Reprinted from

JOURNAL OF Economic Behavior & Organization

Journal of Economic Behavior & Organization
Vol. 31 (1996) 273–291

Bayesian equilibrium in double auctions populated by biased heuristic traders[1]

Karim Jamal[a,*], Shyam Sunder[b]

[a] Department of Accounting and MIS, Faculty of Business, University of Alberta, Edmonton T6G 2R6, Canada
[b] Graduate School of Industrial Administration, Carnegie Mellon University, Pittsburgh, PA 15213-3890, USA

ELSEVIER

ELSEVIER

Journal of Economic Behavior & Organization
Vol. 31 (1996) 273-291

JOURNAL OF
Economic Behavior
& Organization

Bayesian equilibrium in double auctions populated by biased heuristic traders[1]

Karim Jamal[a,*], Shyam Sunder[b]

[a] *Department of Accounting and MIS, Faculty of Business, University of Alberta, Edmonton T6G 2R6, Canada*
[b] *Graduate School of Industrial Administration, Carnegie Mellon University, Pittsburgh, PA 15213-3890, USA*

Received 1 December 1995; revised 27 February 1996

Abstract

We use computer simulation to examine three asset markets with imperfect information. In processing imperfect information, traders in the three markets are bayesian, empirical bayesian, and heuristic (representativeness and anchor-and-adjust) respectively. All three converge to the same bayesian equilibrium – although the latter two converge more slowly – without profit maximization, natural selection, arbitrage, or mutual cancellation of random actions. The results support Becker (1962) and Simon (1973) in that the rationality of the market emerges as a consequence of the market structure, and not from the rationality of individuals.

JEL classification: A12; C11; D44; D81

Keywords: Aggregate market rationality; Bayesian equilibrium; Double auction; Representativeness heuristic; Zero-intelligence trader

1. Introduction

In economic theory competitive or strategic equilibria are derived from optimal choices made by agents. Direct observations of individual intuitive behavior indicate that it systematically and significantly deviates from optimal choice. When working by

* Corresponding author. Tel.: 403 492 5829; e-mail: kjamal@gpu.srv.ualberta.ca.

[1]We gratefully acknowledge the comments of the referee, Stephen Kachelmeier, Charles R. Plott and other participants on an earlier draft of the paper presented at the meeting of Economic Science Association at Tucson, Arizona, in November 1994. We alone are responsible for the paper. Programs to see these simulations live on your own computer can be downloaded through ftp. See the Appendix for details.

intuition, humans use heuristics that often fail to use all the information that is relevant and available. In this paper we ask: Do competitive markets belong to a class of social institutions whose structure causes them to yield outcomes that are substantially insulated from the imperfect judgments and actions of their participants? We address this question by conducting a computational experiment of double auction markets populated by one of three different types of computer traders (algorithms written as computer code): bayesian, empirical bayesian, and biased heuristic traders. The biased heuristic traders employ two well-documented biased heuristics in cognitive psychology called the "representativeness" and the "anchor-and-adjust" heuristics. We find that bayesian equilibrium prediction is better in organizing the data from all three sets of markets than the representativeness heuristic prediction. When all individuals follow the representativeness heuristic, the aggregate market behavior conforms to the prediction of the heuristic initially, but it gradually settles down to the prediction of Bayes' theorem.

Our study follows the methodological tradition initiated by Richard M. Cyert, his colleagues and students some forty years ago to study various aspects of organizational behavior using computer simulation techniques. Cohen and Cyert (1965) classified such simulations into four categories: (1) descriptive studies that are meant to depict the actual behavior of existing organizations, e.g. Cyert et al., 1963; (2) illustrative studies to discover the consequences of specific behavior in organizations, e.g. Cohen et al. (1963); (3) normative simulations to help design better organizations, e.g. Bonini (1963); and (4) man–machine simulations to train people in how to behave in organizations, e.g. Cohen et al. (1960). We have shifted our focus from organizations to markets, and we use computer simulations to discover and compare the market-level consequences of specific types of individual behavior. In their *Behavioral Theory of the Firm*, Cyert and March (1963) shifted the focus of inquiry from outcomes toward processes. Our paper is an attempt to synthesize the outcome orientation of economic theory with the process orientation of theories about individual behavior. Such a synthesis yields a better understanding of both.

2. Markets as task environments

Studies of human problem solving (e.g. Newell and Simon, 1972) indicate that problem solving performance depends on two factors – the information processing demands of the task at hand, and the computational abilities of the agent. Tasks can be ranked along a continuum from well-structured to ill-structured (Simon, 1973). An algorithm can perform a well-structured task with little human intervention. At the other extreme, an ill-structured task requires a human agent to choose a cognitive representation (or frame) from a set of competing representations, and to choose an interpretation of data that has a major effect on the quality of the solution achieved. Under this perspective, the effort required to solve the problem, and the quality of the solution achieved, depend on the structure of the task.

2.1. Ill-structured task environments

Behavioral scientists have generally assumed that market traders are involved in a complex ill-structured task where mental calculation, cognitive effort, and individual utility maximization are necessary conditions for attaining competitive equilibria. The concern about the implications of psychological studies that indicate that individuals often use simple heuristics that generate biased behavior has given rise to two debates. The methodological debate challenges the experimental procedures used by psychologists. The theoretical debate has focused on the relevance of biased individual behavior for market-level outcomes.

The methodological debate has probed the procedures used in psychological experiments. The key issues are a lack of performance-based incentives (Grether, 1980, Jamal and Sunder, 1991); the use of one-shot experiments where inexperienced subjects are asked to perform complex, unfamiliar tasks without an opportunity to get feedback and learn (Grether and Plott, 1979); and the lack of availability of decision aids that might normally be used by agents to perform the task on hand (Edwards and von Winterfeldt, 1986).

The theoretical debate has sought to identify processes that could correct biases in individual behavior. Several candidate processes have been identified in the literature (see Camerer, 1987, Duh and Sunder, 1986, 1993, Anderson and Sunder, 1995) such as: random individual biases cancel out in a market (cancellation); smart traders act as arbitrageurs and determine the marginal trades that influence prices (arbitrageurs); biased traders learn from feedback, by observing smart traders, or by real life work experience (learning); biased traders are driven out of the market by bankruptcy and only the smart traders survive (natural selection). Biased traders may also seek expert advice and thus institutional mechanisms or specialists may arise to assist these traders (see Jamal et al., 1995 for an example where auditors protect market traders from management fraud). Studies conducted to address these issues with human subjects have generally yielded mixed results (Grether, 1980, Camerer, 1987, Duh and Sunder, 1986, 1993). Human subjects who are paid on the basis of their performance and who act in a market environment exhibit the use of various biased heuristics identified in psychological studies. However, the magnitude of individual biases, and their effect on market-level outcomes is attenuated in the presence of incentives and experience (e.g. Anderson and Sunder, 1995).

2.2. *Well-structured task environments*

In Simon's terminology, Becker (1962) can be said to have proposed that several economic tasks are well-structured rather than ill-structured. For example, Becker showed that downward sloping aggregate demand functions can arise even in the absence of utility maximization. Constraints imposed by the task environment (e.g. a budget constraint) are sufficient to yield rational outcomes at an aggregate level, without requiring rational behavior from individuals. The task environment determines the outcome without conscious human intervention in this case.

The debate on individual rationality has obscured the effect of institutions and task environments on outcomes. Becker (1962) cautioned that we need to distinguish between individual and aggregate behavior, and that we should not impute irrationality of individuals to markets, or rationality of aggregate market behavior to individuals. However, a direct and causal link between individual behavior and aggregate market behavior is frequently assumed (Russell and Thaler, 1985). Social institutions such as markets, introduce structure into individual choice behavior and can therefore generate an aggregate behavior that is quite different from the behavior of disconnected individuals (Coleman, 1986). Research in experimental economics (Plott, 1982 and Smith, 1982) indicates that aggregate market prices and allocations of securities are influenced significantly by the type of market institution employed (e.g. english, dutch, or sealed-bid auction). We need to understand how institutions transform individual into aggregate behavior (Coleman, 1986).

