# THE STRATEGIC IMPACT OF PACE IN DOUBLE AUCTION BARGAINING

Steven Gjerstad

Department of Economics
McClelland Hall 401
University of Arizona
Tucson, AZ 85721

March 24, 2003

This paper evaluates performance of human subjects and instances of a bidding model that interact in continuous-time double auction experiments. Asks submitted by instances of the seller model ("automated sellers") maximize the seller's expected surplus relative to a heuristic belief function, and arrive stochastically according to an exponential distribution. Automated buyers are similar. Across experiment sessions we vary the exponential distribution parameters of automated sellers and buyers in order to assess the impact of the relative pace of asks and bids on the performance of both human subjects and the automated sellers and buyers. In these experiments, prices converge and allocations converge to efficiency, yet the split of surplus typically differs significantly from the equilibrium split. In order to evaluate the impact of pace, a statistical model is developed in which the relative performance of sellers to buyers is examined as a function of the profile of types present in each experiment session. This econometric model demonstrates that (1) human buyers outperform human sellers, (2) automated sellers and buyers with a longer expected time between asks or bids outperform faster automated sellers and buyers, and (3) the performance of the faster automated buyers is comparable to that of human buyers.

Keywords: Double auction, experimental economics, bounded rationality

# 1 Introduction

Double auction experiments attain competitive equilibrium (CE) allocations and prices for a wide range of economic environments. Although experiments directly reveal subjects' behavior and their performance, their decision rules and the impact of these rules on individual and aggregate performance are not directly revealed in experiments. Simulations and analytical models partially overcome this limitation: models specify decision rules, realized decisions are observed, and outcomes can be compared to aggregate market performance statistics from experiments, such as efficiency, price variability, and relative performance of sellers and buyers.

Several models of bargaining in the double auction have been formulated in the last 15 years, including Wilson [1987], Friedman [1991], Easley and Ledyard [1993], Gode and Sunder [1993], and Gjerstad and Dickhaut [1998]. Wilson formulates a Bayesian equilibrium model that extends models of bilateral bargaining to the multilateral bargaining of the double auction. Friedman models conditions on final exchanges in a trading period that guarantee that final trades are in the vicinity of equilibrium prices. Easley and Ledyard model a process in which sellers update beliefs about asks which will be accepted by buyers and buyers update beliefs about bids which will be accepted by sellers. Sellers in their model expect that prices in the current period will be no lower than the minimum of the lowest ask and lowest trade price in the previous period and no higher than the maximum of the highest trade price and the highest bid in the previous period. Sellers then randomly select their asks from this interval. Beliefs and actions of buyers in the Easley-Ledyard model are identical to those of sellers. Since asks and bids in their model are drawn uniformly from the price range in the previous period, the interval contracts with positive probability in each period, and prices therefore converge across periods with replicated supply and demand conditions. Gode and Sunder do not attempt to model human behavior in the double auction. Instead, they consider the extreme situation in which sellers and buyers do not update beliefs or learn from previous market activity, but simply choose an individually rational ask or bid. Gode and Sunder demonstrate through simulations that even in this case, allocations are nearly efficient with high probability.

The model by Gjerstad and Dickhaut, like the Easley-Ledyard model, develops a belief for each seller that his ask will be accepted by some buyer, and also develops a belief for each buyer that her bid will be accepted by some seller. These beliefs are formed on the basis of observed market data, including frequencies of asks, bids, accepted asks, and accepted bids. Then buyers and sellers choose actions that maximize their own expected surplus. Asks and bids in this model are submitted at random times, with an exponential distribution that depends on both the expected surplus of the seller or buyer, and on the time remaining in the period. We refer to this model as *Heuristic Belief Learning* (HBL) due to the central role of the belief formation element of the model. Simulations of this model demonstrate that the asks by sellers and the bids by buyers lead to efficient outcomes and stable prices, as in experiments with human subjects. The HBL model of seller and buyer behavior formulated in Gjerstad and Dickhaut [1998] is specified so that instances of the sellers and buyers can be embedded into an experiment with human subjects and performance of the sellers and buyers from the model can be compared directly to human counterparts. This is the premise of the research reported in this paper.

Interaction between model strategies and human counterparts in experiments simultaneously allows researchers to assess the effects of specific elements of the model of behavior, as well as the performance of model behavior relative to human behavior. In particular, when parameters that govern timing distributions for asks and bids by automated sellers and buyers are varied across experiment sessions, we observe their impact on statistical aggregates of human behavior and performance, and thereby identify the impact of this strategic variable. In this paper, a modified version of the original HBL model is described, parameters are specified for the model, and model performance is examined through experiments that involve direct interaction between instances of the model and human subjects. This approach has two methodological advantages. One advantage is that when automated sellers and buyers interact with human sellers and buyers, parameters for automated sellers and buyers can be calibrated so that their performance is similar to that of human subjects. In the experiments reported in this paper, one specification of parameter values for automated sellers and buyers leads to aggregate performance of automated buyers that is similar to performance by human buyers. Another advantage of this methodology is that it expands the types of subject response behaviors that can be measured. Experiments typically examine behavioral response either to changes in the economic institution or game form or to changes in the economic environment (preferences, endowments, or production technologies), but by varying the timing of asks and bids by automated sellers and buyers we can observe changes in statistical aggregates of human response and performance as a function of changes to model parameters that affect behavior by instances of the model.

Analysis of data from these experiments clearly shows that prices converge and that markets converge toward efficient allocations with human subjects, in simulations, and when automated sellers and buyers interact with human subjects.[1] These observations are consistent with previous double auction experiments and with simulations reported in Gjerstad and Dickhaut [1998]. Although it has been noted previously that sellers frequently underperform buyers in double auction experiments – even when the surplus at equilibrium is equally split between sellers and buyers – this observation has not been quantified previously. The magnitude of the difference between performance of sellers and buyers is surprisingly large in experiments, in simulations that include a difference between the pace of asks and the pace of bids, and in hybrid experiments. Perfect price convergence and perfect efficiency guarantee that surplus is evenly split between sellers and buyers in a symmetric market (provided that the equilibrium price is unique), but even with low price variability and small efficiency losses, the difference in relative performance of the two sides of the market can be large.[2] Since the split of surplus between sellers and buyers frequently differs

[1] For brevity, throughout this paper a market conducted with human subjects is referred to as an *experiment*, a market with only automated sellers and buyers is referred to as a *simulation*, and a market that combines human subjects and automated sellers or buyers is referred to as a *hybrid experiment*.

[2] If supply and demand are linear and the slope of demand is the negative of the slope of supply, then the fraction of surplus lost is the square of the fractional difference of price from the equilibrium price, but the difference between the performance of sellers and buyers is approximately four times the fractional difference between the price and the equilibrium price. For example, if the price is 10% below the equilibrium price by (that is, if the price is below the equilibrium price by 10% of the difference between the equilibrium price and the price at which supply is zero), then the efficiency loss is 1% yet buyers' earned surplus exceeds sellers' earned surplus by approximately 40%. This issue is examined in greater detail in Section 4.3.

substantially from an equal split – even when the difference between the equilibrium price and the mean realized price is small – it is important to determine which aspects of the strategic behavior of sellers and buyers lead to the difference in relative performance. To address this question, a statistical model is developed which treats the relative performance of sellers and buyers as a function of the combination of types present in the experimental design. In order to identify the effect of timing decisions, three types are included on each side of the market. Seller types employed in the experimental design are human sellers, slow automated sellers, and fast automated sellers; buyer types are human buyers, slow automated buyers, and fast automated buyers.[3] 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 side of the market that has the slower pace. Consequently, we conclude that the pace of asks by sellers and the pace of bids by buyers are important strategic decisions in the temporally unstructured bargaining of the double auction.

This paper is organized as follows. Section 2 includes a description of the economic environment tested in the experiments, simulations, and hybrid experiments and of the double auction mechanism. Section 2 also includes a summary of and modifications to the *Heuristic Belief Learning* model of seller of buyer behavior. The experiment design is described in Section 3. Section 4 provides analysis of experiment results. Conclusions are drawn in the final section.

## 2 The microeconomic system

Formulation of market experiments is separated into three elements, which correspond to elements of analysis in the mechanism design literature. In the context of these market experiments, the elements are (1) the economic environment induced in the experiments, simulations, and hybrid experiments, (2) the double auction mechanism, and (3) the HBL

[3] The designations "slow" and "fast" are relative, not absolute, although in the analysis in Section 4.5 we are able to benchmark performance of each of the two types of automated sellers and the two types of automated buyers present in simulations and hybrid experiments relative to their human counterparts.

model of behavior. Human behavior is also a factor in experiments and hybrid experiments, but decision rules employed by human subjects are latent (only the actions that result from their decision rules are observed), so evaluation of human behavior is included in Section 4 with analysis of experimental data.

## 2.1 The economic environment

Experiments were conducted with a set $\mathcal{I}$ of subjects and HBL model instances partitioned into a set $\mathcal{I}_\mathcal{S}$ of sellers and a set $\mathcal{I}_\mathcal{B}$ of buyers, with $|\mathcal{I}_\mathcal{S}| = 6$ and $|\mathcal{I}_\mathcal{B}| = 6$. Sellers had a list of unit costs for either 10, 11, or 12 units of a commodity, and buyers had a list of values for 10, 11, or 12 units of this commodity, as in table 1. A seller incurred the cost for a unit each time that a trade was negotiated with some buyer and received currency from the buyer equal to the negotiated price. Each buyer had a currency endowment that was greater than the sum of her unit values, in order that she could purchase all units for which she had a positive unit value, if she chose to do so. Each of the twenty-eight experiments and hybrid experiments reported in this paper had the same basic supply and demand conditions, although there were four variants on this design that differed by constants which were added to all unit values and all unit costs. The base design, referred to as ICV2a, is shown on the left side of figure 1. Costs for each seller and values for each buyer in this design are shown in table 1. Units that could be traded profitably at prices in the equilibrium price range [95, 97] are shown in bold type. A second supply and demand condition, which differs from the ICV2a design by a shift of 12 to each unit cost and each unit value, is shown on the right side of figure 1. The other two versions of the supply and demand schedule used in the experiments, ICV2c and ICV2d, differed from the ICV2a design by increases of 24 and 40 to all unit costs and all unit values.

