SIMPLE AGENTS, INTELLIGENT MARKETS

by

Karim Jamal, Michael Maier, and Shyam Sunder

COWLES FOUNDATION PAPER NO. 1801

COWLES FOUNDATION FOR RESEARCH IN ECONOMICS
YALE UNIVERSITY
Box 208281
New Haven, Connecticut 06520-8281

2016

http://cowles.yale.edu/

Comput Econ
DOI 10.1007/s10614-016-9582-3

Simple Agents, Intelligent Markets

Karim Jamal[1] · **Michael Maier**[1] · **Shyam Sunder**[2]

Accepted: 19 April 2016

Abstract Attainment of rational expectations equilibria in asset markets calls for the price system to disseminate agents' private information to others. Markets populated by human agents are known to be capable of converging to rational expectations equilibria. This paper reports comparable market outcomes when human agents are replaced by boundedly-rational algorithmic agents who use a simple means-end heuristic. These algorithmic agents lack the capability to optimize; yet outcomes of markets populated by them converge near the equilibrium derived from optimization assumptions. These findings point to market structure (rather than cognition or optimization) being an important determinant of efficient aggregate level outcomes.

Keywords Rational expectations · Bounded rationality · Means-end heuristic · Information dissemination · Minimally intelligent agents

JEL Classification C92 · D44 · D50 · D70 · D82 · G14

Our knowledge of the very narrow limits of human rationality must dispose us to doubt that business firms, investors or consumers possess either the knowledge or computational ability that would be required to carry out the rational expectations strategy.

Herbert Simon (1969)

Electronic supplementary material The online version of this article (doi:10.1007/s10614-016-9582-3) contains supplementary material, which is available to authorized users.

✉ Shyam Sunder
shyam.sunder@yale.edu

1 University of Alberta, Edmonton, Alberta T6G 2R6, Canada

2 Yale University, New Haven, CT 06520-8200, USA

Published online: 11 May 2016

The claim that the market can be trusted to correct the effect of individual irrationalities cannot be made without supporting evidence, and the burden of specifying a plausible corrective mechanism should rest on those who make this claim.

Tversky and Kahneman (1986)

The principal findings of experimental economics are that impersonal exchange in markets converges in repeated interaction to the equilibrium states implied by economic theory, under information conditions far weaker than specified in the theory.

Vernon Smith (2008)

1 Introduction

A central feature of economic theory is derivation of equilibrium in economies populated by agents who optimize some well-ordered function such as profit or utility. Although it is recognized that actions of economic agents are subject to institutional constraints and feedback (North 1990), exploration of the extent to which equilibrium arises from characteristics of the institutional environment, as opposed to the behavior of individuals, has been limited; Becker's (1962) derivation of downward slope of demand functions is a notable exception. The normal modeling technique is to ascribe sophisticated computational abilities to a representative agent to solve for equilibrium (Muth 1961). Plott and Sunder (1982, henceforth PS) have shown that markets with uncertainty and asymmetrically distributed information (with two or three states of the world) disseminate information and converge near rational expectations equilibria when populated with profit-motivated human traders. The present paper asks if the PS results can also be achieved by minimally intelligent traders (Gode and Sunder 1993) using the means-end heuristic and reports an affirmative answer.

Simon (1969, Chapter 3) questioned the plausibility of human agents, with their limited cognitive abilities, forming rational expectations by intuition. Accumulated observational evidence on these cognitive limits of individuals shifted the burden of proof and led to calls for evidence that markets can overcome such behavioral limitations (Thaler 1986; Tversky and Kahneman 1986).

Laboratory studies of markets populated by asymmetrically-informed profit-motivated human subjects reveal that their aggregate level outcomes tend to converge near the predictions of rational expectations theory (Forsythe and Lundholm 1990; Forsythe et al. 1982; Plott and Sunder 1988). However, since complex patterns of human behavior can only be inferred, not observed directly, it is difficult to know from human experiments which elements of trader behavior and faculties are necessary or sufficient for various markets to attain their theoretical equilibria.[1] This difficulty has led to claims that the inability of humans to optimize by intuition implies that economic theories based on optimization assumptions are prima facie invalid [for example, Tversky and Kahneman (1986)].

[1] See for example Dickhaut et al. (2012) regarding conditions where markets with human traders are less likely to conform to predicted equilibria.

Such doubts about the achievability of mathematically derived equilibria, when individual agents are not able to perform complex optimization calculations are understandable. From a constructivist point of view (Smith 2008), rational expectations equilibria place heavy demands on individual cognition to learn others' preferences or strategies, and to arrive at unbiased estimates of underlying parameters of the economy by observing markets. In theory, disseminating and detecting information in markets calls for bootstrapping—rational assessments are necessary to arrive in equilibrium and such assessments require observation of equilibrium outcomes. Cognitive and computational demands on individuals to arrive at economic equilibria, especially rational-expectations equilibria, are high, raising doubts about the plausibility of equilibrium models (Simon 1969).

Replacing humans with algorithms allows us to examine whether the use of certain simple heuristics by individual traders is sufficient for attaining rational expectations equilibria (as a proof of concept). Without claiming that human traders actually use such heuristics, it is possible and useful to determine if heuristics making low computation demands on human reasoning might be sufficient for attaining equilibria in a given market environment. Combining Newell and Simon's (1972) means-end heuristic with Gode and Sunder's (1993, 1997) zero-intelligence (ZI) approach, we find and report that markets with uncertainty and asymmetric information attain outcomes approximating rational expectations equilibria, even when they are populated by simple minimally-intelligent adaptive algorithmic traders. Since the statistical distribution of these outcomes is centered near the PS observations of markets with human traders, the convergence of their outcomes to equilibrium can be attributed to the combination of the market structure and the minimal levels of intelligence and adaptive ability built into the trading algorithms. Since these trader faculties are far less demanding than what is assumed in deriving the equilibria, and certainly within the known human capabilities, we infer that the convergence of markets to rational expectations equilibria emerge mainly from the properties of the market and simple and plausible decision heuristics, rather than from complex and sophisticated optimization (Becker 1962; Gode and Sunder 1993; Gigerenzer and Todd 1999; Smith 2008).