Using this perspective, Gode and Sunder's (1993, 1995) and Bosch and Sunder's (1994) computational and human experiments have shown that the allocative efficiency of a double auction derives largely from its structure and not from the agents' motivation, knowledge or learning. In these simulations, "zero-intelligence" (ZI) computer program traders randomly submit bids and offers. ZI traders have no motivation, knowledge, memory or learning; they just behave randomly subject to a budget constraint or opportunity set. Yet the allocative efficiency of markets composed of ZI traders is close to 100%. These results suggest that the structure of the market institution (double auction), rather than individual behavior, may be the primary driver behind the efficiency of aggregate market behavior. This is significant because individual profit maximization and complex cognitive calculation are usually assumed to be key characteristics driving human behavior in markets. Experimental results that do not provide significant monetary rewards are often suspect for not providing sufficient inducement to generate the complex mental calculations required in a market environment (Smith, 1982).

Performance of ZI markets indicates that, contrary to the prevailing view in both economics and psychology, aggregate market-level rationality demands only a relatively low level of information processing of individuals in double auctions. Individuals do not need to be rational, or even follow any judgment heuristics. Random individual behavior, subject to a budget constraint, is sufficient to generate rational aggregate behavior. These results indicate that the double oral auction is a well-structured task that requires little human intervention to achieve problem solving success.

The Gode and Sunder (1993, 1995) and Bosch and Sunder (1994) results suggest that Becker's (1962) argument about aggregate rational behavior emerging from individually irrational behavior is more general than previously thought. English and Allison (1993) also demonstrate how various empirical relationships in both the psychology and economics literature can be derived from the assumption of random response to external task constraint. These results suggest a fascinating hypothesis: markets may attain a bayesian equilibrium even if the individual agents participating in the market do not process information in a manner consistent with Bayes' rule, and all the usual suspects (smart arbitrageurs, cancellation of random responses, learning, experts or natural selection) have valid alibis. In this paper we present results from simulated markets designed to test this hypothesis.

3. Method

The judgment heuristics used by human subjects cannot be controled and selected for. In psychology experiments, various judgment heuristics are activated by manipulating task materials. Results from psychological experiments indicate large variations across individuals in heuristics used and outcomes chosen by human subjects (Tversky and Kahneman, 1986). Human subjects also differ in their preferences and are vulnerable to experimental demand effects. The number and duration of human subject experiments are limited by subject availability, fatigue, and cost. These factors make it difficult to conduct long market experiments to compare the consequences of actual human behavior against the bayesian standard. We circumvent these difficulties by replacing human traders with computer programs.

Computer traders behave in a programmed manner, allowing us to explicitly simulate the use of the representativeness and anchor-and-adjust heuristics. The use of a simulation model provides us with an opportunity to causally explore the factors that lead to market outcomes. Use of computer program traders also enables us to control the length of the experiments and determine how long it takes a market to reach the bayesian equilibrium, if it is attained. Experiments with human subjects generally permit some 15–20 periods; we report results of computational experiments for each of the three types of traders of 100 markets (consisting of 100 periods each) for a total of 30,000 periods.

By using an unchanging group of computer traders programmed to use judgment heuristics that normally lead to biased (non-bayesian) behavior, we explicitly rule out explanations based on evolutionary arguments, arbitrage (few smart individuals), and cancellation of random responses. There are also no markets for information or specialists in our experiment. Traders cannot seek the advice of experts in the experiment.

3.1. Market mechanism

We used a market design similar to the designs used in experiments with human subjects by Grether (1980), Duh and Sunder (1986, 1993), Camerer (1987) and Anderson and Sunder (1995). Computer traders participated in 100-period long double auctions.

At the beginning of each period, each of the *n* computer traders is endowed with some cash and one certificate. The certificate pays a single state-contingent dividend that varies randomly across the *n* computer traders. For example, the dividend in state X (and Y) for trader 1 may be 100 (and 50) points, but the dividend for trader 2 in state X (and Y) may be 70 (and 80) points. The vector of state-contingent dividends for each trader is known privately by each trader. It serves as a constraint on the maximum values that the trader would be willing to pay for a certificate and the minimum value it would be willing to accept. (See Fig. 1 and Table 1; Section 4 gives the detailed parametric structure of these markets.)

Prior to trading, computer traders receive an imperfect public signal (G for green or B for brown) about the state (X or Y) for that period.The signal received is the same across all traders. Computer traders are then free to enter a bid to purchase a security or to make an offer to sell a security. The same buyer or other buyers can subsequently raise the bid. Similarly, any seller can offer to sell a certificate at a given price or reduce the price offered by some other trader. A transaction occurs when the bid and offer prices match or cross. When bids and offers cross, the transaction is recorded at their average. At the end of the period the actual state (X or Y) is revealed, dividends are paid, and profit for each computer trader is computed.

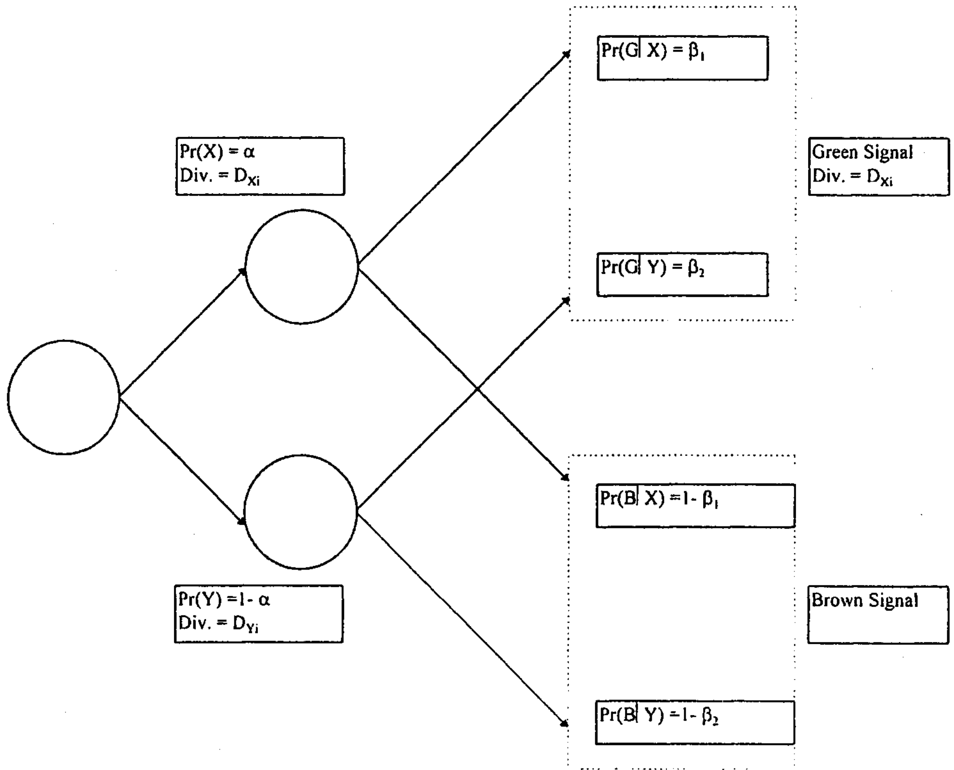

Fig. 1. Uncertainty and information structure of the market.