These supply and demand schedules are symmetric with respect to the equilibrium price $p_e :$ for each unit value $v_i$ on the demand schedule that yields surplus $v_i - p_e$ to the buyer if traded at the equilibrium price, there is a unit cost $c_{i'}$ on the supply schedule that yields surplus $p_e - c_{i'} = v_i - p_e$ if traded at the equilibrium price.

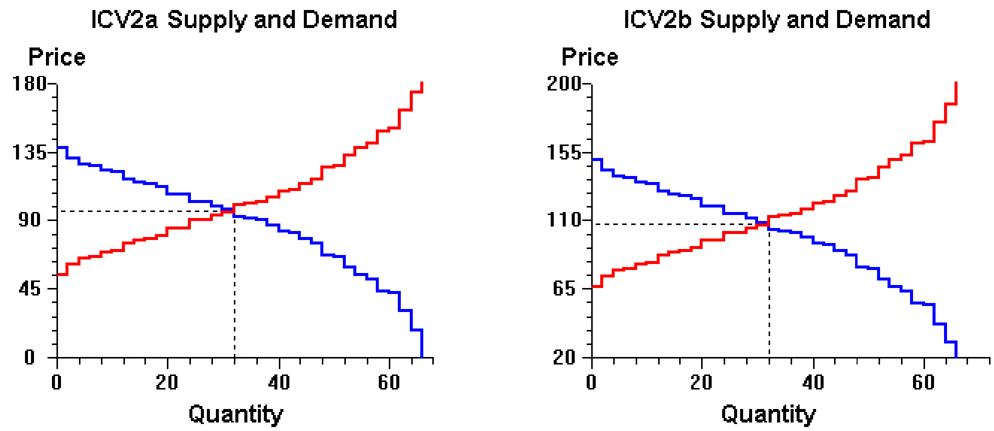


Figure 1: Supply and demand configurations for ICV2a and ICV2b.

| | 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 |

Table 1: Sellers' unit cost and buyers' unit value schedules for ICV2a series.

### 2.2 The double auction mechanism

The double auction is operationally simple, robustly efficient, and leads to stable transaction prices. In this mechanism, any seller may submit an ask at any time during a trading period. An ask is entered in the area labelled "Enter Ask" on the sellers screen display, which is shown in figure 2. This ask represents the seller's current report of the lowest price that he is willing to accept for a unit of an abstract "commodity." Similarly, buyers may submit a bid at any time, which represents the buyer's current report of the highest price that she is willing to pay for a unit of the commodity. If an ask is placed that is at or below the current high bid, a trade results. Similarly, if a bid is placed that meets or exceeds the current low ask, a trade occurs. Sellers and buyers may make any number of asks or bids, and may trade any number of units that is consistent with their endowments. Several

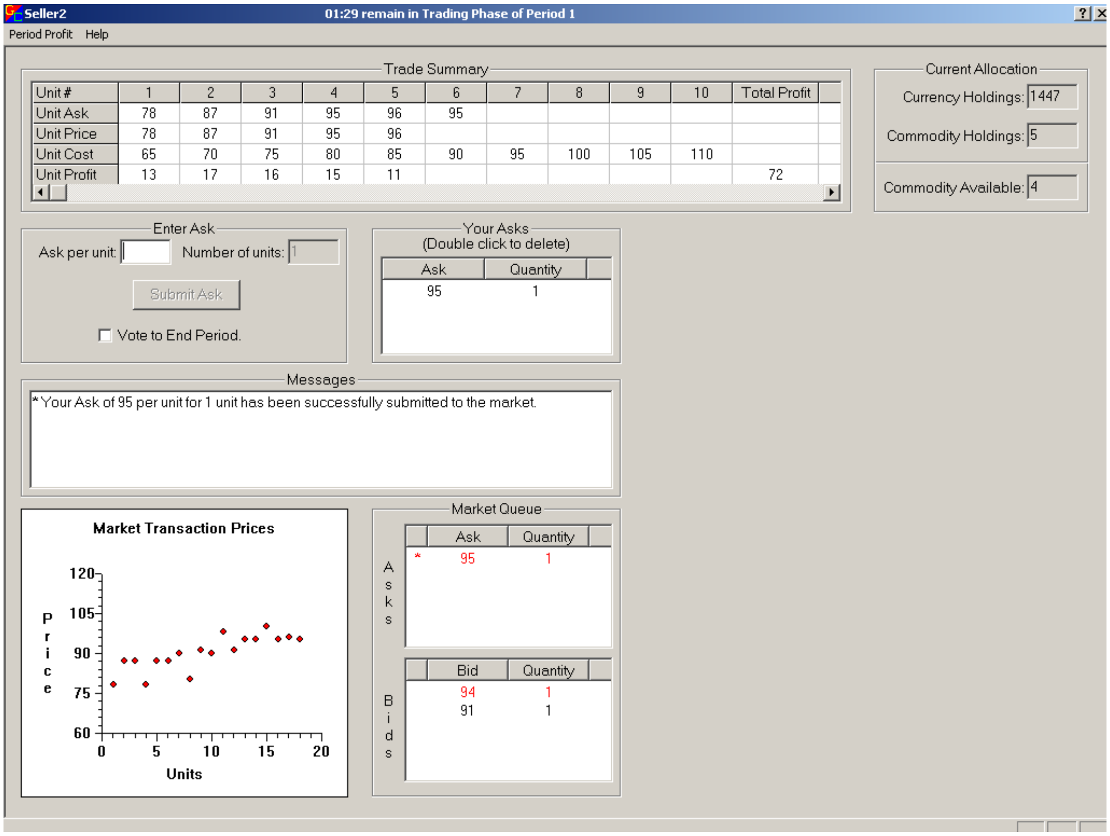


Figure 2: Seller screen with elements of market institution.

specific rules are implemented in the version of the double auction that was used to conduct the experiments, simulations, and hybrid experiments reported here. Of these, the most important is the "spread reduction rule," which requires that each new ask is made at a value that is below the current low ask and each new bid is placed at a higher value than the current high bid. Sellers and buyers have the option to remove any ask or bid that they have previously made, provided the request to remove the ask or bid is received before it results in a trade. Sellers are permitted a single ask at any given time and any new ask by a seller replaces his previous ask if he has one in the market queue. Each ask is the unit price offered by the seller for a single unit: multiple unit trades are not permitted. Similarly, buyers also are permitted a single bid at any given time and that bid is for one unit.

Throughout each trading period, a queue on the screen of each seller displays all current asks and all current bids (shown as the "Market Queue" in figure 2); the screen of each buyer also displays both queues. When a seller successfully enters an ask into the ask queue, he receives a confirmation message in the “Messages” area of the screen display. Similarly, a buyer receives a confirmation message when she enters a bid into the bid queue. When a seller and buyer complete a trade, they both receive a confirmation message, and the trade price is included on a graphical display of all trade prices from the current period, shown as the “Market Transaction Prices” graph. The length of each trading period was known to each seller and to each buyer, and a clock on the screen of each seller and buyer showed the time remaining in the period. Each of these elements of the market institution appears on a seller’s trading screen, as shown in figure 2. Buyers have a similar trading screen. Operations of these elements of the double auction mechanism are explained to subjects in a detailed interactive instruction set.

## 2.3 Heuristic Belief Learning

Each seller in the HBL model forms a belief that his ask will be accepted by some buyer; similarly, each buyer forms a belief that her bid will be accepted by some seller. These beliefs are formed on the basis of observed market activity, including frequencies of asks, bids, accepted asks, and accepted bids. Then each seller chooses an ask and each buyer chooses a bid that maximizes his or her own expected surplus. This model converges quickly to efficient allocations, with stable transaction prices. In the conclusions of their paper, Gjerstad and Dickhaut report two differences between simulation data and data from experiments. One reported difference between the HBL model and human behavior is that in the model asks by sellers and bids by buyers are made initially at values that are closer to equilibrium than those of human subjects (that is, in the HBL model initial asks are too low and initial bids are too high). The other observed difference is that once the low ask reaches recent trade prices, the probability of another ask remains close to one half, whereas most human sellers are much more reluctant to decrease their ask once the low ask reaches recent trade prices. This has the consequence that in the HBL model, trades are frequently initiated by sellers at prices below the average of recent prices. Analogous behaviors by buyers in the HBL model are observed: buyers’ bids are increased too frequently after the high bid reaches the average of recent trade prices. After a summary of the HBL model, these two differences are described in more detail, and two modifications are described which mitigate the effects of these differences between the behavior specified in the model and human behavior.

### 2.3.1 Elements of the Heuristic Belief Learning model

An action by a seller or buyer in the HBL model has three primary elements: a heuristic belief function, maximization of expected surplus relative to this belief, and a choice of the timing of an ask or bid. The heuristic belief function is based on a transformation of empirical acceptance frequencies for asks and bids. We begin with a description of these empirical acceptance frequencies, develop heuristic belief functions based on these acceptance frequencies, and then describe expected surplus maximization and the timing of sellers' asks and buyers' bids.