2 Background and Theory

Instead of assuming sophisticated information processing capabilities and maximization objectives of agents, we can think of market structure constraining human behavior to guide their aggregate level outcomes to the neighborhood of theoretical equilibria. Becker (1962) showed that the downward slope of demand functions arises from individuals having to act within their budget constraints, even if they choose randomly from their opportunity sets. Smith (1962) reported that classroom double auction markets populated by a mere handful of profit-motivated student traders with minimal information arrive in close proximity of Walrasian equilibrium. Moreover, Smith's auction markets had little resemblance to the tâtonnement story often used to motivate theoretical derivations of equilibria.

Gode and Sunder (1993) combined Becker's constrained random choice with Smith's double auctions and reported the results of computer simulations of simple double auctions populated by “zero intelligence” (henceforth ZI) algorithmic traders who bid or ask randomly within their budget constraints (i.e., buyers do not bid above their private values and sellers do not ask below their private costs). Although these traders do not remember, optimize, maximize profits, or learn, simulated markets populated by such traders also reach the proximity of their theoretical equilibria, especially in their allocative efficiency. In simple double auctions without uncertainty or information asymmetry, theoretical equilibria are attainable with individuals endowed with only minimal levels of intelligence (not trading at a loss). Jamal and Sunder (1996) extended the results to markets with shared uncertainty with algorithmic agents using means-end heuristic (henceforth M-E), developed by Newell and Simon (1972). In the current paper, we examine whether the ZI results reported in the literature also hold in more complex rational expectation markets where there is state uncertainty, variability in the number of states, and differently informed traders. The empirical results from PS experiments with human subjects form a benchmark for comparison with those of our algorithmic traders.

Substitution of human subjects used in traditional laboratory markets by algorithmic agents using M-E heuristic has the advantage of helping us gain precise control of traders’ information processing and decision making (i.e., “cognitive”) abilities. Holding trader “cognition” constant at a specified level allows us to explore the properties of outcomes of market structures and environment (also, see Angerer et al. 2010; Huber et al. 2010). In contrast, we can neither observe nor hold invariant the strategies used by human traders. The use of algorithmic traders enables us to run longer computational experiments, randomize parameters in the experimental setting, and conduct replications without significantly more time or money.

The paper is organized in five sections. The third section describes a simple M-E heuristic used by minimally-intelligent algorithmic traders in a double auction market. In the fourth section, we implement this heuristic in a market where some traders have perfect insider information (while others have no information) and compare the simulation results with the data from the profit-motivated human experiments reported by PS. The fifth section presents implications of the findings and some concluding remarks.

3 Simple Agents, Market Environment and Experimental Design

Simon (1955) proposed bounded rationality as a process model to understand and explain how humans, with their limited knowledge and computational capacity, behave in complex settings. He postulated that humans develop and use simple heuristics to seek and attain merely satisfactory, not optimal, outcomes. To understand human problem-solving Newell and Simon (1972) developed General Problem Solver (GPS). They adduced a large body of data which show that GPS is a robust model of human problem-solving in a wide variety of tasks and environments. The key heuristic used by GPS is means-ends analysis (M-E or the heuristic of reducing differences). Gigerenzer and Todd (1999) have focused on the usefulness and effectiveness of fast and frugal heuristics like M-E in human life, whereas Tversky and Kahneman (1974) have documented a similar heuristic which they labeled anchor-and-adjust.

GPS recognizes knowledge states, differences between knowledge states, operators, goals, sub-goals and problem solving heuristics as entities. GPS starts with an initial (or current) knowledge state and a goal or desired knowledge state. GPS then selects and applies operators that reduce the difference between the current state and the goal state. The M-E heuristic for carrying out this procedure can be summarized in four steps: (i) compare the current knowledge state a with a goal state b to identify difference d between them; (ii) find an operator o that will reduce the difference d in the next step; (iii) apply the operator o to the current knowledge state a to produce a new current knowledge state a' that is closer to b than a; and (iv) repeat this process until the current knowledge state a' is acceptably close to the goal state b. Knowledge states of traders can be represented as aspiration levels that adjust in response to experience (Simon 1956). The M-E heuristic for a trader thus requires a mechanism for setting an initial aspiration level, and a method for adjusting these levels in light of experience (e.g., Jamal and Sunder 1996). In Appendix A we outline the algorithm used by our program.

3.1 Market Environment

Markets examined here are defined by four elements: (i) uncertainty, (ii) distribution of information, (iii) security payoffs, and (iv) rules of the market. Following PS, we examine markets for securities with either two (X and Y) or three (X, Y, and Z) states of the world, where each state S_i occurs with a known probability π_i. One half of the traders in the markets ($n = 6$) are informed about the realized state before trading starts each period, while the other half ($n = 6$) are uninformed. At the beginning of each period, each trader of type j *(j = 3 types in our experiment)* is endowed with two units of a security which pays a single state-contingent dividend D_{Sj} at the end of the trading period. There are no cash constraints. There are three types of traders and each trader type gets a different dividend in a given state. The rules of the double auction are as follows: after a bid or ask is generated (see Sect. 3.2 and Appendix A for details on algorithm for generating bids and asks), the highest bid price is compared to the lowest ask price. If the bid price is equal to or greater than the ask price, a trade occurs. The recorded transaction price is set to be equal to the midpoint between the bid and ask prices.

3.2 Implementing the M-E Heuristic

In the first of the two implementation steps, each agent's initial knowledge state (aspiration level) is set equal to the expected value of the payoff based on its private information.[2] The second step implements the idea that subjects without perfect information make gradual adjustments by applying weight $\gamma (0 \leq \gamma \leq 1)$ to the newest observed price P_t, and weight $(1 - \gamma)$ to their Current Aspiration Level (CAL_t). This

[2] See Appendix A for more details.

process can be represented as a first-order adaptive process:

$$CAL_{t+1} = (1 - \gamma)\, CAL_t + \gamma\, P_t. \tag{1}$$

If CAL_0 is the initial value of CAL_t, by substitution,

$$CAL_{t+1} = (1-\gamma)^{t+1} CAL_0 + \gamma((1-\gamma)^t P_1 + (1-\gamma)^{t-1} P_2 + \ldots + (1-\gamma)\, P_{t-1} + P_t). \tag{2}$$

In the context of markets organized as double auctions (where both buyers and sellers can actively propose prices to transact at), these two elements of the M-E heuristic—setting an initial aspiration level and gradually adapting it in light of observed transaction prices—constitute the entire heuristic activity of the agent.[3]