Table 1
Design parameter for simulations

Parameter	Value or distribution
Number of markets	100
Number of periods in each market	100
Number of traders in each market	20
Number of iterations per period	3,000
Dividend in state X (D_{Xi})	~U(0,140)
Dividend in state Y (D_{Yi})	~U(60,200)
Base rate prob. (X) (α)	~U(0.3,0.7)
Prob. (Signal G | X) (β_1)	~U(0.7,0.9)
Prob. (Signal G | Y) (β_2)	~U(0.1,0.3)
Transaction price adaptive parameter (γ)	~U(0.01,0.1)
End-of-period dividend adaptive parameter (δ)	~U(0.02,0.2)

3.2. *Trading strategies*

Three types of trading strategies were simulated, bayesian and empirical bayesian as controls, and biased heuristic as the treatment. Following Gode and Sunder (1993), all three are "zero-intelligence" (ZI) strategies in the sense that they pick bids or ask randomly from the no-loss feasible set. The following paragraphs give the general description of the strategies.

Trading strategy 1 is a bayesian strategy. After observing the imperfect public signal (G or B) about the realized state (X or Y) at the beginning of each trading period, these traders compute a posterior bayesian expected value from their dividends, prior probabilities of X and Y, and the conditional probabilities of signals G and B. If we use the notation:

Prob. (StateX) $= \alpha$,	Prob. (StateY) $= 1 - \alpha$
Prob. (Green | StateX) $= \beta_1$,	Prob. (Brown | StateX) $= 1 - \beta_1$
Prob. (Green | StateY) $= \beta_2$,	Prob.(Brown | StateY) $= 1 - \beta_2$
Dividend for *i*th investor (State X) $= \mathrm{D_{Xi}}$,	Dividend for *i*th investor (State Y) $= \mathrm{D_{Yi}}$

Bayesian trader computes the posterior expected value conditional on the observed signal:

$$\begin{aligned}\text{Exp. Value (Dividend} \mid \text{Signal G)} &= \mathrm{D_{xi}}(\alpha\beta_1/\alpha\beta_1 + (1-\alpha)\beta_2) \\ &\quad + \mathrm{D_{Yi}}((1-\alpha)\beta_2/(\alpha\beta_1 + (1-\alpha)\beta_2)), \\ \text{Exp. Value (Dividend} \mid \text{Signal B)} &= \mathrm{D_{xi}}(\alpha(1-\beta_1)/(\alpha(1-\beta_1) + (1-\alpha)(1-\beta_2))) \\ &\quad + \mathrm{D_{Yi}}((1-\alpha)(1-\beta_2)/(\alpha(1-\beta_1) \\ &\quad + (1-\alpha)(1-\beta_2))). \qquad (1)\end{aligned}$$

These expected values determine the feasible range of bid and ask messages for this trader. For bids, the feasible range is from 0 to the bayesian expected value (ceiling). For offers, the range is from the bayesian expected value (floor) to 200. Bids and offers are generated randomly, distributed independently, identically, and uniformly over these ranges. Note that this is a stationary, not a learning strategy. Bayes' theorem is applied only to determine the posterior expected value of dividends conditional on the signal, and this value is used as the limit for zero-intelligence trading strategy. The actual market transactions do not affect the bidding behavior of this trader in any way. Consequently, we should expect the behavior in the later periods of a market to be the same as in the first period.

Trading strategy 2 is an empirical bayesian strategy. At the beginning of the first period of each market, these traders start with a diffuse prior over the two states of the world, X and Y, ignoring the given base rate (α) for the market. Instead, they accumulate and use the realized relative frequencies of each state conditional on the two signals from prior periods of a market. At the end of the t'th period of a market, let

Number of Periods with G Signal $= \mathrm{N_G}$
Number of Periods with B Signal $= \mathrm{t} - \mathrm{N_G}$
Number of G Periods with State X $= \mathrm{NG_X}$
Number of G Periods with State Y $= \mathrm{N_G} - \mathrm{NG_X}$
Number of B Periods with State Y $= \mathrm{N_{BY}}$
Number of B Periods with State X $= \mathrm{t} - \mathrm{N_G} - \mathrm{N_{BY}}$

The empirical bayesian trader uses these frequencies to compute expected dividend conditional on the observed signals:

$$\begin{aligned}\text{Exp. Value (Dividend} \mid \text{Signal G)} &= \mathrm{D_{xi}}(\mathrm{N_{GX}}/\mathrm{N_G}) + \mathrm{D_{yi}}((\mathrm{N_G} - \mathrm{N_{GX}})/\mathrm{N_G}), \\ \text{Exp. Value (Dividend} \mid \text{Signal B)} &= \mathrm{D_{xi}}(\mathrm{t} - \mathrm{N_G} - \mathrm{N_{BY}})/(\mathrm{t} - \mathrm{N_G}) + \mathrm{D_{Yi}}(\mathrm{N_{BY}})/(\mathrm{t} - \mathrm{N_G}). \qquad (2)\end{aligned}$$

These conditional expected values are the ceilings (floors) for ZI bids (offers) made by the empirical bayesian traders. Unlike bayesian traders' expected values, these expected values change each period with the accumulation of relative frequencies. The behavior of empirical bayesian traders can be quite erratic in the early periods of a market. However, as observations are accumulated over the periods, the law of large numbers should lead to convergence toward the same bayesian expected values of type 1 traders. Many researchers have argued that human subjects are not bayesian in experiments because they do not accept the base rates given to them by the experimenter. The empirical bayesian trader operationalizes the idea that human behavior would become bayesian if experiments could be run for a long enough period of time, and the traders respond on the basis of the actual relative frequencies experienced during the experiment (e.g. Libby, 1985 and Butt, 1988).

Trading strategy 3 combines two simple heuristics. Experimental studies of individual behavior suggest that individuals use the "representativeness" and the "anchor-and-adjust" heuristics to make judgments and decisions. According to Tversky and Kahneman (1974), the representativeness heuristic is used by human agents to generate an initial cognitive representation; an agent who uses the representativeness heuristic determines the subjective probability that an object A belongs to class B, or A originates from B, based on the degree to which A resembles B. In our context, using the notation given above, if β_1 were more than 0.5 and β_2 were less than 0.5, agents following the representativeness heuristic will infer that the state is X(Y) whenever they see signal G(B). The use of the representativeness heuristic to make probability judgments leads to serious errors because the heuristic does not incorporate all factors that should affect probability judgments. One critical factor ignored by this heuristic is the prior probability, or base rate frequency, of the possible outcomes.[2]

An anchor-and-adjust strategy is used by human agents in sequential tasks where beliefs have to be updated as new information is received. An agent who uses the anchor-and-adjust heuristic will use an initial probability value as the anchor (or starting point) and then make adjustments as new information is received. If we denote the anchor probability as p, the amount of adjustment as k, and the final outcome from the anchor-and-adjust process as $S(p)$, the anchor-and-adjust model can be denoted as $S(p)=p+k$. The anchor probability can come from a variety of sources including the formulation of the problem, memory or can be explicitly suggested by an agent. A key feature of the adjustment process is that it is insufficient (compared to the bayesian adjustment process) so that the final probability is biased toward the initial (anchor) value.

Type 3 computer traders begin by setting an initial current aspiration level or "CAL" (Simon, 1981) based on the signal received. Using the representativeness heuristic, these traders overreact to the signal and treat it as if it were perfect. For example, if the first signal is G, each trader makes a judgment about which state is most likely to generate signal G based on the diagnosticity of signal G (representativeness). Suppose all traders

[2]Kahneman and Tversky (1972) conducted a study where subjects were asked to assess the probability that various personality descriptions belonged to an engineer rather than a lawyer. The results indicated that the subjects ignored the base rates and treated the descriptions as being perfectly deterministic. These initial findings have been replicated in a variety of contexts (Bar-Hillel, 1980, Nisbett and Ross, 1980).

conclude that State X will be realized when signal G is observed. Each trader i sets his/her initial current aspiration level (CAL) equal to his/her dividend value in State X (D_{Xi}). If the first signal is B, each trader makes a judgment about which state is most likely to generate signal B. Suppose all traders conclude that State Y will be realized when signal B is observed. All traders set their initial current aspiration level (CAL) equal to their respective dividend values in State Y (D_{Yi}). In each state (e.g. State X), dividend values differ across traders creating an opportunity for a profitable trade even if all traders believe that they are in State X.