**Empirical acceptance frequencies**

A seller's heuristic belief function is based on empirical acceptance frequencies for asks within the seller's memory. For each ask $a$ that has occurred within the history maintained by the seller, the empirical acceptance frequency for $a$ is defined as the number of acceptances of asks at $a$, divided by the number of asks that have been placed at $a$. Let $A(a)$ denote the number of asks that have been made at $a$, and let $TA(a)$ denote the number of those asks that have been accepted (or taken) by some buyer. For each seller $i \in \mathcal{I}_{\mathcal{S}}$, the empirical frequency function is

$$\check{p}_i(a) = \frac{TA(a)}{A(a)}.$$

The empirical acceptance frequency function for a buyer is similar. Let $B(b)$ be the bids at $b$ and let $TB(b)$ be the taken bids at $b$. The empirical frequency function for buyer $i \in \mathcal{I}_{\mathcal{B}}$ is

$$\check{q}_i(b) = \frac{TB(b)}{B(b)}.$$

Figure 3 shows both the frequencies of bids (as the solid line) and of accepted bids (as the dashed line) in intervals of width 5 from 50 to 100 during period 1 of the experiment ICV2a-DA-010606b. This experiment has the supply and demand shown on the left side of figure 1, which has an equilibrium price range $p_e = [95, 97]$. The mean price during period 1 was 92.8, and half of the bids in the interval $[91, 95]$ were accepted.

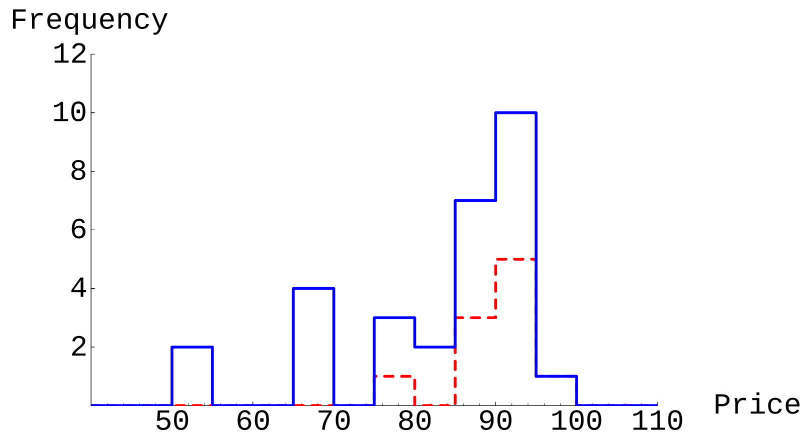


Figure 3: Bids and accepted bids in period 1 of experiment ICV2a-DA-010606b.

There are several problems that are encountered in an effort to use the empirical frequency function $\check{q}_i(b)$ as the belief function for a buyer. The large action set for buyers combined with the small number of data points after a small number of bids produces an irregular belief function. In the example in figure 3, after 29 bids there are intervals on which $\check{q}_i(b)$ is undefined (even when bids are clustered into intervals of width 5), and on intervals where the empirical belief function is defined, it is not monotonic. This is undesirable because a seller is more likely to accept a higher bid, so a buyer's belief function should be monotonically increasing. Similar considerations apply to the empirical acceptance frequency function of each seller.[4]

**Heuristic belief functions**

In order to address these issues, the data that generate a seller's empirical frequency function $\check{p}_i(a)$ and a buyer's empirical frequency function $\check{q}_i(b)$ are transformed into *heuristic belief functions* $\widehat{p}_i(a)$ and $\widehat{q}_i(b)$ which are monotonic, and are defined for any number of asks, bids, and trades. This construction consists of four steps.

[4] Beliefs and decision rules of sellers and buyers in the HBL model are symmetric. There are several instances throughout this paper where a description of a seller's (a buyer's) belief, ask, or a property of his belief function or ask (her belief function or bid) is provided. In each of these cases there is an analogous statement for the opposite side of the market that is mentioned, but repetition of the symmetric argument is frequently less detailed when the analogous argument is obvious.

**Step 1** In the first step, the initial belief is defined. Let $M_i$ be an upper bound on the range for the belief of agent $i \in \mathcal{I}$. Buyer $i \in \mathcal{I}_{\mathcal{B}}$ 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 $i \in \mathcal{I}_{\mathcal{S}}$ 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 experiments, for sellers $M_i \equiv 2^{\,2+[\log_2(\max\{c_i\})]}$. For buyers, $M_i \equiv 2^{\,1+[\log_2(\max\{v_i\})]}$ where $[\log_2(\max\{v_i\})]$ is the greatest integer less than or equal to $\log_2(\max\{v_i\})$.

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

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

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 $i \in \mathcal{I}_{\mathcal{B}}$ is defined at each possible bid $b \in \mathcal{D}_{m,i}$ as

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

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 should reinforce the seller's belief that an ask at $a$ will be accepted, as should any bid at or above $a$ that has been made. Any ask at a value at or below $a$ that has not been accepted should decrease 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 $d \geq sa$ and setting $\widehat{q}_i(d) = 0$ for all $d \leq sb$.

**Step 4** The beliefs $\widehat{p}_i(a)$ and $\widehat{q}_i(b)$ are extended from $\mathcal{D}_{m,i}$ to $[0, M_i]$ by cubic spline interpolation. For each $d_k \in \mathcal{D}_{m,i}$, 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.

With this definition, Gjerstad and Dickhaut show that each seller's belief function $p_i(a)$ is monotonically decreasing, and each buyer's belief function $q_i(b)$ is monotonically increasing. Figure 4 shows a buyer's belief at the end of period 1 from experiment ICV2a-DA-010606b. In the simulations and hybrid experiments reported in Section 4, beliefs of the automated sellers and buyers are formed using the history of asks and bids that lead up to the last 12 trades ($m = 12$). The heuristic belief function $q_i(b)$ in figure 4 should be compared to the empirical belief function $\check{q}_i(b)$ in figure 3, since the motivation for the definition of the heuristic belief function is to develop a transformation of the data from the empirical belief function into a monotonic function.

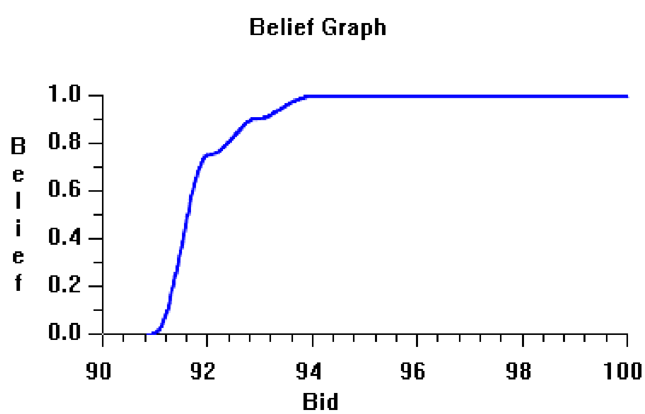


Figure 4: Buyer's belief at end of period 1 in experiment ICV2a-DA-010606b.

**Expected surplus maximization**

For the belief function $p_i(a)$, the expected surplus that seller $i$ with unit cost $c_i$ attains when he places an ask $a$ is $E[S_i(a)] = (a - c_i)\, p_i(a)$. The expected surplus maximizing ask for this seller is $a_i^* = \arg\max E[S_i(a)]$. At this expected surplus maximizing ask $a_i^*$, the expected surplus is $S_i^* = E[S_i(a_i^*)] = (a_i^* - c_i)\, p_i(a_i^*)$. The expected surplus function, the expected surplus maximizing bid, and the maximum expected surplus for each buyer $i \in \mathcal{I}_{\mathcal{B}}$ is defined similarly.

**Timing of asks and bids**

Asks and bids are made at random times by sellers and buyers in the HBL model, according to exponential distributions. Assume that a total of $\kappa - 1$ asks or bids have occurred in a trading period and the times of these events are $\{t_1,\, t_2,\, t_3,\, \ldots,\, t_{\kappa-1}\}$. Immediately after the ask or bid at time $t_{\kappa-1}$, each agent $i \in \mathcal{I}$ calculates a time $t_{\kappa_i}$ to wait until his next ask or bid. For each agent $i \in \mathcal{I}$, the probability distribution for the wait time $t_{\kappa_i}$ is

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

The parameter $\lambda_i$ depends on the buyer's current maximum expected surplus $S_i^*$, on the current time $t_{\kappa-1}$ that has elapsed in the trading period, and on the total time $T$ in the trading period. With this specification of the wait time for each agent, $Pr[t_{\kappa_i} < t_{\kappa-1}] = 0$, and the wait time until the next ask or bid is $t_\kappa = \min\{t_{\kappa_i}\}_{i\in\mathcal{I}}$.

The specification of $\lambda_i(S_i,\, t_{\kappa-1},\, T)$ in the HBL model is

$$\lambda_i(S_i^*, t_{\kappa-1}, T) = \frac{\beta_i\,(T - \alpha_i\, t_{\kappa-1})}{S_i\, T} \tag{2}$$

where $\alpha_i \in (0,\, 1)$, $\beta_i \in (0,\, \infty)$, and $S_i^*$ is the maximum expected surplus for agent $i$.

