: the mean time until a bid increases once the high bid reaches the mean of recent prices.

If the current standing ask sa is less than $\bar{p}_m$ then a bid $b = sa$ results in a trade with probability one, so $\tilde{S}_i^* \geqslant v - sa > v - \bar{p}_m \geqslant (v - \bar{p}_m)q(\bar{p}_m) = \bar{S}_i$. Consequently, $\tilde{\lambda}_i < \lambda_i$: the mean time until a bid decreases when the low ask is below the mean of recent prices.

B.3. Assessment of modifications

These two features of human behavior in the double auction – initial bids below equilibrium and reluctance to increment bids above recent prices – lead to a difference in the original HBL model between the mean price of trades initiated by sellers ($\bar{p}_a$) and the mean price for trades initiated by buyers ($\bar{p}_b$). This is because asks by human sellers are typically decreased to values just above or at the mean of recent trades, and bids by human buyers are increased to prices just below or at the mean of recent trade prices. A trade between a human seller and human buyer occurs when either the buyer accepts the seller's ask (which is typically no lower than the mean of recent trade prices) or a seller accepts the buyer's bid (which, similarly, is typically no higher than the mean of recent trade prices). Consequently, on average trades between human sellers and human buyers that are initiated by sellers occur at higher prices than trades initiated by buyers.

As a result of bids that frequently exceed recent prices and asks below recent prices, this statistic is reversed in the model. The two modifications above are intended to address this problem. Examination of the price difference between trades initiated by sellers and those initiated by buyers gives a measure of the extent to which the issue has been resolved. Table 10 shows the effects of these two

Table 10
Difference between sellers' and buyers' accepted price proposals ($\bar{p}_a - \bar{p}_b$)

	$c=0$	$c=3$	$c=5$	$c=7$	$c=8$
Myopic objective	−2.29	−1.21	−0.89	−1.02	−1.04
Modified objective	−1.44	−0.45	−0.27	**−0.15**	**−0.04**

modifications on the difference between prices initiated by sellers and those initiated by buyers for groups of 10 simulations in each of the 10 cells of the table. In all cells except the two with $c = 7$ and $c = 8$ and the modified objective, it is possible to reject the hypothesis that the mean difference is positive. Across five baseline sessions with human subjects the mean of this statistic was 0.45, and for these baseline sessions, it is not possible to reject the hypothesis that the mean difference is positive. It is clear from this table that the model mimics this aspect of human behavior more effectively as a result of the combination of the modifications described above, and that either of these two modifications independently does not resolve this discrepancy between model and human performance. In simulations and hybrid sessions, $c = 5$.

Appendix C

Before trading begins, a seller goes through an instruction set that describes elements of the seller screen and their operations. These elements include: (1) the seller's endowment and allocation of currency and commodity; (2) the seller's induced unit costs; (3) input that the seller provides during the session; (4) processing of input (by the double auction mechanism) from the seller, as well as from other sellers and from buyers; and (5) the determination of profits for the seller. Buyers' instructions are analogous to sellers' instructions. The most important principle adhered to in our implementation of instructions is that inputs of asks by sellers and bids by buyers are initiated by subjects, in order to avoid creation of a reference point that would influence asks or bids once the session begins.

Each seller and each buyer views a total of 27 screens, and makes inputs on five of these screens. There are six questions that each seller and each buyer must answer correctly in order to proceed. Most of the screen space is devoted to the actual screen display the subject will use during the experiment session: a text box describes this screen display. The total length of the instructions is 2,898 words, which is equivalent to approximately six pages of text. Subjects typically complete instructions in 15–35 min.

The instruction summary below refers frequently to elements of the seller's screen, which is shown in Fig. 2. The instructions for a buyer are exactly analogous. Direct experimenter interaction with subjects was kept to a minimum whenever possible, including sign-in, seating, and payment.

Screen 1. The subject's earnings are based on subject's decisions and the decisions of other participants.

Screen 2. The subject will be a seller throughout an experiment session that lasts 15 trading periods.

Screen 3. The subject should not communicate with or distract others. The subject's data are anonymous.

Screen 4. Payment will be made at an exchange rate of \$0.015 per unit of experiment currency accumulated. (This rate was £0.015 at the University of York and \$0.012 at the IBM Research Center.) Payment will be made anonymously at the conclusion of the experiment.

Screen 5. There is a set of interactive instructions that follow this screen. Subject will know that they either have or have not completed instructions based on the status message at the top of their screen.

Screen 6. Each period of the session consists of a 'Trading Phase' of 180 s and a 'Review Phase' of 30 s. A clock at the top of the screen ticks down to the end of each phase.

Screen 7. The seller begins the period with currency and commodity. The current balance of each is shown on the seller's screen.

Screen 8. The seller has a list of unit costs, one for each of the units that the seller can offer for sale. The cost for a unit is incurred only if the unit is sold. Unit profit is calculated as the difference between the sale price for the unit and the cost of the unit.

Screen 9. Total profit is the sum of the profit earned on individual units. Payment at the end of the session is proportional to the total profit earned. Unit costs are initialized at the beginning of each period; unsold units are not carried over to the next period. If a unit is sold at a price below its cost, a loss is incurred on that unit. (Gains are shown in black type; losses are shown in red type.)