Type 3 computer traders use these state contingent initial current aspiration level (CALs) as ceilings (floors) for their bids (offers) in ZI mode of trading as described above for Type 1 and Type 2 traders. For bids, the feasible range is from 0 to the current aspiration level (ceiling). For offers the range is from the current aspiration level(floor) to 200.

After each transaction, Type 3 traders update their aspiration level using a simple anchor-and-adjust (i.e. first order adaptive) process:

$$\text{Revised aspiration level} = (1-\gamma)(\text{CAL}) + (\gamma^{*}\,\text{Transaction Price}), \tag{3}$$

using a given value of the adaptive parameter γ for the simulation (see Table 1). Trading resumes with the updated CALs serving as constraints on the opportunity sets of the traders until the next transaction occurs, and this cycle is repeated until the end of the period. At the end of a period the state is revealed publicly to all traders. At the end of the period, traders adjust their aspiration level by incorporating the actual dividend received:

$$\text{End of period aspiration level} = (1-\delta)(\text{CAL}) + (\delta^{*}\,\text{dividend realized}), \tag{4}$$

using a given value of the adaptive parameter δ for the simulation (see Table 1). At the end of each period, the traders' holdings are refreshed to their initial endowments to begin the next period. However, the end-of-the-period CAL is retained for use in the following period.

It is important to note that the Type 3 computer traders are not bayesian in processing information. The traders select their initial CAL by overreacting to an imperfect signal as if it were perfect (representativeness). The traders select messages randomly, subject to limits imposed by this CAL, until a transaction occurs. After each transaction, traders adjust their aspiration level in the direction of the transaction price using a simple anchor-and-adjust procedure. Traders also adjust their aspiration level at the end of the period in the direction of the dividend realized using a similar anchor-and-adjust procedure. These computer traders engage in very simple computations using judgment heuristics documented in psychological studies of individual behavior.

4. Design

Design parameters for the simulations are give in Table 1. We ran 100 markets. In every market, an independent realization of uniformly distributed random variables is obtained for the parameters listed in Table 1. There is no correlation among the parameter values used in the computational experiment. Once a set of parameter values had been chosen, three 100-period markets were run, one for each of the three types of traders. In each market there were 20 traders who traded a two-state, single period asset. Before trading was started in each period, a random state of the nature (X or Y), and a random signal (green or brown) were drawn with the following probabilities (see Fig. 1):

As parameters α, β_1, β_2, γ and δ, and the individual dividends change from market to market, so do the equilibrium prices. The bayesian equilibrium price is determined by assuming that each trader values the asset at its bayesian posterior expected value, and perfect competition prevails in the market. The representativeness heuristic equilibrium price is determined by assuming that, after observing the signal, each trader values the asset at the dividend associated with the state that generates that signal with a higher probability (see Duh and Sunder, 1986):

$$\begin{array}{ll} \text{If } \beta 1 > \beta 2, & \text{Green signal} \Rightarrow \text{State X and dividend } D_{Xi} \\ & \text{Brown signal} \Rightarrow \text{State Y and dividend } D_{Yi}; \\ \text{If } \beta 1 < \beta 2, & \text{Green signal} \Rightarrow \text{State Y and dividend } D_{Yi} \\ & \text{Brown signal} \Rightarrow \text{State X and dividend } D_{xi}. \end{array} \tag{5}$$

5. Experimental results

5.1. Prices

Since the equilibrium prices vary from market to market, aggregate presentation of the results must be normalized. For each period t of each market m, we calculate the mean transaction price P (dropping subscripts m and t). This period-wise mean transaction price is normalized to obtain P^n:

$$P^n = 100(P - P_b)/(P_r - P_b) \tag{6}$$

where P_b and P_r are the bayesian and the representativeness competitive equilibrium prices respectively for the signal observed during that period. If P is equal to the representativeness competitive equilibrium price, the normalized price is 100; if P is equal to the bayesian competitive equilibrium price, the normalized price is 0. P outside these two equilibrium prices yields a normalized price outside the 0–100 range. Since the bayesian equilibrium prices under the two signals are always bracketed by the two representativeness equilibrium prices, a normalized price below zero means that observed prices deviate from these equilibria in the direction of the mean of the equilibrium prices under the two signals. This normalized period-wise mean transaction price can be aggregated across the 100 markets to explore and compare the aggregate market behavior under the three types of traders.

Given the random choice of market parameters, P_r and P_b can be arbitrarily close to each other. Since their difference appears in the denominator of the normalized mean price, it is possible that the first moment of the distribution of P^n may not exist. We therefore report the order statistics (period-by-period median and inter-quartile range) of P^n across the 100 markets in Fig. 2 (for a total of 30,000 market periods).