For $\lambda_i(\cdot)$ as specified in equation (2), the mean time until agent $i$ with the CDF in equation (1) places his ask or bid is $\lambda_i$. The mean time decreases in $S_i^*:$ agents with greater maximum expected surplus are more anxious to trade. The parameter $\beta_i$ is a linear scale factor for the timing decision. Parameters $\{\beta_i\}_{i\in\mathcal{I}}$ are varied across experiment sessions in order to determine the effect of timing decisions on performance.[5] The factor $\alpha_i$, which is multiplied by the current elapsed time in the trading period, affects the extent to which the pace of a seller's asks or a buyer's bid increases as the period progresses. When $t = 0$, $\frac{(T - \alpha_i\, t)}{T} = 1$, and as $t \to T$, $\frac{(T-\alpha_i\, t)}{T} \to 1 - \alpha_i$. Surplus on units sold late in the trading period is typically much lower than it is early in the trading period. The factor $\frac{(T-\alpha\, t)}{T}$

[5] In hybrid experiments, all parameters $\{\beta_i\}_{i\in\mathcal{I}}$ are either $\beta_i = 250$ ("fast") or $\beta_i = 400$ ("slow"). Simulations were conducted with slow sellers and slow buyers, with slow sellers and fast buyers, and with fast sellers and slow buyers.

balances the effect of this difference by increasing the pace as the period progresses. In both the simulations and the hybrid experiments reported in Section 4, the value of $\alpha$ is 0.95.

### 2.3.2 Model modifications

Two aspects of the behavior of buyers and sellers in the model differ from human behavior and two modifications to the model have been developed that mitigate these two differences and thereby improve model performance. Sellers in the HBL model typically make initial asks near recent prices and subsequently they frequently place asks below recent prices. These two aspects of seller behavior in the model frequently result in prices proposed by sellers that are below the average of recent prices. In experiments (that is, in markets with only human subjects), a price below the average of recent prices typically occurs when a seller accepts a buyer's bid that was below the average of recent prices. Analogous differences between bids by buyers in the HBL model and bids by human buyers are observed. To summarize, in experiments trades at prices above the mean of recent prices are typically initiated by sellers and trades at prices below the mean of recent prices are typically initiated by buyers, whereas this pattern is reversed in the HBL model. Consequently, if we consider the difference $\bar{p}_a - \bar{p}_b$, where $\bar{p}_a$ is the mean price for trades initiated by sellers and $\bar{p}_b$ is the mean price for trades initiated by buyers, this difference is typically positive in experiments, and it is typically negative in simulations of the HBL model. Following the description of the modifications, we assess the effectiveness of the modifications by comparing values of $\bar{p}_a - \bar{p}_b$ in the original HBL model, the modified HBL model, and experiments.

**Modification to decision rule**

Human buyers realize that a bid below equilibrium can be made at low cost, since bids can be replaced. Consequently, human buyers frequently make bids below recent trade prices, and bids progress up toward recent trade prices. Figure 3 in Section 2.3.1 shows all bids in period 1 of experiment ICV2a-DA-010606b. Twenty-three of twenty-nine bids were below the average price during the period shown in the figure (and only four of those twenty-three bids were accpeted). Five of the six bids placed at values above the mean price were accepted. From the figure, we see 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 given bid 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 that a buyer will have 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

$$\begin{aligned} E[S_i(b)] &= (v-b)\,\{q(b) + (1-q(b))\,q(b) + \cdots + (1-q(b))^{n-1}\,q(b)\} \\ &= (v-b)\,(1-(1-q(b))^n). \end{aligned} \tag{3}$$

The aspect of model performance that motivates this change to the decision rule is that initial asks are too close to recent prices. The change to the decision rule should lead to a decrease in the expected surplus maximizing bid for the modified decision rule. Proposition 1 shows that the bid that maximizes expected surplus is lower when the surplus function is $\tilde{S}_i(b) = (v-b)\,q(b)^{1/n}$, and we have tested the model with this modification.[6]

**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 $\tilde{S}(b) = (v-b)\,q(b)^{\frac{1}{n}}$. 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 proof is provided in Appendix A.

Assertion (1) in the proposition 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, $\tilde{b}^* < b^*$.

[6] The function $\tilde{S}(b)$ has been used as the alternative objective function. This form for the objective was selected based on an attempt to employ an alternative to the myopic expected surplus function prior to formulation of the continuation value argument described above and characterized in equation (3).

**Modification to timing**

The second problem with the model is that bids frequently exceed recent prices. This has also been addressed with a simple modification: when current maximum expected surplus for a buyer exceeds the expected surplus that the buyer would attain at the average of recent prices, we decrease the expected time until the buyer submits another bid. 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 $\bar{S}_i \equiv \max\{0, E[S_i(\bar{p}_m)]\} = \max\{0, (v_i - \bar{p}_m)\, q(\bar{p}_m)\}$. The timing modification is implemented by taking the ratio of expected surplus at the average price $\bar{S}_i$ divided by the current maximum expected surplus $S_i^*$ for the buyer, and raising this ratio to a positive power $c$. The modified specification of $\lambda_i(S_i^*, t_{\kappa-1}, T)$ is

$$\tilde{\lambda}_i(S_i^*, \bar{S}_i, t_{\kappa-1}, T) = \begin{cases} \frac{\beta_i\,(T-\alpha_i\, t_{\kappa-1})}{S_i^*\, T} \left(\frac{\bar{S}_i+1}{S_i^*+1}\right)^c, & \text{if } \bar{S}_i > 0, \\ \frac{\beta_i\,(T-\alpha_i\, t_{\kappa-1})}{S_i^*\, T}, & \text{if } \bar{S}_i \leq 0, \end{cases} \tag{4}$$

where $\alpha_i \in (0,\, 1)$, $\beta_i > 0$, and $S_i^*$ is the maximum expected surplus of agent $i$.

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 $S_i^* \geq v - sa > v - \bar{p}_m \geq (v - \bar{p}_m)\, q(\bar{p}_m) = \bar{S}_i$. Consequently $\tilde{\lambda}_i < \lambda_i$: the mean time until a bid decreases when the low ask is below the mean of recent prices.

**Assessment of modifications**

The two observations about human behavior described above lead to a difference 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 typically 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 (as described in the first paragraph of Section 2.3.2), this statistic is reversed in the model. The

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

Table 2: Difference between sellers' and buyers' accepted price proposals ($\bar{p}_a - \bar{p}_b$).

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 2 shows the effects of these two modifications on the difference between prices initiated by sellers and those initiated by buyers for groups of ten simulations in each of the ten 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 experiments with human subjects the mean of this statistic was 0.45, and for this sample 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.

## 3 Experiment design

The economic environment, described in Section 2.1, consists of six buyers and six sellers. 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. Buyers also have three schedules, assigned similarly. This was done so that in "balanced" hybrid experiments – which include human sellers, automated sellers, human buyers, and automated buyers – there are human seller and automated seller counterparts in the market that have identical cost schedules and there are also human buyer and automated buyer counterparts in the market with identical value schedules.

The experiment design includes two factors. These are the fraction of human sellers and the fraction of human buyers, and the pace of asks by the automated sellers and the pace of bids by automated buyers. The fraction of agents is varied by implementing two conditions. In one condition, referred to as balanced, half of the buyers and half of the sellers were instances of the modified HBL model, and the other half on each side of the market were human subjects. In the "unbalanced" condition, either all of the buyers or all of the sellers were HBL model instances and human subjects filled all roles on the side of the market opposite the HBL model. This is represented in figure 5, where the horizontal axis is the percentage of human buyers, and the vertical axis is the percentage of human sellers. The origin represents 30 simulations (no human buyers or sellers). The upper right represents five baseline (all human) experiments. The fourteen observations in the center (ten with slow buyers and four with faster buyers, as explained below) are the balanced condition. The nine off-diagonal observations are the two unbalanced conditions, which are also split into experiments with a slower version of automated bidders and experiments with a faster version.

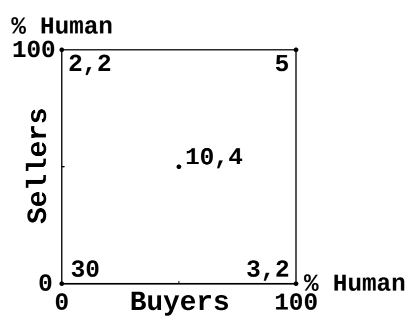


Figure 5: Design for human-model interaction experiments.

For the second treatment variable, we varied the parameters $\{\beta_i\}_{i\in\mathcal{I}}$ that govern the pace at which sellers and buyers in the model submit their asks and bids. The model specifies that asks by sellers and bids by buyers are submitted randomly according to exponential distributions with the parameters $\{\tilde{\lambda}_i\}_{i\in\mathcal{I}}$ in equation (4). Two values of the linear parameters $\beta_i$ in that equation were tested. These were slow automated sellers and buyers with $\beta_i = 400$ and fast automated sellers and buyers with $\beta_i = 250$. Since sellers' beliefs after several periods are approximately $p(a) = 1$ for $a$ at the equilibrium, and since the difference between the equilibrium price and the lowest cost is 42 for seller 1, the value of $\tilde{\lambda}_1$ for seller 1 with $\beta_i = 400$ is approximately 9.52 seconds in early periods, which is the expected length of time into a new period when seller 1 places his first ask. Since seller 1 and buyer 1 have symmetric costs and values, buyer 1 also has an expected wait time of 9.52 seconds under these assumptions. When $\beta_i = 250$, the expected time until the first ask by seller 1 or the first bid by buyer 1 is approximately 5.95 seconds if $p(a) = 1$ at the equilibrium price.