Screen 10. An ask is entered in the 'Enter Ask' area of the seller screen. The subject enters an ask at this point.

Screen 11. The ask entered by the subject appears in the Unit Ask row of the Trade Summary display. Commodity Holding remains unchanged, but Commodity Available is reduced by one unit. The ask also appears in the 'Your Asks' display area.

Screen 12. Asks by all sellers and bids by all buyers appear in the 'Market Queue' display.

Screen 13. A new ask is generated randomly to simulate an ask by another seller. The simulated ask is higher than the subject's own ask, to illustrate the Ask Improvement rule. (The new ask is generated randomly from one to 10 above the seller's own ask, though the subject is not informed of the process that generates the new ask.) The subject is also informed that if he makes a new ask at this point, the new ask will replace his current ask.

Question 1. The subject has a quiz question appear at this point. The subject, after seeing a description of the Ask Improvement rule, is prompted to state what is the highest ask that can be submitted at this point.

Question 2. After the subject answers Question 1 correctly, he is asked to state the number of asks that will be in the Market Queue if he submits a new ask (to test comprehension of the ask replacement rule).

Screen 14. The subject is prompted to enter a new ask that improves on the ask submitted by the simulated seller.

Screen 15. The new ask replaces the seller's previous ask. The Messages and Trade Summary displays are updated.

Screen 16. The subject is prompted to remove his ask by double clicking on the ask in the 'Your Asks' display. Changes to the seller screen are described, including updates to the Commodity Available, Your Asks, Market Queue, Trade Summary, and Messages displays.

Screen 17. The seller is prompted to enter a new ask.

Screen 18. Updates to the seller's screen that result from the new ask are reviewed.

Screen 19. A trade occurs when a seller's ask is at or below the current best bid, or a buyer's bid meets or exceeds the best ask. A bid from a simulated buyer is generated that meets the seller's ask, so that the seller trades with the simulated buyer.

Screen 20. Unit Price and Unit Profit entries in the Trade Summary table are described. Changes to the Currency Holdings and Commodity Holdings areas are described. The Market Transaction Prices graph update is described.

Screen 21. The price determination rule is reviewed and the subject is informed that after two questions, a bid will be simulated that results in a trade with the current low ask.

Question 3. The subject is asked whether a trade will result if a buyer submits a bid below the current low ask.

Question 4. The subject is asked what the trade price will be if a bid is submitted that meets or exceeds the current low ask.

Screen 22. A bid is simulated that generates another trade. (This trade is between a simulated seller and a simulated buyer, so the subject only sees public information regarding the trade, i.e., its price on the Market Transaction Prices graph.)

Screen 23. A new simulated bid appears in the Market Queue. The price determination rule is reviewed once more and the subject is asked two more questions.

Question 5. The seller is asked whether a trade will result if he submits an ask that exceeds the current high bid.

Question 6. The seller is asked what the trade price will be if he submits an ask that is below the current high bid.

Screen 24. The subject is prompted to enter an ask that is at or below the current high bid, in order to produce a new trade.

Screen 25. Changes to the Current Allocation, Trade Summary, and other screen displays that result from the most recent trade are reviewed.

Screen 26. The seller is informed of the Vote to End Period option, and the unanimity rule that triggers an early end to the trading phase of the current period.

Screen 27. A Period Profit window appears during the review phase of each period.

Screen 28. Subjects are cautioned that the ask by another seller and the bids by other buyers in the instructions were simulated and that these may not be similar to the responses by buyers and by other sellers during the experiment. The subject is informed that he has now completed the instructions and trading will begin when all subjects have completed their instructions.

References

Cason, T.N., Friedman, D., 1996. Price formation in double auction markets. Journal of Economic Dynamics and Control 20, 1307–1337.

Easley, D., Ledyard, J.O., 1993. Theories of price formation and exchange in double oral auctions. In: Friedman, D., Rust, J. (Eds.), The Double Auction Market: Institutions, Theories, and Evidence. Addison-Wesley, Reading, MA, pp. 63–97.

Friedman, D., 1991. A simple testable model of price formation in the double auction market. Journal of Economic Behavior and Organization 15, 47–70.

Gjerstad, S., Dickhaut, J., 1998. Price formation in double auctions. Games and Economic Behavior 22, 1–29.

Gode, D.K., Sunder, S., 1993. Allocative efficiency of markets with zero intelligence traders: market as a partial substitute for individual rationality. Journal of Political Economy 101, 119–137.

Rubinstein, A., 1982. Perfect equilibrium in a bargaining model. Econometrica 50, 97–110.

Smith, V.L., 1982. Microeconomic systems as an experimental science. American Economic Review 72, 923–955.

Smith, V.L., Williams, A.W., 1982. The effects of rent asymmetries in experimental auction markets. Journal of Economic Behavior and Organization 3, 99–116.

Williams, A.W., Smith, V.L., 1984. Cyclical double-auction markets with and without speculators. Journal of Business 57, 1–33.

Wilson, R.B., 1987. On equilibria of bid–ask markets. In: Feiwel, G.W. (Ed.), Arrow and the Ascent of Modern Economic Theory. New York University Press, New York, pp. 375–414.