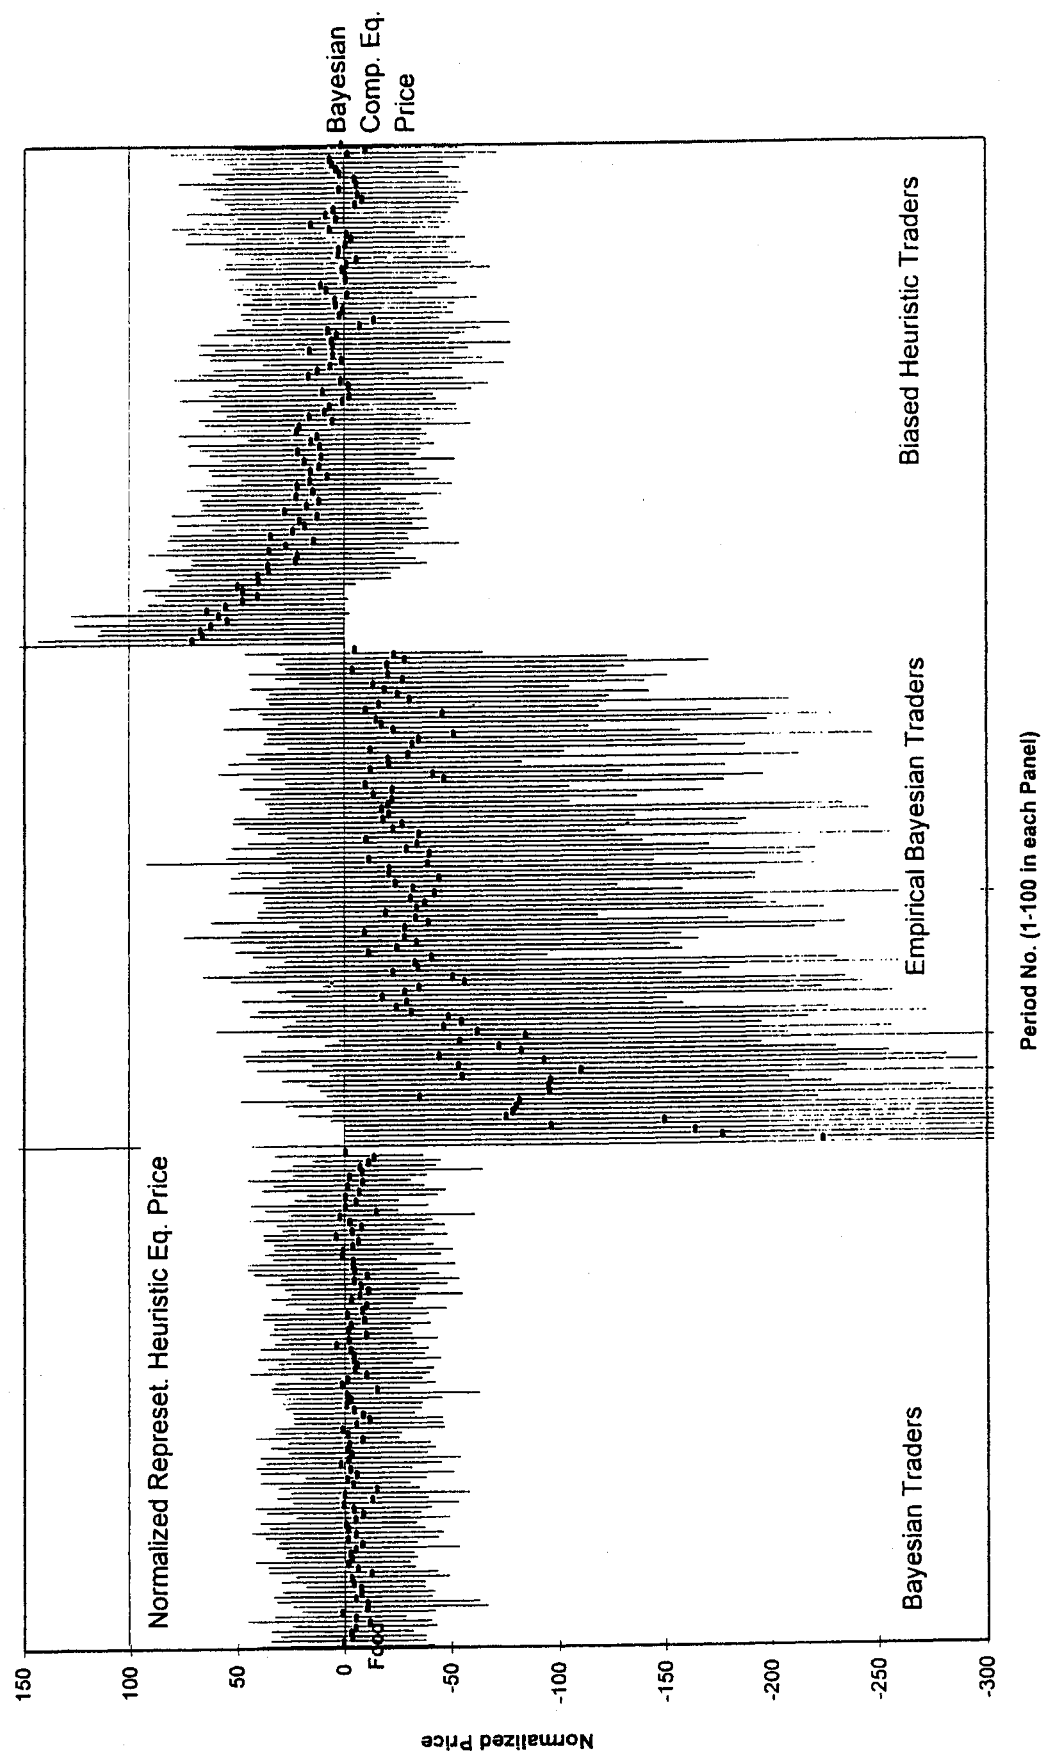

Fig. 2. Median and inter-quartile range of normalized period-wise average transaction prices.

Three panels of Fig. 2, one for each type of investor, summarize the price data. The first panel is for bayesian traders. Horizontal lines at 0 and 100 are the bayesian and representative heuristic competitive equilibrium respectively. A vertical line for each period indicates the inter-quartile range of period-wise average transaction price across the 100 different markets; the small dark rectangular marker on the vertical line is the median. As we would have expected of the control group of bayesian traders, the median lies close to the bayesian norm of 0. This suggests that the Gode and Sunder (1993) results about the behavior of double auctions with ZI traders also hold for the case of uncertainty and imperfect information. Since bayesian traders act as zero-intelligent bidders using bayesian expectation as constraints, transaction prices show a fairly high variance. Since the bayesian expectation does not change over periods, this panel shows, as expected, a stationary pattern of behavior over the 100 periods.

The transaction prices of a market composed of only type 2 (empirical bayesian) traders are shown in the second panel of Fig. 2. The median of the normalized mean prices begins substantially below the bayesian equilibrium price (at about – 250). This means that during the first period, the mean transaction price tended to lie close to the average of the bayesian competitive equilibrium prices for the two signals (G and B). We expected, and observe, that a market composed of only empirical bayesian traders will start below 0 normalized price, behave erratically at the beginning of a market, but will gradually converge, under the force of the law of large numbers, to the bayesian equilibrium price, after they accumulate enough experience of the frequency of each state conditional on the respective signal.

The results shown in the second panel suggest that even if all traders were empirical bayesian traders, the results of an experiment conducted with human traders would generate data which suggest that the market is not bayesian because human subjects rarely get the opportunity to experience ten or more occurrences of any individual signal. These results suggest that experiments with human subjects who do not accept the base rates, but can keep track of relative frequencies, would require exposing the subjects to about twenty experiences with each signal to obtain prices in the neighborhood of the bayesian equilibrium.

The transaction prices of a market composed of only type 3 biased heuristic traders are shown in the third panel of Fig. 2. The data show that in the first period of markets, the median of normalized trading prices is about 70, which is much closer to the representativeness than the bayesian equilibrium prediction mark. That is not surprising because all the traders in this market use the representativeness heuristic to determine the boundaries for their zero-intelligence messages (bids and offers) in the first period of each market. However, as they anchor-and-adjust this initial assessment in light of the market transaction prices, and the end-of-period dividends, there is a progressive tendency for the median of average transaction prices to move ever closer to the bayesian equilibrium mark (normalized to zero in the figure).

Of the three panels of Fig. 2, the third is the only one in which prices exhibit any significant tendency to lie in the approximate direction of the representative norm of 100. Yet, these representative heuristic assessments themselves generate market transaction prices that gradually pull the traders' assessments in the direction of the bayesian norm through the anchor-and-adjust heuristic. The magnitude of this pulling force is arbitrary (depending on the values of γ and δ we chose for the simulation, see Table 1), but the direction – toward the bayesian norm – is unmistakable. Biased heuristics of individual traders do not directly translate into biases in aggregate market level outcomes.

5.2. Allocative efficiency

Allocative efficiency of markets depends on who holds the assets, and not on the prices. Efficiency is calculated as the ratio of the bayesian expected value of the assets for traders who actually hold them at the end of trading each period to the bayesian expected value of the assets under the equilibrium allocation. Three panels of Fig. 3 show the period-by-period median and inter-quartile range of allocative efficiency of the three sets of markets with three different types of traders. It is hardly surprising that the median efficiency of markets with bayesian traders is virtually 100% throughout. Even with zero-intelligence traders it shows little dispersion because in an asset market, inefficient allocation can be improved upon by resale of security to another trader who values the security even more. For empirical bayesian traders, median efficiency starts around 92% and gradually rises close to 100% with dispersion of a few percent. For markets with biased heuristic traders, median efficiency lies between 93% and 95%. It could be raised by a few points, but not much more, by doubling the number of iterations in these simulations with ZI traders. The inter-quartile ranges indicate much greater dispersion of efficiency in the 80–100% range.

5.3. Discussion

A comparison of the trading outcomes of bayesian traders and biased heuristic traders shows that the latter begin trading at prices close to the representativeness equilibrium (of 100), whereas bayesian traders begin trading at prices close to the bayesian equilibrium (of 0) or at prices below the bayesian equilibrium (for empirical bayesian traders). However, there is a clear pattern of trading activity that shows that the trading prices of biased heuristic traders and empirical bayesian traders tend toward the bayesian equilibrium price.