The HBL model specifies the length of time that each seller will wait before he places an ask, and it also specifies the length of time that each buyer will wait before she places her bid. Once a new ask or bid arrives, each seller recalculates his new expected surplus maximizing ask, as well as the length of time that he will wait. Similarly, each buyer calculates a new bid and wait time. The realized ask or bid is the minimum over all sellers and buyers of these wait times. In order to compare the expected or realized length of time between messages from the HBL model to the observed time between messages from experiments, we consider the minimum wait time among instances of the HBL model. The minimum of $n$ independent exponential random variables with parameters $\{\tilde{\lambda}_i\}_{i\in\mathcal{I}}$ is an exponential random variable with parameter $\tilde{\lambda} \equiv \left(\sum_i \tilde{\lambda}_i^{-1}\right)^{-1}$. For the costs and values in table 1, at the beginning of each period, if beliefs are focused at the equilibrium for both sides of the market and $\beta_i = 400$, $\tilde{\lambda} = \frac{1}{1.08}$ so the mean time until the first ask or bid is approximately 0.926 second. In the five baseline experiments, periods $6-15$, the mean time until the first bid was 0.921 second.

A more detailed comparison between the timing decisions of human subjects and the timing actions for the HBL model could be carried out by examining what each instance of the HBL model would do at each point in time during an experiment, where the data (asks, bids, acceptances, and the timing of these actions) is fed to the instances of the model from the data generated by human subjects in experiments. The brief comparison above follows this method for the beginning of each trading period, under the assumption that beliefs are focused on the equilibrium value in periods $6-10$. In this limited case, and under this assumption about belief functions, the mean time until the first ask or bid from an instance of the HBL model is similar to the mean time until the first ask or bid by a human subject.

**Design matrix**

Table 3 summarizes the experiment design. Each row indicates a specific treatment (that is, the number of groups of automated sellers and buyers, including their pace, and the number of groups of human sellers and buyers), and the number of experiment sessions or simulations for the specified treatment. As described in Section 2.1, there are two identical sets of seller parameters and two identical sets of buyer parameters. The numbers in columns two through four indicate the number of buyer sets of each type and the numbers in columns five through seven indicate the number of seller sets of each type. Since there are six buyers and six sellers in each market, with two identical sets of buyers and two identical sets of sellers, the numbers in columns two through four sum to two in each row, as do the entries in columns five through seven. Rows in table 3 represent the same data as figure 5 : numbers of sessions in each row correspond to the numbers of sessions in the figure from top left to bottom right.

| Number of sessions | Human buyers | Slow automated buyers | Fast automated buyers | Human sellers | Slow automated sellers | Fast automated sellers |
|---|---|---|---|---|---|---|
| 2 | 0 | 2 | 0 | 2 | 0 | 0 |
| 2 | 0 | 0 | 2 | 2 | 0 | 0 |
| 5 | 2 | 0 | 0 | 2 | 0 | 0 |
| 10 | 1 | 1 | 0 | 1 | 1 | 0 |
| 4 | 1 | 0 | 1 | 1 | 0 | 1 |
| 10 | 0 | 2 | 0 | 0 | 2 | 0 |
| 10 | 0 | 2 | 0 | 0 | 0 | 2 |
| 10 | 0 | 0 | 2 | 0 | 2 | 0 |
| 3 | 2 | 0 | 0 | 0 | 2 | 0 |
| 2 | 2 | 0 | 0 | 0 | 0 | 2 |

Table 3: Experiment design matrix with number of sessions.

The rank of the design matrix is four, so in the regressions of Section 4.5, estimates of the contribution of each of the types to the relative performance of sellers and buyers are done with pairwise comparisons first between human subjects and slow automated sellers and buyers, and then between human subjects and fast automated sellers and buyers.

# 4 Analysis

As in previous analyses of double auction simulations and experiments, approximate price convergence and convergence of allocation to efficient outcomes both occur in the experiments, hybrid experiments, and simulations. In view of this convergence, it is surprising though that there are frequently large differences between the ratio of sellers' realized surplus earnings (as a fraction of their equilibrium surplus) to the buyers' realized surplus earnings (again, as a fraction of their equilibrium surplus).[7] Moreover, the difference between an individual seller's earnings and his equilibrium earnings is highly correlated with the difference between the earnings of other sellers and their equilibrium earnings. This correlation justifies a statistical model in which the relative performance sellers and buyers is modelled as a function of the seller and buyer types present in each experiment session. The main result of the paper is established through this performance analysis by regressing the relative performance of sellers to buyers in each session on the composition of types present in the session, where these types are human human buyers, slow automated buyers ($\beta_i = 400$), fast automated buyers ($\beta_i = 250$), human sellers, slow automated sellers ($\beta_i = 400$), and fast automated sellers ($\beta_i = 250$). This analysis establishes several facts: (1) performance of automated sellers and automated buyers is similar, (2) slow automated buyers outperform human buyers, (3) human buyers outperform human sellers, and (4) performance of fast automated buyers is similar to human buyers. The last of these observations is fortuitous, because it allows us to draw one other important conclusion: since performance of fast automated buyers is comparable to performance of human buyers, and since slow automated buyers outperform human buyers, we can conclude that at least locally in the vicinity of the pace of human buyers, buyer performance is enhanced by a reduction to the pace of bids.

The remainder of this section details these arguments by first establishing, in Section 4.1, convergence of prices and then, in Section 4.2, convergence to efficient allocations. Following the sections on convergence, an analysis is developed in Section 4.3 that establishes that relative surplus differences between sellers and buyers can be substantially larger than effi-

[7] In particular, if the price standard deviation is zero and the market is efficient, then the surplus split between sellers and buyers would be the equilibrium split if the equilibrium price is unique.

ciency losses. Section 4.4 demonstrates the strong relationship between an individual seller's performance and the performance of all other sellers. A similar regression is performed to establish that the same result holds for buyers. The main result of the paper is established in Section 4.5, where the relative performance of sellers to buyers is examined as a function of the types present in an individual experiment session.

### 4.1 Price convergence

Let $s_n$ be the standard deviation of price in period $n$. Price convergence is examined by estimating a model of the form $s(n) = an^{-b}$. This can be expressed as a linear model of the form $\ln s(n) = \ln a - b \ln n + \epsilon_n$. For this model, the estimate of $a$ represents the initial level of variability of transaction price, and the estimate of $b$ represents the rate at which prices converge.

| Treatment | $\hat{a}$ | $\hat{b}$ | $R^2$ | $F$ | $\hat{s}(15)$ |
|---|---|---|---|---|---|
| Simulation (original) | 4.083 | 0.086 | 0.030 | 4.546 | 3.23 |
| Simulation (modified) | 2.195 | 0.181 | 0.092 | 14.94 | 1.34 |
| Balanced (slow) | 4.215 | 0.508 | 0.420 | 107.3 | 1.06 |
| Balanced (fast) | 4.995 | 0.525 | 0.523 | 63.7 | 1.21 |
| Unbalanced (automated buyers) | 5.357 | 0.489 | 0.250 | 19.29 | 1.42 |
| Unbalanced (automated sellers) | 4.670 | 0.581 | 0.420 | 52.78 | 0.97 |
| Baseline (all human) | 13.019 | 0.928 | 0.638 | 128.5 | 1.05 |

Table 4: Price convergence in simulations, hybrid experiments, and experiments.

Results of regressions with this model, pooled across simulations, hybrid experiments, and experiments, are summarized in table 4. There are several patterns that are readily apparent from the regression results. First, both the original HBL model and the modified HBL model converge much faster in simulations than human subjects do in experiments. Moreover, there is much less reduction in variability across periods in these simulations than there is in experiments. Finally, each of the hybrid experiments has both estimates $\hat{a}$ and $\hat{b}$ that are between those of the experiments and those with human subjects. Of primary importance though is the fact that prices are very stable after several periods in simulations, in hybrid experiments, and in experiments, as measured by $\hat{s}(15)$, the estimate of price variability in the final period.

## 4.2 Efficiency

Efficiency can also be modelled with exponential convergence. If the efficiency level in period $n$ is $e(n)$, and if the rate at which foregone surplus reduces is exponential in the number of periods that have elapsed, then $1 - e(n) = an^{-b}$. For this model, significant convergence trends occur in both the hybrid experiments and the experiments. Simulations do not demonstrate a trend of convergence to efficiency, yet the initial efficiency level, which is approximately 99.4%, is similar to the figure that is attained in hybrid experiments and experiments by period 15, $\hat{e}(15)$.

| Treatment | $\hat{a}$ | $\hat{b}$ | $R^2$ | $F$ | $\hat{e}(15)$ |
|---|---|---|---|---|---|
| Simulation (original) | 0.0059 | -0.0297 | 0.0021 | 0.317 | 0.9936 |
| Simulation (modified) | 0.0057 | 0.0201 | 0.0005 | 0.074 | 0.9946 |
| Balanced (slow) | 0.0377 | 0.7575 | 0.2594 | 41.3 | 0.9952 |
| Balanced (fast) | 0.0847 | 1.2624 | 0.5499 | 70.8 | 0.9972 |
| Unbalanced (automated buyers) | 0.0165 | 0.6325 | 0.1921 | 13.7 | 0.9970 |
| Unbalanced (automated sellers) | 0.0877 | 1.0552 | 0.3569 | 40.5 | 0.9950 |
| Baseline (all human) | 0.0686 | 0.7806 | 0.3251 | 35.2 | 0.9917 |

Table 5: Efficiency in simulations, hybrid experiments, and experiments.