3.2.1 Minimally Intelligent Algorithmic Agents

Algorithmic agents use their "current aspiration level" (*CAL*) to implement a ZI strategy after Gode and Sunder (1993); they bid randomly chosen prices below, and ask randomly chosen prices above, their aspiration levels. Traders draw a uniformly distributed random number between 0 and an upper limit of 1. If the number drawn is less than or equal to 0.5, the trader generates a bid; if the number drawn is greater than 0.5, the trader generates an ask. The bid amount is determined by drawing a second random number between a lower bound of 0 and an upper bound of the individual trader's *CAL*. If this bid exceeds the current high bid, it becomes the new high bid. Correspondingly, if the action is an ask, its amount is determined by generating a second random number in the range between the lower bound of the trader's *CAL* and the upper bound of 1. This newly generated ask becomes the new current low ask if it is less than the existing current low ask. These random draws from uniform distributions are generated independently. The algorithmic agents are myopic, making no attempt to anticipate, backward induct, or theorize about the behavior of other traders. They simply use the knowledge of observable past market events (transaction prices) to estimate their opportunity sets, and choose randomly from these sets.

These markets are populated in equal numbers by traders of each payoff type of whom 50 % are (and 50 % are not) informed about the realized state of the world. The informed algorithmic traders begin by setting their initial *CAL* using the perfect signal they have about the realized state of the world for any given trader type j:[4]

$$\begin{aligned} &\text{If realized state } = X,\ CAL_X = D_{Xj} \\ &\text{If realized state } = Y,\ CAL_Y = D_{Yj} \end{aligned} \tag{3}$$

[3] Previous attempts to model individual human behavior has used processes similar to equation 2 (Carlson 1967; Carlson and Okeefe 1969).

[4] For 3-state markets, if realized state = Z, $CAL_Y = D_{Zj}$.

The uninformed traders of type j use their unconditional expected dividend value to set their initial *CAL* using the prior state probabilities:[5]

$$CAL_j = Pr\,(X)^*\,(D_{Xj}) + Pr\,(Y)^*\,(D_{Yj}) \tag{4}$$

Since they know the state with certainty, informed traders do not update their *CAL*s in response to observed transactions; they learn nothing about the state of the world from transaction prices.[6] Uninformed traders of every dividend type, however, update their *CALs* after each transaction using the M-E heuristic (i.e., first-order adaptive process) given in equation [1] above.

CAL updating is done with a randomly chosen value of the adaptive parameter γ for the simulation (see the Experimental Design below). Submission of bids and asks continues with the updated *CALs* serving as constraints on the opportunity sets of traders until the next transaction occurs, and this process is repeated for 5,000 cycles each period. At the end of each period the realized state is revealed to all traders, dividends are paid to their accounts, and each trader's security endowment is refreshed for the following period. The uninformed algorithmic traders carry their end-of-period *CAL* forward and use it as the starting point in the following period.[7] Since our traders have minimal intelligence, they do not learn by observing other's behavior or make generalizations across markets. They act in a myopic way at all times to help us examine the sufficiency of using such a strategy for attaining economic equilibria.

In the following period, informed traders again get a perfect signal about the state of the world and set their $CAL = D_{Xj}$ (or D_{Yj}) depending on whether the signal received is *X or Y* (or *Z* in 3-state markets). The uninformed traders use their end-of-period *CAL* from the preceding period as CAL_0 to trade and to generate CAL_1 after the first transaction, and so on.

3.3 Experimental Design

We use the market design parameters from the PS (1982) human experiment for the present simulations (see Table 1). We ran 50 replications of four markets numbered 2, 3, 4 and 5 as reported by PS (1982) human experiment (three states in Market 5, and two in the other three markets).[8] The participants were freshly endowed with two

[5] For 3-state markets, $CAL_j = Pr\,(X)^*\,(D_{Xj}) + Pr\,(Y)^*\,(D_{Yj}) + Pr(X)^*(D_{Zj})$.

[6] The informed traders could, for example, learn that in some states market prices are higher than their own dividend in that state, and thus raise their *CAL* to that higher level. Human traders, presumably, make this adjustment but our algorithmic traders are not allowed to make such adjustments. We should not, therefore, expect the markets with these minimally-intelligent agents to behave identically to the human markets.

[7] It would have been possible for the agents to keep track of the prices associated with each realized state and use this information in subsequent periods. In the spirit of minimal intelligence, our agents do not do so, and uninformed agents simply carry forward their *CAL* from the end of one period to the beginning of the next period. The *CAL* of informed agents responds to a perfect signal about the state realized in each period and is not dependent on experience in previous periods.

[8] Plott and Sunder (1982) found that the information structure of their Market 1 was too complex for it to reach rational expectations equilibrium in less than a dozen periods. Accordingly, we have not tried to replicate that information structure and market in the present simulations.

Table 1 Simulation parameters

Market	Corresponding markets in PS (1982) experiment	State	Probability	Dividends for each trader type			RE predictions price (allocation to)[a]	PI predictions price (allocation to)[a]
				Type I	Type II	Type III		
2	Plott and Sunder (1982), Market 2	X	0.333	0.1	0.2	0.24	0.24(III)	0.266(I_u)
		Y	0.667	0.35	0.3	0.175	0.35(I)	0.35(I_i)
3	Plott and Sunder (1982), Market 3	X	0.4	0.4	0.3	0.125	0.4(I)	0.4(I_i)
		Y	0.6	0.1	0.15	0.175	0.175(III)	0.22(I_u)
4	Plott and Sunder (1982), Market 4	X	0.4	0.375	0.275	0.1	0.375(I)	0.375(I_i)
		Y	0.6	0.1	0.15	0.175	0.175(III)	0.21(I_u)
5	Plott and Sunder (1982), Market 5	X	0.35	0.12	0.155	0.18	0.18(III)	0.212(I_u)
		Y	0.25	0.17	0.245	0.1	0.245(II)	0.245(II_i)
		Z	0.4	0.32	0.135	0.16	0.32(I)	0.32(I_i)

Plott and Sunder (1982) conducted an experiment with profit oriented human traders (half informed about the state, and half uninformed) to ascertain whether they traded at prices (and quantities) predicted by rational expectations models. Table 1 shows the parameters used in the experiment and the predictions about price and which trader type should hold securities in these markets. Our simulation uses the same parameters as those used in the PS experiment. Traders have two tokens each available for trade, and no cash constraints.