A market composed solely of biased heuristic traders tends repeatedly (in a 100 market simulation run) toward the bayesian equilibrium price without any arbitrageurs, natural selection, or cancellation of random responses. The early transaction prices are substantially closer to the representativeness equilibrium mark of 100, and differ from the bayesian equilibrium mark, allowing an arbitrageur (or specialist) the opportunity to make a substantial profit in the early periods of such a market. However, even in a market composed solely of biased heuristic traders, and no arbitrageurs, the opportunity to earn arbitrage profits diminishes rapidly with time.

Experiments with human subjects generally run for some 10–20 periods in which human subjects may experience a particular signal some eight to twelve times. The results of our computational experiment suggest that one of the reasons human experiments may generate mixed results may be because they are terminated prematurely. Our results show that even if all traders were empirical bayesians, results obtained from an experiment where the traders were exposed to each signal five times would generate results that would appear as if the market does not converge to the bayesian equilibrium.

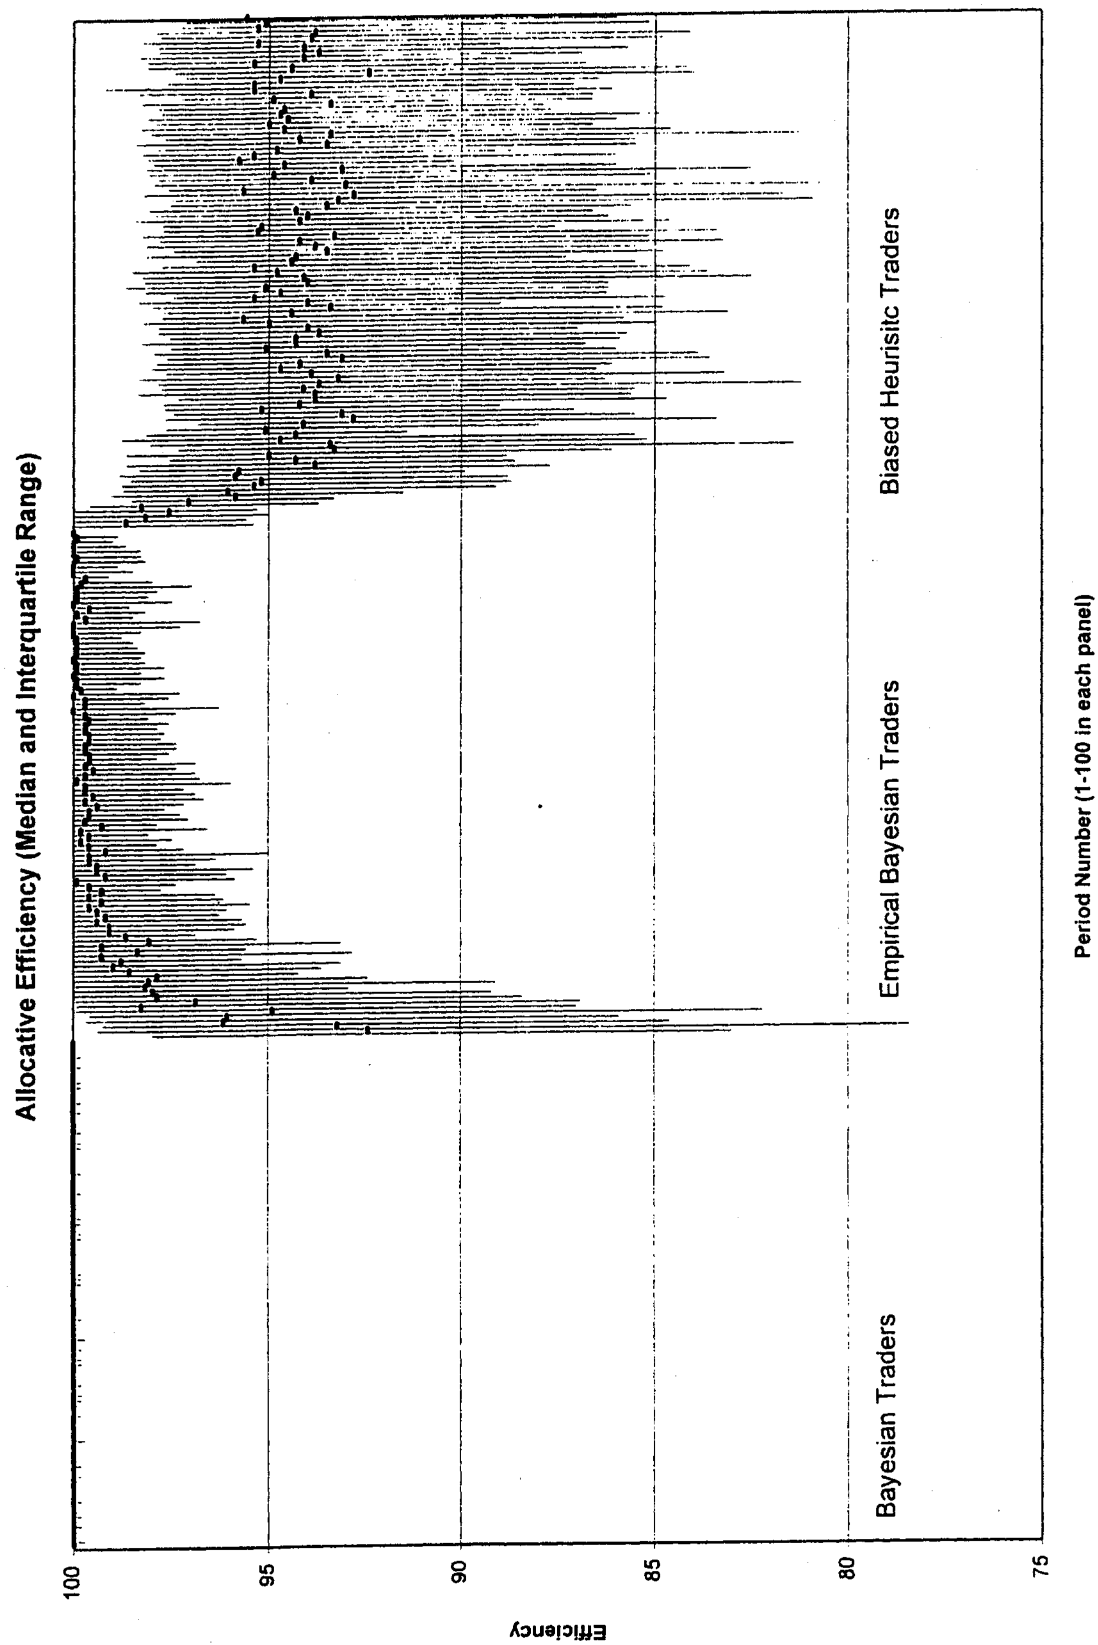

Fig. 3. Allocative eficiency (median and inter-quartile range).

6. Conclusions

Convergence to the bayesian equilibrium price occurs in markets populated solely by biased heuristic traders. This result is obtained without the presence of various forces normally credited in the economics literature with making human subjects rational. There are no smart traders in our market who make marginal trades (no arbitrageurs). There is also no learning of new information processing heuristics in our market. All traders begin by using biased information processing heuristics documented in psychological experiments. These traders continue to use these biased heuristics throughout the computational experiment. The traders do not learn any new information processing heuristics and the biased heuristics are not altered or corrected in any manner. There are also no evolutionary forces operating in our experiment to drive biased traders out of the market.

Our results generalize and extend Gode and Sunder's (1993, 1995) findings in a market with certainty to a market with uncertainty. We find that the information processing demands of double auctions are quite low even when there is uncertainty about the state of nature. Markets populated by traders who use simple adaptive heuristics can attain a bayesian equilibrium even though the heuristics are myopic and ignore base rate information. Double auctions seem to possess a structure in which the bayesian equilibrium can be achieved even if all individual traders are biased in their information processing activities.