## 4.3 Surplus split

### Theoretical analysis

The mean price in experiments and hybrid experiments typically differs from the equilibrium price, and this difference leads to large differences in the relative performance of buyers and sellers. Nevertheless, the surplus loss in these markets is typically very small. We examine two benchmark cases, which are depicted in figure 6, to demonstrate that the measure of convergence typically used in market experiments does not imply an equal division of surplus. Analyses for these two cases are both based on the assumptions that supply and demand are both linear, and the slope of demand is the negative of the slope of supply.

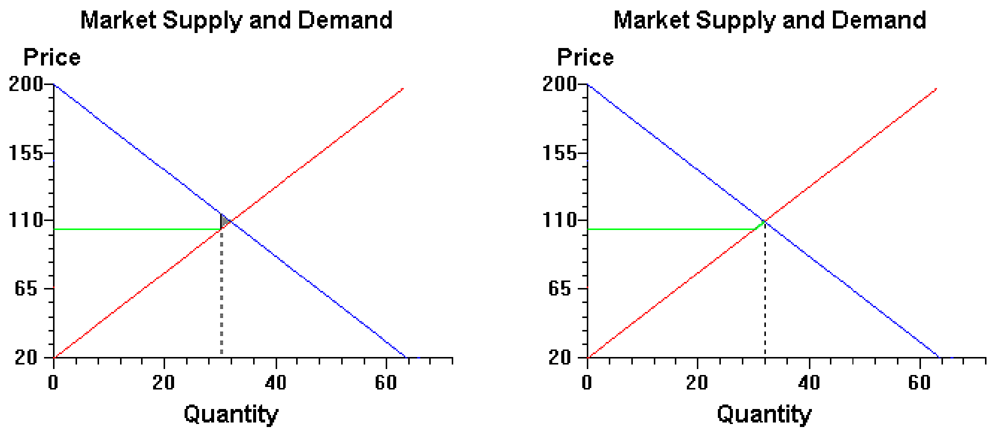


Figure 6: Surplus losses and earnings differences at non-equilibrium prices.

In the first benchmark case, shown on the left side of figure 6, all trades occur at a single price (price variability is zero), and efficiency is high, yet the split of surplus is highly uneven. Specifically, when all trades occur at a single price $p(\alpha) = \alpha\, p_1 + (1-\alpha)\, p_e$, where $S(p_1) = 0$, then the fraction of surplus lost is only $\alpha^2$, but the ratio of surplus earned by buyers to surplus earned by sellers is $R_1(\alpha) = 1 + 4\,\frac{\alpha}{1-\alpha}$. In the other benchmark case, shown on the right side of figure 6, there is complete efficiency, yet even with low price variability, the split of surplus is also uneven. In this case, all trades occur at a single price $p(\alpha)$ until the supply is exhausted at that price, and subsequent trades – for prices $p \in (p(\alpha), p_e)$ – occur at prices $p$ along the supply schedule. In this case, we find that the price standard deviation, normalized by the range from the highest value to the lowest cost, is approximately 1/7 the value of $\alpha$. Yet the ratio of buyer to seller surplus in this case is $R_2(\alpha) = \frac{(2-(1-\alpha)^2)}{(1-\alpha)^2}$, which is even greater than in the first benchmark case.

For example, if the fractional difference of the price from equilibrium is $\alpha = 0.07$, then the surplus loss is only 0.49% in the first benchmark case (shown as the shaded triangle), but the ratio of buyers' surplus to sellers' surplus is 1.301. In the second benchmark case, with $\alpha = 0.07$ the normalized price variability is 0.52% and the ratio of buyers' to sellers' surplus is 1.312. This observation has important consequences for the interpretation of the results of market experiments. Since large differences in earnings are typical, and are consistent with both price stability and high efficiency, strategic interactions between sellers and buyers play a more important role in competitive markets than we have realized.

Table 6 shows the theoretical surplus split as a function of price differences from equilibrium for the ICV2 design used in the experiments, hybrid experiments, and simulations. The surplus split and efficiency calculations in the table are made under the assumptions that all trades take place at a single price, and all trades occur that increase or at least leave constant the surplus of both seller and buyer, i.e., under the assumptions of the first benchmark case above.

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

Table 6: Surplus and efficiency as a function of price deviation from equilibrium.

**Realized split**

Figure 7 shows the natural logarithm of the ratio of buyers' to sellers' surplus as a function of the difference between the average transaction price and the equilibrium price for three types of simulations. The triangles on the upper left show the observations from 10 simulations with fast sellers ($\beta_i = 250$ for $i \in \mathcal{I}_{\mathcal{S}}$) and slow buyers ($\beta_i = 400$ for $i \in \mathcal{I}_{\mathcal{B}}$). The 10 squares in the center of the figure are from simulations with slow buyers and slow sellers. The remaining 10 observations are from simulations with fast buyers and slow sellers.

While there is some variability of relative performance in these simulations within each

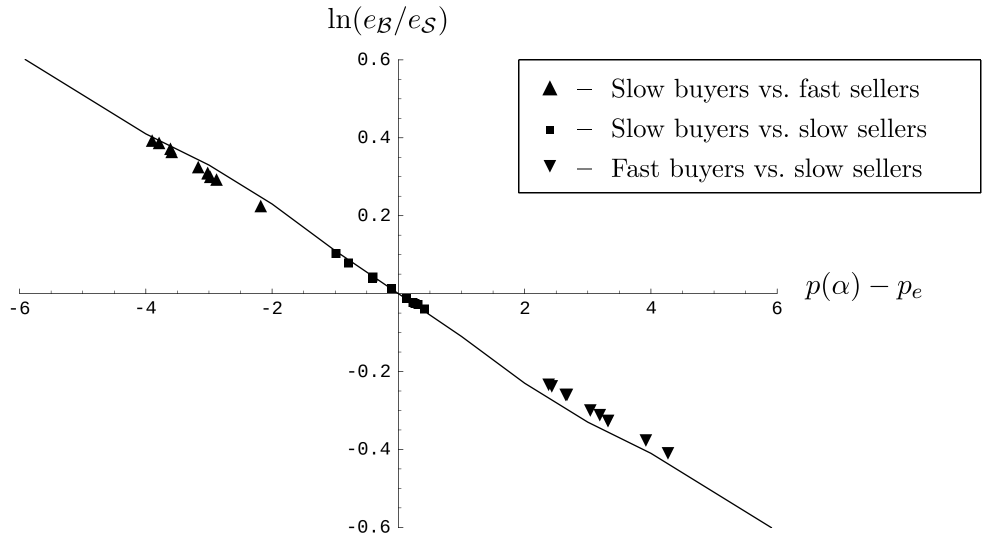


Figure 7: Surplus ratio and price difference from equilibrium in simulations.

group of simulations, a difference between the pace of asks and bids has a pronounced effect on the relative performance of sellers and buyers. This pattern is also present in hybrid experiments, as shown in figure 8. As in figure 7, there are three types of data shown in figure 8. Hybrid experiments with human sellers and automated buyers are shown with upward triangles, each of which is in the quadrant with a positive value for $\ln(S_{\mathcal{B}}/S_{\mathcal{S}})$ and a negative value of $p(\alpha) - p_e$. Hybrid experiments with human sellers and automated buyers are shown with downward triangles, each of which is in the quadrant with a negative value for $\ln(S_{\mathcal{B}}/S_{\mathcal{S}})$ and a positive value of $p(\alpha) - p_e$. Balanced hybrid experiments are shown as squares: of these 13 are in the upper left and one is in the lower right quadrant.

Finally, figure 9 shows the relationship between the ratio of buyers' to sellers' surplus from five experiments. In the experiments, there is also a tendency for sellers to underperform buyers. Performance ratio statistics from experiments and hybrid experiments do not have the degree of regularity observed in the simulations of figure 7, yet there is a consistent pattern between each of the six types employed in the experiment design (human buyers and sellers, slow automated buyers and sellers, and fast automated buyers and sellers) and the division of surplus between buyers and sellers.

Automated sellers and automated buyers have equal bargaining capability since they are

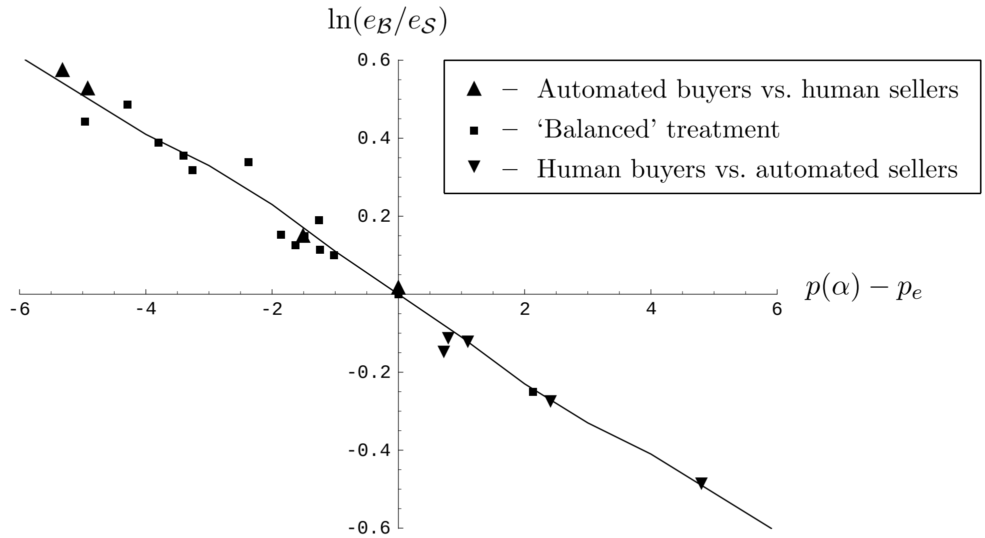


Figure 8: Surplus ratio and price difference in hybrid experiments.

specified symmetrically, but as figure 7 shows, slow automated sellers have an advantage over fast automated buyers and vice versa. In addition, both automated sellers and automated buyers outperform humans. The third pattern is that human buyers outperform human sellers. In order to demonstrate these patterns of relative performance, a statistical model is developed that regresses the performance of sellers versus buyers on the types present in each experiment. Coefficients for each type are then interpreted as contributions to the performance of sellers relative to buyers in an experiment. In order to carry out statistical analysis that treats the performance of sellers relative to buyers as a function of the types present in each experiment, 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 modelled as dependent on the profile of types (e.g., slow automated sellers, fast automated sellers, and human sellers) present in an experiment. The next subsection demonstrates the high positive correlation of performance among sellers and among buyers.