[a] Allocation code: I, II, and III for all traders of types I, II, and III respectively. I_i for informed traders of type I, I_u for uninformed traders of type I, and similarly for informed and uninformed traders of types II and III

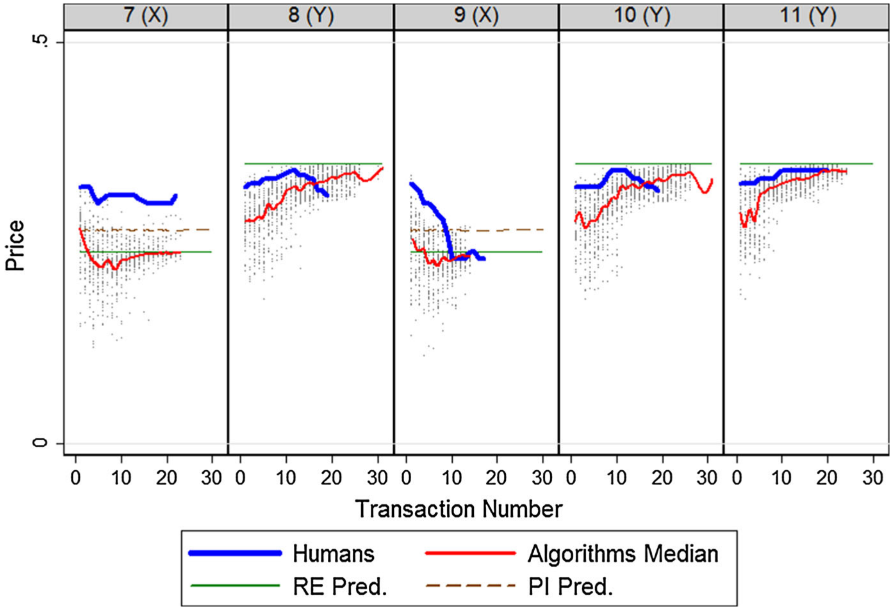

Fig. 1 The price paths in Market 2 of Plott and Sunder (1982) for periods where participants have asymmetric information (heavy *blue line* for mean price in markets with human traders; medium *red line* for median of 50 replications of simulated markets with algorithmic traders). Each *black dot* in the "cloud" is an observed transaction price in the simulated markets plotted by transaction sequence number. The *green straight line* and the *brown broken line* depict the rational expectation (RE) and prior information (PI) predicted equilibrium prices for the respective periods (the two prices are identical under State Y). *Source*: Time Chart of Prices in a 2 State Market Versus Human Trader Data from Market 2 of Plott and Sunder (1982)

securities every period and have no cash constraint. For each of the 50 replications, the adjustment parameter γ was randomly and independently drawn from a uniform distribution $U(0.05, 0.15)$.[9] In each market, there are 12 traders who traded single-period securities. A random state of nature—X, Y, (or Z in case of 3-states)—was drawn at the start of each period to match the actual realizations observed in the PS's markets. Except for a few initial periods (when no trader was informed), and in some final periods (when all traders were informed), six of these 12 traders had perfect insider information and the other six were uninformed. For consistency and ease of reference, we identify these markets using the same numbers as used by PS.[10]

[9] These ranges have been used in previous market simulation studies (Gode and Sunder 1993, 1997; Jamal and Sunder 1996) and have no normative content per se.

[10] In this paper we only report periods where six traders in the market are informed and the other six are uninformed. We have also simulated periods where all traders were informed, or all were uninformed. The results are not qualitatively different from human participants reported in PS. Full simulation results, including all periods with informed/uninformed traders are available as supplementary material at http://dx.doi.org/10.1007/s10614-016-9582-3. This website also gives an outline of the code, and allows visitors to see the charts of market outcomes.

4 Experimental Results

4.1 Price, Volume and Efficiency

Figure 1 shows the time chart of prices observed in five asymmetric information periods of Market 2 populated with profit-motivated human traders (heavy blue curve) reported in PS against the background of two theoretical (RE—solid green horizontal line) and Walrasian (PI—dashed brown horizontal line) predictions for respective periods. The red curve of medium thickness plots the median of prices from 50 replications (shown as a cloud) of the same market with M-E heuristic algorithmic traders. The adaptive parameter γ is randomly and independently drawn each period from a uniform distribution $U(0.05, 0.15)$ and is identical across all traders. Six of the 12 traders have perfect insider information and the other six are uninformed. Allocative efficiency and trading volume are shown numerically for each period in Table 2.

Figure 1 indicates: (i) In State X(with low RE price of 0.24 in periods 7 and 9), transaction prices of both human traders (blue curve) and algorithmic traders (red curve) approach the RE equilibrium level from above. (ii) In State Y (with higher RE price of 0.35 in periods 8, 10 and 11), transaction prices of both human traders and algorithmic traders generally approach and get close to the equilibrium level from below. (iii) As shown in Table 2 for Market 2, in State X (low RE price) periods, average trading volume for human traders across the two periods is 19.5 while the average volume for algorithmic traders is 17.5. The allocative efficiency of human trader markets across the two X periods is 63.5 %, while efficiency of the simulated markets is 80.3 %. Note that allocative efficiency arises from having the appropriate number of securities acquired by the appropriate type of traders as specified by rational expectations equilibrium. Efficiency levels (below 100 %) arise when the wrong type of traders hold some of the securities. In State Y (high RE price) periods, human traders' average volume is 19.3 (vs. 23.7 for algorithmic traders) and human trader efficiency is 100 %, while algorithmic traders achieve efficiency levels of 98.7 %. The direction and volume of trading is close to the predictions of RE equilibrium.

There are also important differences between the convergence paths for human and simulated markets: the convergence of prices to RE predictions with human traders is tighter and progressively faster in later periods; algorithmic simulations exhibit little change from early to later realizations of the same State (X or Y). Efficiency results also show human subjects improving over time (when State is X), whereas markets populated with algorithmic traders show less improvement over time.