Appendix A

A.1. Viewing the simulation live on your own computer

You can get a better sense of the dynamics of the markets reported in this article by looking at the actual simulations on your own computer by following the instructions given in this Appendix. We also describe what you will see on your computer screen.

How to get the software:

Download a copy of file simul1.exe and readme.txt from ftp.andrew.cmu.edu through anonymous ftp and store it on your hard drive or a floppy.

Anonymous ftp to:	ftp.andrew.cmu.edu
Change directory by using command:	cd /pub/gsia/ss8a/jamal.sunder.1
Download explanation of the program by command:	get readme.txt
Change to binary mode of transfer by issuing command:	bin
Download simulation program by issuing command:	get simul1.exe

Terminate your ftp session by typing: quit

Hardware requirements:

IBM/clone 386 or better with VGA or Enhanced VGA monitor.

How to run the simulation:

1. Enter the directory in which file simul1.exe is stored, type simul1, and press the enter key.
2. In response to the first question, type 1, 2 or 3, depending on which simulation you wish to run(1 for bayesian traders, 2 for empirical bayesian traders, or 3 for biased heuristic traders) and hit the return key.
3. The program prompts you for the number of markets you wish to run, and the number of periods in each market. Enter a positive integer in response to each question, and press the enter key.
4. If you want the computer to pause after each full screen, respond 1 (yes), 0 (no) otherwise. It will always pause at the end of the summary screens.
5. All other parameters are picked by the computer as preset values or random draws. Sit back and enjoy the show. If you chose the pause option, press the return key whenever the simulation pauses to allow you to review the screen.

What you see on the screen:

1. On the left hand of your simulation screen is a box with demand and supply functions. The quantity axis has the length of 19 (1 unit for each of the 19 traders in the market). The price axis is 0–200. The firm lines are the bayesian demand and supply functions; broken lines are the demand and supply functions derived from the assumption that the traders use the representativeness heuristic. Green (brown) supply and demand functions correspond to periods in which the imperfect signal is Green (Brown). Please note that these are asset markets, and all traders are buyers as well as sellers; hence the symmetry of the demand and supply functions.
2. Firm horizontal green (brown) lines extending across the screen define the bayesian equilibrium price for Green (Brown) periods. Broken horizontal green (brown) lines extending across the screen define the representativeness heuristic equilibrium price for Green (Brown) periods.
3. Each transaction price is plotted to the right of the demand/supply box in green (brown) circles for periods when the signal is green (brown). At the end of a period, a vertical line is drawn to separate the data for one period from the next. When the screen if filled, the program pauses for you to hit the enter key on your keyboard before proceeding further (if you choose the pause option).
4. If you are running the biased heuristic traders, the circle for each transaction price is preceded by two vertical lines in green (brown) indicating the range of the individual current aspiration levels (CALs for bidding and asking) of all traders in the market. In addition, a horizontal tick mark on these lines indicates the mean of individuals CALs.

5. After completion of all periods of a market, summary statistics of that market appear on three screens. The program pauses at the end of each screen for you to hit the enter key before proceeding to the next screen. Hit enter key at the end of the final screen to terminate the program.
6. Please note that the three summary screens in this simulation give you a market-by-market summary. In contrast, the figures in the published paper aggregate the results of the first two summary screens across one hundred such independently run markets.
7. The first summary screen shows the period-by-period mean and range of transaction prices for the market just completed. The color corresponds to the signal for the period and order is chronological. The demand/supply box and the horizontal lines for equilibrium price ranges are included for the purpose of comparison.
8. The second summary screen shows the efficiency of each period of the market you simulated on a 0–100% grid in chronological order. Color of the circle corresponds to the signal for the period.
9. The third summary screen shows the mean and standard deviation of i'th transaction across all green (brown) periods of the market you just finished simulating. Since there is a total of 19 traders in the market, each with endowment of one security, the expected trading volume is 10. Some periods will have fewer or more than 10 trades. Therefore this screen will show the mean and standard deviations of 1st through about 12th transaction price across all green (all brown) periods. Note that even if you run each market for, say, 100 periods, you will only see about 12–13 green and an equal number of brown means and standard deviations on this summary screen. The purpose of these statistics is to see if there is a tendency for transaction prices within trading periods to move in the direction of any of the benchmark equilibrium predictions indicated by the horizontal lines. The demand/supply box and the horizontal lines for equilibrium price ranges are included for the purpose of comparison.
10. After the third summary screen, the simulation continues to the next market if you have asked for more.

Parameters of the simulation:

Distributions of various parameters of the simulation are given in the text of the paper. Other parameters are as follows:

Number of traders	19
Number of iterations each period	5,000

References

Anderson, Matthew J. and Shyam Sunder, 1995, Professional traders as intuitive Bayesians, Organizational Behavior and Human Decision Processes, 64(2), 185–202.

Bar-Hillel, Maya, 1980, The base rate fallacy in probability judgments, Acta Psychologica, 44, 211–233.

Becker, Gary S., 1962, Irrational behavior and economic theory, Journal of Political Economy, 70, 1–13.

Bonini, Charles P., 1963, Simulation of Information and Decision Systems in the Firm, (Prentice-Hall, Englewood Cliffs, NJ).

Bosch, Antoni and Shyam Sunder, 1994, Tracking the invisible hand: Convergence of double auctions to competitive equilibrium, Working Paper, (Carnegie Mellon University, Pittsburgh, PA).

Butt, Jane, 1988, Frequency judgments in an audit-related task, Journal of Accounting Research, 26, 315–330.

Camerer, Colin F., 1987, Do biases in probability judgments matter in markets? Experimental evidence, American Economic Review, 77, 981–997.

Cohen, Kalman J. and Richard M. Cyert, 1965, In: James G. March, Ed., Simulation of Organizational Behavior, Handbook of Organizations, (Rand McNally, Chicago).

Cohen, Kalman J., Richard M. Cyert, William R. Dill, Alfred A. Kuehn, Merton H. Miller, Thedore A. Van Wormer and Peter R. Winters, 1960, The Carnegie Tech management game, Journal of Business, 33, 303–327.

Cohen, Kalman J., Richard M. Cyert, James G. March and Per O. Soelberg, 1963, A general model of price and output determination, In: R.M. Cyert and J.M. March, Eds., A Behavioral Theory of the Firm, (Prentice-Hall, Englewood Cliffs, NJ), Chapter 8.

Coleman, James S., 1986, Psychological structure and social structure in economic models, The Journal of Business, 59(4), 365–369.

Cyert, Richard M. and James G. March, 1963, A Behavioral Theory of the Firm, (Englewood Cliffs, Prentice-Hall, NJ).

Cyert, Richard M., James G. March and Charles G. Moore, 1963, A specific price and output model, In: R.M. Cyert and J.G. March, Eds., A Behavioral Theory of the Firm, (Englewood Cliffs, Prentice-Hall, NJ), Chapter 7.

Duh, Rong R. and Shyam Sunder, 1986, Incentives, learning and processing of information in a market environment: An examination of the base rate fallacy, In: Shane Moriarty, Ed., Laboratory Market Research, (University of Oklahoma Press, Norman, OK).

Duh, Rong R. and Shyam Sunder, 1993, Economic agent as an intuitive bayesian: Experimental evidence, Cuadernos Economicos, 54, 101–127.

Edwards, Ward and D. von Winterfeldt, 1986, On cognitive illusions and their implications, In: H.R. Arkes and K.R. Hammond, Eds., Judgment and Decision Making: An Interdisciplinary Reader, (Cambridge University Press, New York).

English, J. and J. Allison, 1993, Chance and rationality, Psychological Review, 100(3), 535–545.