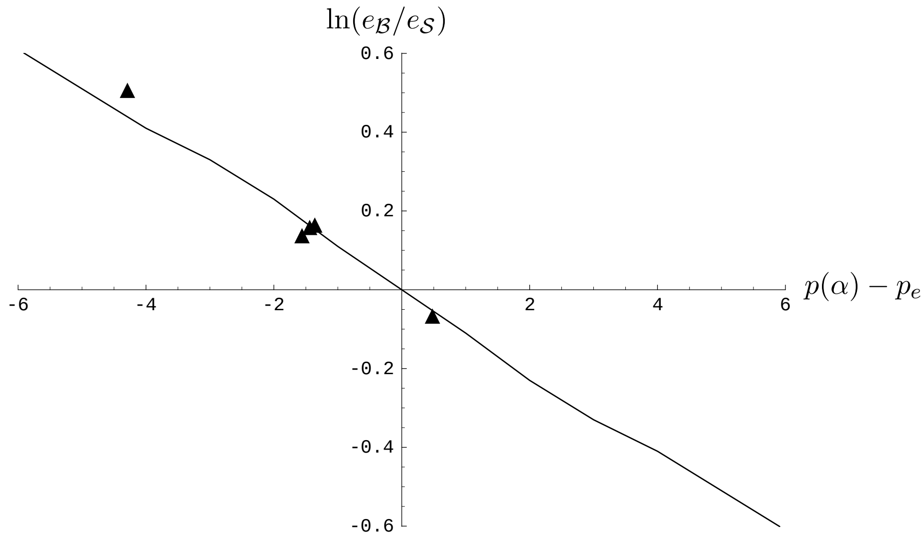


Figure 9: Surplus ratio and price differences in experiments.

## 4.4 Correlation

The performance of sellers and of buyers, when plotted across trading periods, exhibits a high degree of correlation. Figure 10 shows the performance of three automated buyers and three automated sellers across 15 periods of the hybrid experiment ICV2a-DA-020330b1. In the first period, and from periods 5 through 15, each of the three automated buyers had earnings that exceeded their equilibrium earnings, and each of the three automated sellers had earnings below their equilibrium earnings. This situation is typical as the regression model below demonstrates.

Let $e_i$ be the ratio of the surplus earned by seller $i$ to the equilibrium surplus for seller $i$, and let $e_{-i}$ be the ratio of the surplus earned by all sellers other than $i$ to the equilibrium surplus for these other sellers. If the surplus in equilibrium is the same for each buyer and each seller, the sum of the normalized performance measures $e_i - 1$ is equal to the market efficiency minus one. For the case of the experiment parameters in table 1, the sum of the normalized performance measures is approximately zero provided the earned surplus for each seller and buyer is positive. Consequently, a reasonable null hypothesis is that the values of the performance measures are independent and identically distributed on the set $\{e : \sum_{i=1}^{12} e_i - 1 = 0\}$. In this case, if the coefficients $\{e_j : j \in \mathcal{I}_{\mathcal{B}} \setminus \{i\}\}$ sum

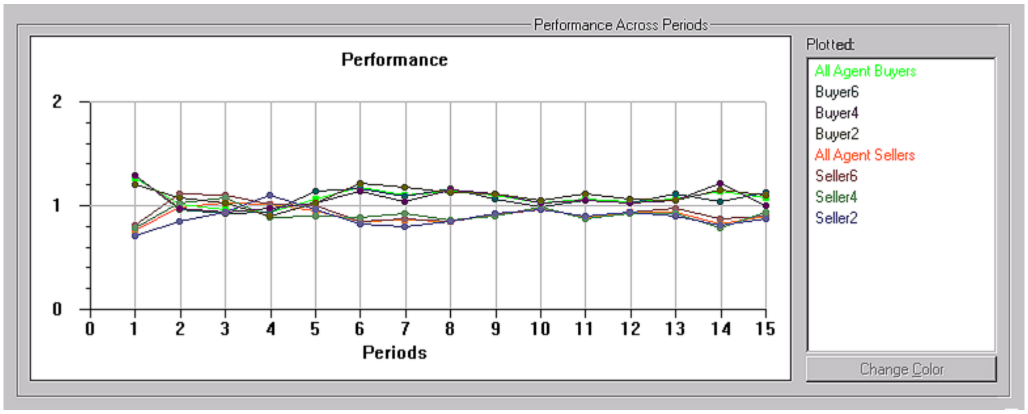


Figure 10: Performance of sellers and buyers across 15 experiment periods.

to $e_{-i}$ then $E[e_i] = -\frac{1}{7}e_{-i}$. To test this hypothesis, we examine the regression model $e_i - 1 = \alpha_0 + \alpha_1(e_{-i} - 1) + \epsilon_i$. Table 7 shows the results of this regression model for each of the design treatments.

| Treatment | Type | $\alpha_0$ | $\alpha_1$ | $R^2$ | $F$ |
|---|---|---|---|---|---|
| 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: Performance correlation among sellers and among buyers.

Based on the results of these regressions, with all estimated slope coefficients closer to one than to $-\frac{1}{7}$, it is appropriate to model the relative performance of sellers to that of buyers as a function of the distribution of types present in an individual experiment session. This relative performance model is described and analyzed in the next subsection.

### 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 experiment session. Relative performance $r$ in this context is defined as

$$r \equiv \frac{e_{\mathcal{S}}}{e_{\mathcal{B}}}$$

where $e_{\mathcal{S}}$ is the surplus attained by sellers divided by the equilibrium surplus of sellers, and $e_{\mathcal{B}}$ is defined similarly.

Assume that relative performance is a function $r(X)$ from $\mathbb{R}_+$ to $\mathbb{R}_+$ where $X$ is a random variable that is a measure of the relative pace of asks versus bids. In order to have the function $r(X)$ 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 $X = e^Y$ and $Y \sim N(\mu, \sigma)$. Let $B = \{b_1, b_2\}$ and $S = \{s_1, s_2\}$ denote the two buyer and two seller groups, and 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 $X = e^Y$ that is a measure of relative pace 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 $E[X] = 1$, in which case the relative performance of sellers and buyers is equal.

The regression model is

$$\ln r(X_k) = \alpha_0 + \alpha_h\, s_h + \alpha_a\, s_a - \beta_h\, b_h - \beta_a\, b_a + \epsilon_k \tag{5}$$

where $X_k$ is the relative pace of sellers to buyers in experiment session $k$. The variables $s_h$ and $s_a$ indicate the number of human seller groups and the number of slow automated seller groups present in a session, where each of these variables takes the value 0, 1, or 2 and $s_h + s_a \leq 2$. (The number of fast automated seller groups in a session is $s_a^f = 2 - s_h - s_a$.) Similarly, $b_h$ and $b_a$ are the number of human buyer and slow automated buyer groups. The negative sign on the buyer terms have been chosen so that the strength of seller and buyer coefficients can be compared directly, since the negative sign adjusts for the fact that strong buyers have the opposite impact from strong sellers on the ratio of seller performance to buyer performance.

Estimated coefficients from the regression model are shown in table 8. The $R^2$ statistic is 0.6979, so that the model is capable of explaining a large fraction of the variability in relative performance. The $F$ statistic is $F = 30.6$, which has significance at any reasonable level (e.g., 0.001).

These coefficients can be interpreted in both relative and absolute terms. In relative terms, the presence of human sellers ($s_h > 0$) has a detrimental effect on the ratio of seller surplus to buyer surplus, and this effect is significant. Human buyers have no effect on the relative performance statistic (since $\beta_h$ is not significantly different from zero); relatively, they are stronger than human sellers. Slow automated buyers and slow automated sellers have comparable effects on performance. Finally, both human sellers and human buyers are weaker than slow automated sellers or buyers. In the case of automated sellers, their presence has a significant positive impact on the relative performance of sellers. Human buyers have a neutral impact, but slow automated buyers contribute negatively to the relative performance of sellers. Consequently, slow automated buyers are stronger than human buyers. In absolute terms, the expected change in relative surplus of sellers to buyers is $e^{-0.157}$ (approximately a 15.5% reduction) when one group of human buyers is replaced by one group of slow automated buyers.

| | $\alpha_0$ | $\alpha_h$ | $\alpha_a$ | $\beta_h$ | $\beta_a$ |
|---|---|---|---|---|---|
| Estimate | -0.003 | -0.117 | 0.142 | -0.001 | 0.157 |
| SE | 0.066 | 0.036 | 0.028 | 0.034 | 0.028 |

Table 8: Estimates for humans versus slow automated sellers and buyers.