Replication of the additional 2-state markets (Markets 3 and 4) with different parameters (see Figs. 2, 3 and the two middle sections of Table 2) show essentially the same pattern of convergence except that in State Y (with low RE price) human traders have a tendency to converge quickly to the RE price, especially in later periods (not coming from above or below) whereas the paths with algorithmic traders depend on history in the previous period (because the *CAL* of the uninformed traders is carried forward from previous periods). If the previous period is State X (high RE price) the simulation converges from above; if the previous period is State Y (low RE price), the simulation converges from below the RE price. As expected, algorithmic

Table 2 Number of transactions (efficiency levels in percentages) with six informed traders by market and period

Market 2 Period (State)						7(X)	8(Y)	9(X)	10(Y)	11(Y)			Avg.(X)	Avg.(Y)		Avg. (All)
Human Data	Trans (Eff)					22(57)	19(100)	17(70)	19(100)	20(100)			19.5(63.5)	19.3(100)		19.4(85.4)
Simulation (Average of 50 Reps)	Trans (Eff)					19(78)	25(99)	16(83)	25(99)	21(98)			17.5(80.3)	23.7(98.7)		21.2(91.4)
Market 3 Period (State)		3(Y)	4(X)	5(Y)	6(Y)	7(X)	8(Y)	9(X)	10(Y)				Avg.(X)	Avg.(Y)		Avg. (All)
Human Data	Trans (Eff)	15(79)	19(100)	15(88)	14(89)	19(100)	14(98)	15(100)	15(99)				17.7(100)	14.6(90.6)		15.8(94.1)
Simulation (Average of 50 Reps)	Trans (Eff)	14(87)	25(100)	12(81)	14(87)	25(100)	12(81)	25(100)	12(80)				25.0(100)	12.8(83.3)		17.4(89.5)
Market 4 Period (State)				5(Y)	6(X)	7(Y)	8(Y)	9(X)	10(Y)	11(X)	12(Y)	13(X)	Avg.(X)	Avg.(Y)		Avg. (All)
Human Data	Trans (Eff)			17(92)	23(100)	17(95)	12(93)	20(100)	14(94)	21(100)	18(94)	21(100)	21.3(100)	15.6(93.6)		18.1(96.4)
Simulation (Average of 50 Reps)	Trans (Eff)			14(90)	25(100)	12(81)	14(88)	25(100)	12(80)	24(100)	12(81)	24(100)	24.5(100)	12.8(83.9)		18.0(94.1)
Market 5 Period (State)			4(X)	5(X)	6(Y)	7(Z)	8(Z)	9(Y)	10(Y)	11(X)	12(Y)	13(Z)	Avg.(X)	Avg.(Y)	Avg. (Z)	Avg. (All)
Human Data	Trans (Eff)		15(82)	16(94)	17(87)	20(100)	23(100)	21(100)	20(100)	18(87)	18 (100)	16(100)	16.3(87.7)	19.0(96.8)	19.7(100)	18.4(95)
Simulation (Average of 50 Reps)	Trans (Eff)		14(93)	16(95)	22(99)	23(99)	24(98)	16(87)	21(97)	13(87)	22(99)	23(99)	14.3(91.5)	20.3(95.4)	23.3(98.8)	19.4(95.3)

Plott and Sunder (1982) conducted an experiment with profit oriented human traders to ascertain whether they traded at prices (and quantities) predicted by rational expectations models. Table 2 shows the number of transactions and efficiency levels attained by human traders, as well as simulated algorithmic traders who use a simple linear heuristic to update aspiration levels. The number of transactions and efficiency of markets with simulated and human traders are qualitatively comparable across state realizations in the four markets

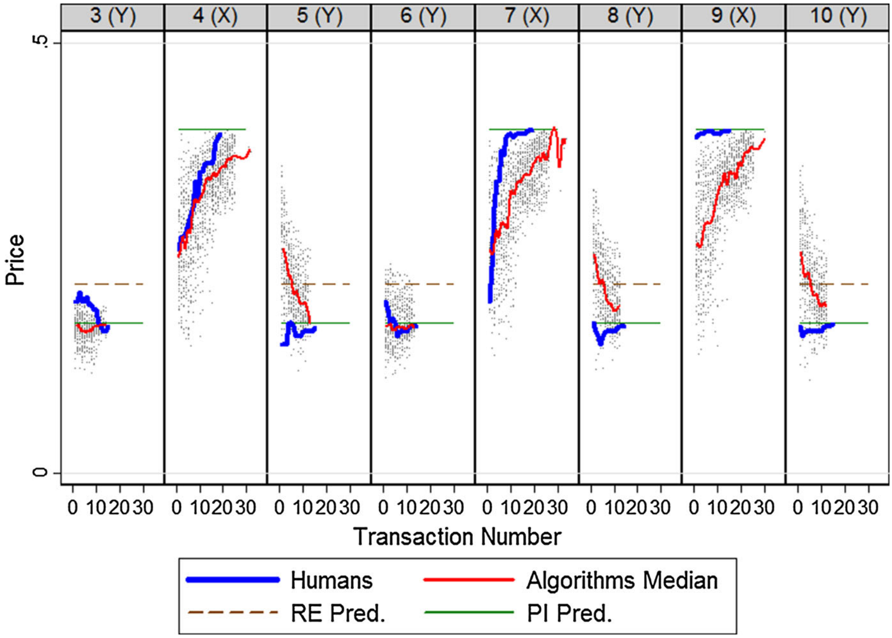

Fig. 2 The price paths in Market 3 of Plott and Sunder (1982) for periods where participants have asymmetric information (*heavy blue line* for mean price in markets with human traders; medium *red line* for median of 50 replications of simulated markets with algorithmic traders). Each *black dot* in the "cloud" is an observed transaction price in the simulated markets plotted by transaction sequence number. The *green straight line* and the *brown broken line* depict the rational expectation (RE) and prior information (PI) predicted equilibrium prices for the respective periods (the two prices are identical under State Y). *Source*: Time Chart of Prices in a 2 State Market Versus Human Trader Data from Market 3 of Plott and Sunder (1982)

traders adjust slowly and learn myopically without any global awareness of equilibrium prices.

Figure 4 displays data for a three-state market reported by PS with human traders, and an identical market replicated for this paper with algorithmic traders. The solid green horizontal line indicates the rational expectations (dashed brown line for PI) equilibrium price for the respective periods. Allocative efficiency and trading volume for Market 5 are shown numerically for each period in the bottom section of Table 2.