Grether, David M., 1980, Bayes' rule as a descriptive model: The representativeness heuristic, Quarterly Journal of Economics, 95, 537–557.

Grether, David M. and Charles R. Plott, 1979, Economic theory of choice and the preference reversal phenomenon, American Economic Review, 69(4), 623–638.

Gode, Dhananjay J. and Shyam Sunder, 1993, Allocative efficiency of markets with zero intelligence traders: Market as a partial substitute for individual rationality, Journal of Political Economy, 101, 119–137.

Gode, Dhananjay J. and Shyam Sunder, 1995, What makes markets efficient?: A statistical approach, Working Paper, (Carnegie Mellon University, Pittsburgh, PA).

Jamal, Karim and Shyam Sunder, 1991, Money vs Gaming: Effects of salient monetary payments in double oral auctions, Organizational Behavior and Human Decision Processes, 49, 151–166.

Jamal, Karim, Paul E. Johnson and R. Glen Berryman, 1995, Detecting framing effects in financial statements, Contemporary Accounting Research, 12(1), 85–105.

Kahneman, Daniel and Amos Tversky, 1972, Subjective probability: A judgment of representativeness, Cognitive Psychology, 3, 430–454.

Libby, Robert, 1985, Availability and the generation of hypotheses in analytical review, Journal of Accounting Research, 23, 648–667.

Newell, Allen and Herbert A. Simon, 1972, Human Problem Solving, (Prentice-Hall, Englewood Cliffs, NJ).

Nisbett, Richard E. and Lee Ross, 1980, Human Inference: Strategies and Shortcomings in Social Judgment, (Prentice-Hall, Englewood Cliffs, NJ).

Plott, Charles R., 1982, Industrial organization theory and experimental economics, Journal of Economic Literature, 20, 1485–1527.

Russell, R. and R. Thaler, 1985, The relevance of quasi-rationality in competitive markets, The American Economic Review, 5(75), 1071–1082.

Simon, Herbert A., 1973, The structure of ill structured problems, Artificial Intelligence, 4, 181–201.

Simon, Herbert A., 1981, The Sciences of The Artificial, 2nd edn., (MIT Press, Cambridge, MA).

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

Tversky, Amos and Daniel Kahneman, 1974, Judgment under uncertainty: Heuristics and biases, Science, 185, 1124–1131.

Tversky, Amos and Daniel Kahneman, 1986, Rational choice and the framing of decisions, The Journal of Business, 59(4), 251–278.

INSTRUCTIONS TO AUTHORS

(1) Papers must be in English.

(2) Papers for publication should be sent in triplicate to:

Professor Timothy Cason
Assistant Managing Editor
Department of Economics
University of Southern California
LOS ANGELES, CA, 90089–0253, U.S.A.

Unsolicited manuscripts must be accompanied by a submission fee of US$ 35.00. Cheques should be sent together with the manuscript and made payable to the *Journal of Economic Behavior and Organization.*

Submission of a paper will be held to imply that it contains original unpublished work and is not being submitted to publication elsewhere. The Editor does not accept responsibility for damage or loss of papers submitted. Upon acceptance of an article, author(s) will be asked to transfer copyright of the article to the publisher. This transfer will ensure the widest possible dissemination of information.

(3) Submission of accepted papers as *electronic manuscripts*, i.e., on disk with accompanying manuscript, is encouraged. Electronic manuscripts have the advantage that there is no need for rekeying of text, thereby avoiding the possibility of introducing errors and resulting in reliable and fast delivery of proofs. The preferred storage medium is a 5.25 or 3.5 inch disk in MS-DOS format, although other systems are welcome, e.g., Macintosh (in this case, save your file in the usual manner; do not use the option 'save in MS-DOS format'). Do not submit your original paper as electronic manuscript but hold on to the disk until asked for this by the Editor (in case your paper is accepted without revisions). Do submit the accepted version of your paper as electronic manuscript. Make absolutely sure that the file on the disk and the printout are identical. Please use a new correctly formatted disk and label this with your name; also specify the software and hardware used as well as the title of the file to be processed. Do not convert the file to plain ASCII. Ensure that the letter 'l' and digit '1', and also the letter 'O' and digit '0' are used properly, and format your article (tabs, indents, etc.) consistently. Characters not available on your word processor (Greek letters, mathematical symbols, etc.) should not be left open but indicated by a unique code (e.g., gralpha, ⟨alpha⟩, @, etc., for the Greek letter α). Such codes should be used consistently throughout the entire text; a list of codes used should accompany the electronic manuscript. Do not allow your word processor to introduce word breaks and do not use a justified layout. Please adhere strictly to the general instructions below on style, arrangement and, in particular, the reference style of the journal. Manuscripts should be typewritten on one side of the paper only, double-spaced with wide margins. All pages should be numbered consecutively. Titles and subtitles should be short. References, tables, and legends for figures should be typed on separate pages. The legends and titles on tables and figures must be sufficiently descriptive such that they are understandable without reference to the text. The dimensions of figure axes and the body of tables must be clearly labeled in English.

(4) A short note on author's current position, affiliation, and other relevant biographical information, including publications, should also be provided.

(5) The first page of the manuscript should contain the following information: (i) the title; (ii) the name(s) and institutional affiliation(s) of the author(s); (iii) an abstract of not more than 100 words; (iv) at least one classification code according to the Classification System for Journal Articles as used by the *Journal of Economic Literature*; in addition, up to five key words should be supplied. A footnote on the same sheet should give the name and present address of the author to whom proofs and reprint order should be addressed.

(6) Acknowledgements and information on grants received can be be given before the References or in a first footnote, which should not be included in the consecutive numbering of footnotes.

(7) Important formulae (displayed) should be numbered consecutively throughout the manuscript as (1), (2), etc. on the right-hand side of the page. Where the derivation of formulae has been abbreviated, it is of great help to referees if the full derivation can be presented on a separate sheet (not to be published).

(8) Footnotes should be kept to a minimum and be numbered consecutively throughout the text with superscript arabic numerals.

(9) The References should include only the most relevant papers. In the text, references to publications should appear as follows: 'Smith (1969) reported that ...' or 'This problem has been a subject in literature before [e.g. Smith (1969, p. 102)]'. The author should make sure that there is a strict 'one-to-one correspondence' between the names (years) in the text and those on the list. At the end of the manuscript (after any appendices), the complete references should be listed as:

For monographs

Tobin, James, 1971, Essays in economics, vol. 1, Macroeconomics (North-Holland, Amsterdam).

For contributions to collective works

Glejser, Herbert, 1970, Predictive world models, in: Colette Duprez and Etienne Sadi Kirschen, eds., Megistos, A world income and trade model for 1975 (North-Holland, Amsterdam) 3–16.

For periodicals

Goldfeld, Stephen M. and Richard E. Quandt, 1973, A Markov model for switching regressions, Journal of Econometrics, 1, 3–15.

(10) Illustrations should be provided in triplicate (1 original drawn in black ink on white paper +2 photocopies). Care should be taken that lettering and symbols are of a comparable size. The drawings should not be inserted in the text and should be marked on the back with figure numbers, title of paper, and name of the author. All graphs and diagrams, should be referred to as figures and should be numbered consecutively in the text in arabic numerals. Graph paper should be rules in blue and grid lines to be shown should be inked black. Illustrations of insufficient quality which have to be redrawn by the publisher will be charged to the author.

(11) All unessential tables should be eliminated from the manuscript. Tables should be numbered consecutively in the text in arabic numerals and typed on separate sheets.

Any manuscript which does not conform to the above instructions may be returned for the necessary revision before publication. Page proofs will be sent to the authors. Corrections other than printer's errors may be charged to the author. 25 reprints of each paper are supplied free; additional reprints are available at cost if they are ordered when the proof is returned.