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

$$\ln r(X_k) = \alpha_0 + \alpha_h\, s_h + \alpha_a\, s_a^f - \beta_h\, b_h - \beta_a\, b_a^f + \epsilon_k \tag{6}$$

where $s_a^f$ and $b_a^f$ are the numbers of fast automated seller and fast automated buyer groups. Coefficients for the regression formulation in equation (6) are shown in table 9.

| | $\alpha_0$ | $\alpha_h$ | $\alpha_a^f$ | $\beta_h$ | $\beta_a^f$ |
|---|---|---|---|---|---|
| Estimate | -0.028 | -0.258 | -0.142 | -0.158 | -0.157 |
| SE | 0.041 | 0.032 | 0.028 | 0.032 | 0.028 |

Table 9: Estimates for humans versus fast automated sellers and buyers.

The most important additional comparison in table 9 is between human buyers and fast automated buyers. These coefficients are similar, and both have small standard errors. Consequently, the timing specification for fast automated buyers yields performance that is similar to human performance. Moreover, from the regression in equation (5) we know that slow automated buyers outperform human buyers. Taken together these observations imply that the performance of automated buyers improves as their pace decreases, when the pace is specified so that automated buyers and human buyers have comparable performance.

These regressions demonstrate that, aggregated across experiment 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, leads to 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.

## 5 Conclusions

The notion of price taking behavior is strongly associated with competitive equilibrium, yet this paper demonstrates that substantial deviations from equilibrium earnings are consistent with the stable prices and approximate efficiency that are typical in double auction experiments, and the deviations from equilibrium earnings are driven to a large extent by the strategic impact of 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 prominantly in the interactions of market participants after low levels of price variability are attained. Once prices have stabilized, 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. Thus, there is a substantial role for strategic behavior even under competitive outcomes, and this paper, through use of automated strategies, demonstrates that the pace of sellers' asks and buyers' bids is a key element of such strategic behavior in temporally unstructured bargaining.

## References

[1] Easley, David, and John Ledyard (1993). "Theories of Price Formation and Exchange in Double Oral Auctions," in *The Double Auction Market: Institutions, Theories, and Evidence* (Daniel Friedman and John Rust, Eds.), pp. 63 - 97. Addison-Wesley.

[2] Friedman, Daniel (1991). "A Simple Testable Model of Price Formation in the Double Auction Market," *Journal of Economic Behavior and Organization* **15**, 47 - 70.

[3] Gjerstad, Steven and John Dickhaut (1998). "Price Formation in Double Auctions," *Games and Economic Behavior* **22**, 1-29.

[4] Gode, Dhananjay, and S. Sunder (1993). "Allocative Efficiency of Markets with Zero Intelligence Traders: Market as a Partial Substitute for Individual Rationality," *Journal of Political Economy*, **101**, 119-37.

[5] Wilson, Robert B. (1987). "On Equilibria of Bid-Ask Markets," in *Arrow and the Ascent of Modern Economic Theory* (G.W. Feiwel, Ed.), pp. 375 - 414. New York University Press.

**Appendix A**

**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 $\tilde{S}(b) = (v-b)\,q(b)^{\frac{1}{n}}$. 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** 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 $\tilde{S}'(b) = (v-b)\,\frac{1}{n}q(b)^{\frac{1}{n}-1}\,q'(b) - q(b)^{\frac{1}{n}}$. 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

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

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

Next consider assertion (2). Let $q_k^* = q(b_k^*)^{\frac{1}{n}}$ and let $\tilde{q}_k^* = q\left(\tilde{b}_k^*\right)^{\frac{1}{n}}$. Then $(\tilde{q}_k^*)^n = q\left(\tilde{b}_k^*\right)$. Suppose that $(v-\tilde{b}_k^*)\,\tilde{q}_k^* > (v - \tilde{b}_{k-1}^*)\,\tilde{q}_{k-1}^*$. Since $\tilde{b}_{k-1}^*$ maximizes $(v-b)\,\tilde{q}(b)$ on $[d_{k-1}, d_k)$, $(v-\tilde{b}_k^*)\,\tilde{q}_k^* > (v-b_{k-1}^*)\,q_{k-1}^*$. Also, $q_k^* > \tilde{q}_k^*$ so $(v-\tilde{b}_k^*)\,q_k^* - (v-b_{k-1}^*)\,q_{k-1}^* > 0$. Then

$$\begin{aligned}(v-b_{k-1}^*)(q_k^*-q_{k-1}^*) &> (\tilde{b}_k^*-b_{k-1}^*)\,q_k^* \\ (v-b_{k-1}^*)(q_k^*-q_{k-1}^*)\sum_{i=0}^{n-1}(q_k^*)^{n-1-i}(q_{k-1}^*)^i &> (\tilde{b}_k^*-b_{k-1}^*)\,q_k^*\sum_{i=0}^{n-1}(q_k^*)^{n-1-i}(q_{k-1}^*)^i \\ (v-b_{k-1}^*)((q_k^*)^n-(q_{k-1}^*)^n) &> (\tilde{b}_k^*-b_{k-1}^*)\,q_k^*\sum_{i=0}^{n-1}(q_k^*)^{n-1-i}(q_{k-1}^*)^i \\ (v-b_k^*)(q_k^*)^n-(v-b_{k-1}^*)(q_{k-1}^*)^n &> (b_{k-1}^*-b_k^*)(q_k^*)^n \\ &\quad +(\tilde{b}_k^*-b_{k-1}^*)\,q_k^*\sum_{i=0}^{n-1}(q_k^*)^{n-1-i}(q_{k-1}^*)^i \\ (v-b_k^*)(q_k^*)^n-(v-b_{k-1}^*)\,(q_{k-1}^*)^n &> (\tilde{b}_k^*-b_k^*)(q_k^*)^n + \\ &\quad (\tilde{b}_k^*-b_{k-1}^*)\sum_{i=1}^{n-1}(q_k^*)^{n-i}(q_{k-1}^*)^i. \qquad \text{(A.1)}\end{aligned}$$

If the expression on the right side of equation (A.1) is positive, then

$$(v - b_k^*)(q_k^*)^n > (v - b_{k-1}^*)(q_{k-1}^*)^n$$

or

$$(v - b_k^*)q(b_k^*) > (v - b_{k-1}^*)\, q(b_{k-1}^*).$$

So claim (2) follows if

$$(\tilde{b}_k^* - b_{k-1}^*) \sum_{i=1}^{n-1} (q_k^*)^{n-i} (q_{k-1}^*)^i > (b_k^* - \tilde{b}_k^*)(q_k^*)^n. \tag{A.2}$$

Choose $\alpha \in (0,\, 1)$ so that $\tilde{b}_k^* = \alpha\, b_k^* + (1-\alpha) b_{k-1}^*$. Then equation (A.2) is equivalent to

$$\alpha \sum_{i=1}^{n-1} (q_k^*)^{n-i} (q_{k-1}^*)^i > (1-\alpha)\,(q_k^*)^n. \tag{A.3}$$

Let $\beta = \frac{q_{k-1}^*}{q_k^*}$. Then equation (A.3) is equivalent to

$$\alpha > \frac{(1-\beta)\beta^n}{1-\beta^{n+1}}.$$

For all $n \geq 2$ and for all $\beta \in (0,\, 1)$, if $\alpha \geq \frac{1}{n+1}$ then $\alpha > \frac{(1-\beta)\beta^n}{1-\beta^{n+1}}$. Consequently, the claim follows if it can be shown that $\tilde{b}_k^* > \frac{1}{n+1}\, b_{k-1}^* + \frac{n}{n+1}\, b_k^*$.

The first-order condition for $\tilde{b}_k^*$ is $\frac{1}{n}\,(v - \tilde{b}_k^*)\, q^{'}(\tilde{b}_k^*) = q(\tilde{b}_k^*)$ or $v - \tilde{b}_k^* = \frac{n\, q(\tilde{b}_k^*)}{q^{'}(\tilde{b}_k^*)}$. The value of $v - \tilde{b}_k^*$ is smallest for $v = d_{k+1}$ so the value of $q^{'}(b)$ is largest at $\tilde{b}_k^*$ when $v = b_{k+1}$. Therefore $\tilde{b}_k^*$ takes on its smallest value in $[d_k,\, d_{k+1})$ when $v = d_{k+1}$ so it is sufficient to prove the assertion for this case. In this case, a direct calculation shows that $\tilde{b}_k^*$ and $b_k^*$ lie in $(0.607625\, d_k + 0.392375\, d_{k+1},\, d_{k+1})$. This calculation involves several steps. First, there is at most one root of $S_i^{'}(b)$ in $[d_k,\, d_{k+1})$, and there is one root if and only if $q(d_k) \leq \frac{11}{27}\, q(d_{k+1})$. In this case, the interior local maximum $S_i(b_k^*)$ at $b_k^*$ is greater than $S_i(d_k)$ (which may be a local maximum on $[d_k,\, d_{k+1})$) if $q(d_k) \leq 0.345615\, q(d_{k+1})$. For values of $q(d_k)$ and $q(d_{k+1})$ that satisfy this inequality, $b_k^* > 0.607625\, d_k + 0.392375\, d_{k+1}$.

Consequently, $\tilde{b}_k^* > \frac{1}{3}\, d_k + \frac{1}{3}\, d_{k+1}$. Since $d_k > b_{k-1}^*$, it follows that $\tilde{b}_k^* > \frac{1}{3}\, b_{k-1}^* + \frac{2}{3}\, d_{k+1}$. Finally, since $d_{k+1} > b_k^*$, the claim that $b_k^* > \frac{1}{3}\, b_{k-1}^* + \frac{2}{3}\, b_k^*$ is verified. ■