Figure 4 indicates: (i) In State Z (with high RE price of 0.32), for both human (blue line) and algorithmic traders (red line) transaction prices approach and get close to the RE equilibrium level from below. (ii) In State Y (with RE price of 0.245 in the middle of the other two states), transaction prices also generally approach and get close to the equilibrium level from below in both human and simulated markets. (iii) In State X transaction prices generally approach from below, the only exception occurs in Period 11 when the market converges from above in both human and simulated markets. It appears that moving from a high equilibrium price state to a lower price state may cause convergences from above. Otherwise, both humans and our simulated traders tend to approach the equilibrium price from below. (iv) Trading volume in all three states is generally greater than the predicted volume of 16 trades. For human traders

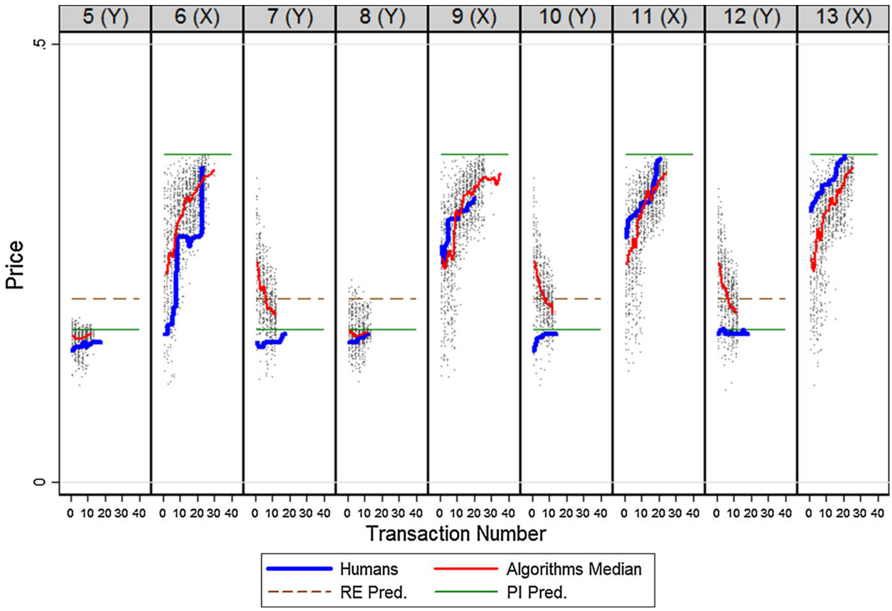

Fig. 3 The price paths in Market 4 of Plott and Sunder (1982) for periods where participants have asymmetric information (*heavy blue line* for mean price in markets with human traders; medium *red line* for median of 50 replications of simulated markets with algorithmic traders). Each *black dot* in the "cloud" is an observed transaction price in the simulated markets plotted by transaction sequence number. The *green straight line* and the *brown broken line* depict the rational expectation (RE) and prior information (PI) predicted equilibrium prices for the respective periods (the two prices are identical under State Y). *Source*: Time Chart of Prices in a 2 State Market Versus Human Trader Data from Market 4 of Plott and Sunder (1982)

volume tends to range from 15–23 trades, whereas algorithmic traders, volume ranges from 14–24 trades. (v) In all periods of State Z (high RE price), allocative efficiency for human traders is 100 % whereas algorithmic traders achieve 98.8 % efficiency. In State Y (middle RE price) periods, allocative efficiency of human traders averages 96.8 % (100 % efficiency in all periods except the first realization of State Y) whereas algorithmic traders achieve 95.4 % efficiency and do not achieve 100 % efficiency in any individual period. In State X (low RE price) periods, allocative efficiency of human traders averages 87.7 % whereas algorithmic traders achieve 91.5 % efficiency. Table 2 shows volume and efficiency numerically. Again, it is clear that outcomes of markets with profit-motivated human and minimally intelligent algorithmic traders exhibit the same central tendencies of convergence towards the predictions of rational expectations models. Apparently, the structural constraints of the market rules and Newell and Simon's (1972) simple means-end heuristics are sufficient to yield this result even as the number of states in the market increases from two states to three.

4.2 Convergence and Statistical Comparisons

To assess price convergence to the rational expectations equilibrium, we report results of a procedure used by Gode and Sunder (1993) who regressed the root mean squared

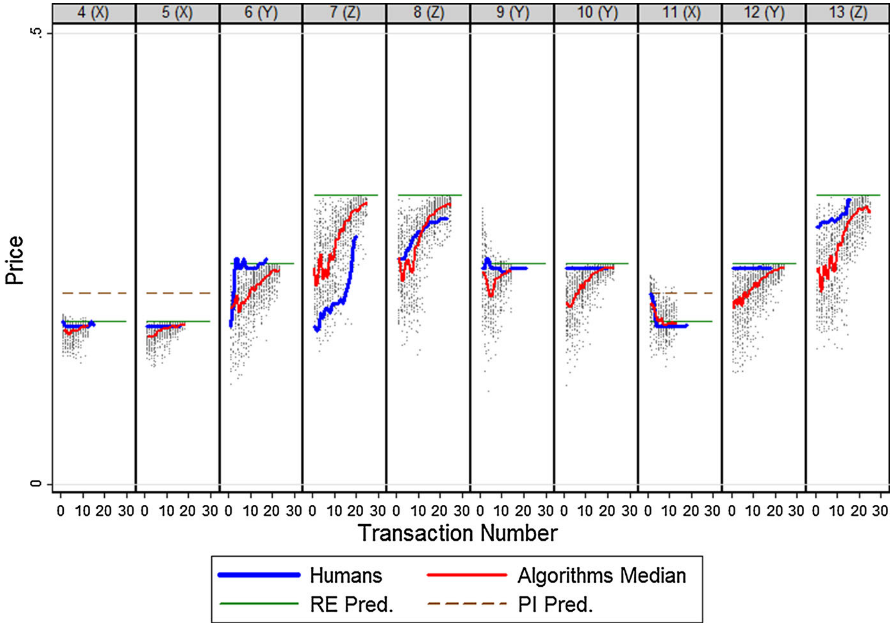

Fig. 4 The price paths in Market 5 of Plott and Sunder (1982) for periods where participants have asymmetric information (*heavy blue line* for mean price in markets with human traders; *medium red line* for median of 50 replications of simulated markets with algorithmic traders). Each black dot in the "cloud" is an observed transaction price in the simulated markets plotted by transaction sequence number. The *green straight line* and the *brown broken line* depict the rational expectation (RE) and prior information (PI) predicted equilibrium prices for the respective periods (the two prices are identical under States Y and Z). *Source:* Time Chart of Prices in a 3 State Market Versus Human Trader Data from Market 5 of Plott and Sunder (1982)

deviation between transaction and RE equilibrium prices on the natural logarithm of the transaction sequence number within a period. If prices move towards RE levels over time, the slope coefficient of this regression should be less than zero. Four panels of Fig. 5 show the behavior of this root mean square deviation over time for the four human and simulated market pairs. Results of ordinary least squares regressions of MSD on log of transaction sequence number in human and simulated markets are shown in two triplets in each panel (slope, p-value, and R^2)[11] respectively. Three of the four human (with the exception of Market 2), as well as all four simulated markets exhibit significant convergence to RE equilibrium, and the zero-slope hypothesis is rejected in favor of negative slope alternative at $p < 0.000$ for the seven of the eight (human and simulated) markets. About 80 % of the reduction in the deviation from RE equilibria is explained by the log of transaction sequence number. Figure 5 shows that root mean squared deviation of transaction from RE equilibrium prices tends towards zero.

Across all 32 periods of the four markets, the difference between the trading volume and efficiency (Table 2; charted in Figs. 6, 7) of human and simulated markets is not

[11] We report results using the same format as Plott and Sunder (1982) so our simulation results can be compared with the human experiment results.

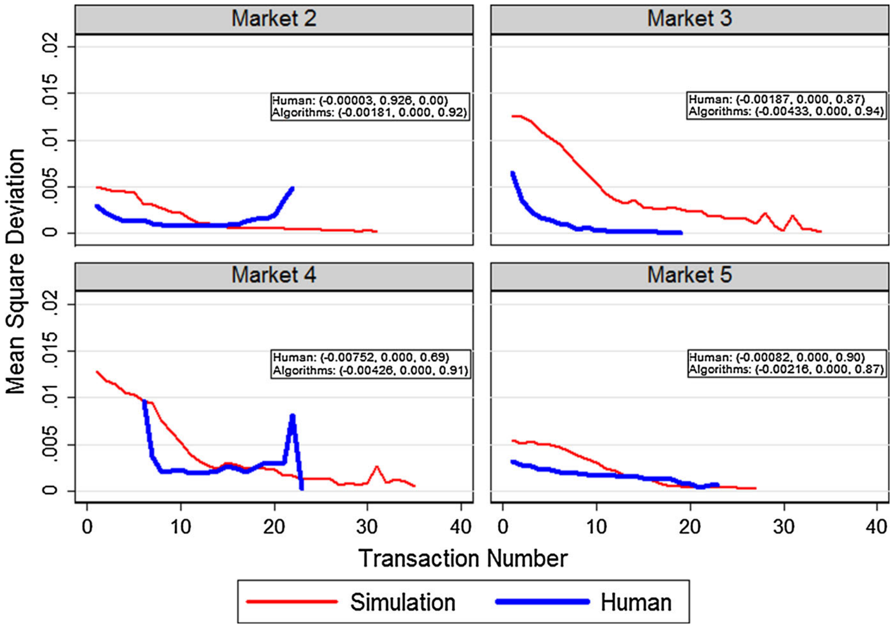

Fig. 5 The progression of mean squared deviation of observed prices from RE equilibrium prices with respect to transaction sequence numbers (*heavy blue line* for price in markets with human traders; *medium red line* for algorithmic traders). In human Market 4, the first five root mean squared deviations exceed 0.02 (for a maximum of 0.145 for transaction 3), and are out-of-scale chosen for the y-axis. Ordinary Least Squares regression ($MSD = \alpha + \beta$ *log Transaction No.*) estimates of β, p-value and R^2 for human and algorithmic markets are shown numerically in boxes inside each chart (e.g., in market 5: $\beta = -0.00082$, p-value = 0.001 and R^2 = 0.90 for human markets). *Source*: Mean Squared Deviation of Observed Prices from RE Equilibrium Prices

statistically different [average volume of simulated market is about one trade greater than for human markets with t-statistic of 1.35 and the average efficiency of simulated markets is 1.6 % lower than that of markets with human traders (t-statistic of -1.08)]. There is no significant difference between the volume and efficiency of markets with human traders as opposed to the median of algorithmic traders. The inference is not that these simple algorithms capture all or even most of the behavior of the humans; that is not true. However, when seen through the perspective of aggregate market outcomes—prices, allocations, trading volume, and efficiency—in their central tendency, these simple heuristics noisily mimic the human subject convergence to RE equilibria in these markets.

4.3 Sensitivity Analysis

We conduct sensitivity analysis to examine outcomes with a varied number of informed traders for markets 2, 3 and 4.[12] In each market, we provided information to one, two,

[12] Due to the structure of Market 5 we are not able to decrease the number of informed traders.

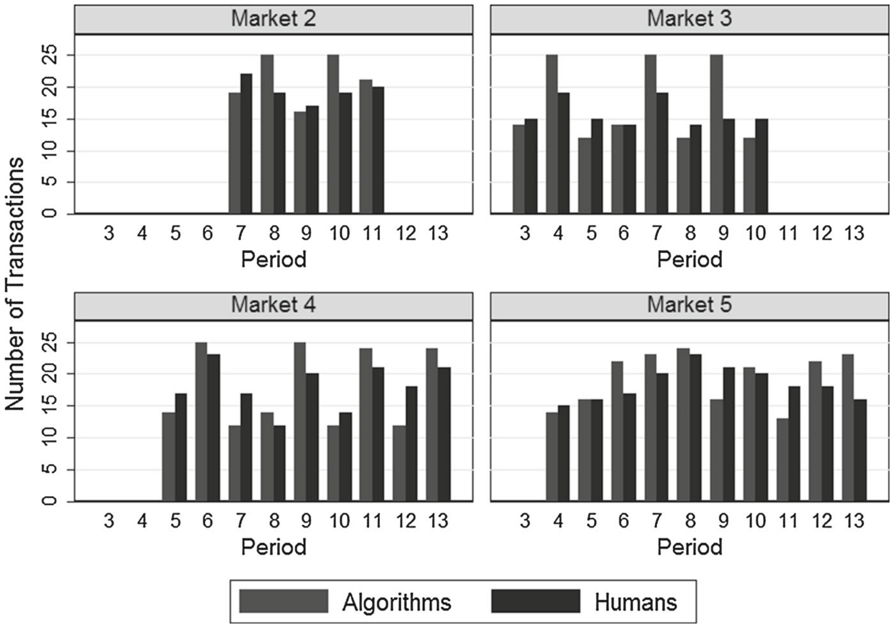

Fig. 6 Average number of transactions for algorithm traders versus human traders of Plott and Sunder (1982)

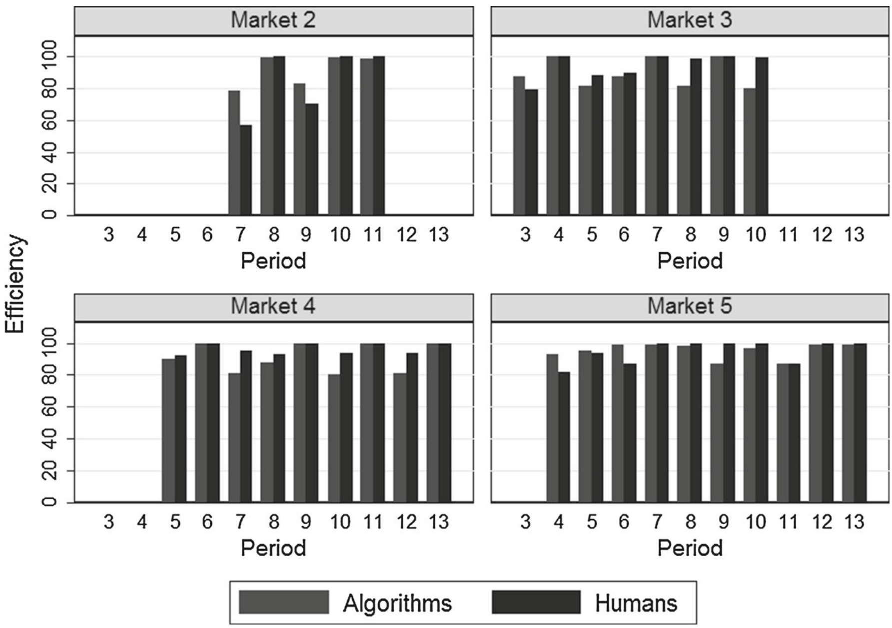

Fig. 7 Average efficiency of algorithm traders vs. human traders of Plott and Sunder (1982)

or three traders of each type. Since there are three types of traders in each market, the minimum number of informed traders is three. All of the remaining parameters were remained unchanged. We ran these simulations twice, once with 5000 iterations and again with 10,000 iterations. This was done to see if the number of iterations was a limiting factor. Table 3 provides the results of the sensitivity analysis.

For the 5000 iteration simulation runs, efficiency levels reported in Table 3 range from 80 to 88 % whereas the comparable efficiency levels with 6 informed traders in Table 2 range from 89.5 to 95 %. The average efficiency levels drop by about 7 % when the number of informed traders of each type is reduced from two to one. Increasing the number of iterations to 10,000 as reported in Table 3 yields an efficiency range of 82 to 89 %; there is not much improvement obtained by increasing the number of iterations available to trade.

We also conducted a simulation in which we increased the number of each type of informed trader to three (for a total of nine informed traders out of a total of 12 traders). Efficiency levels with nine informed traders range from 89.5 to 94 % which is essentially the same as the range obtained with six informed traders (89.5–95 % in Table 2).[13] These results suggest that the presence of even very few informed traders (one of each type in our case) may be sufficient for this market to approach rational expectations equilibria. Additional increases in the number of informed traders (from one informed trader of each type to two) improves market performance a little; however, gains from increasing the number of informed traders flatten out quickly and there is little further improvement from increasing the number of informed traders of each type from two or three.

We note that in the high equilibrium price state (Y in Market 2 and X in Markets 3 and 4; see Table 1), each market achieves close to 100 % efficiency with both human and algorithmic traders (with 3, 6 or 9 insiders—see Fig. 8). We conjecture that in the high price state, informed traders are buyers who have no budget constraint so they can keep bidding up the price until all feasible trades have occurred. In the low equilibrium price state, both human and algorithmic traders have lower efficiency levels, generally close to 80 % on average; see Fig. 8). We conjecture that these lower efficiency levels occur due to the restriction on short-selling in our simulations, particularly in the low-priced state when the informed traders are sellers rather than buyers. Since there are only three informed traders in each market, this means that there are a total of six tokens held by informed traders. Once the informed traders have sold all their tokens, there are generally no further trades available since the CALs of the uninformed traders are usually higher than the prior transaction price. As a result, the informed traders cannot take advantage of this price discrepancy and drive the market price towards the RE equilibrium.

5 Discussion and Concluding Remarks

We have presented evidence that individual behavior, modeled by simple means-end heuristics and minimal-intelligence, is sufficient to yield market-level outcomes

[13] Table 4 for the results with nine out of 12 informed traders is available as supplementary material at (http://dx.doi.org/10.1007/s10614-016-9582-3).

Table 3 Number of transactions (efficiency levels in percentages) by market and period for replications with only three informed traders

Market 2 Period (State)						7(X)	8(Y)	9(X)	10(Y)	11(Y)			Avg.(X)	*Avg.(Y)*	*Avg. (All)*	*Avg. 6 Inf. Traders*
5000 Iterations (Average of 50 Reps)	Trans (Eff)					19(46)	23(97)	8(71)	20(95)	15(91)			13.4(58.5)	19.1(94.3)	17.0(80.0)	21.2(91.4)
10,000 Iterations (Average of 50 Reps)	Trans (Eff)					21(43)	30(100)	8(71)	26(100)	20(95)			14.5(57)	25.3(98.3)	21.0(81.7)	
Market 3 Period (State)		3(Y)	4(X)	5(Y)	6(Y)	7(X)	8(Y)	9(X)	10(Y)				Avg.(X)	Avg.(Y)	Avg. (All)	Avg. 6 Inf. Traders
5000 Iterations (Average of 50 Reps)	Trans (Eff)	9(74)	32(100)	6(80)	6(80)	27(99)	6(81)	17(97)	6(81)				25.3(99)	6.6(79.2)	13.6(86.5)	17.4(89.5)
10,000 Iterations (Average of 50 Reps)	Trans (Eff)	9(75)	33(100)	6(80)	6(80)	28(100)	6(80)	25(100)	6(81)				25.0(100)	12.8(79.2)	17.4(87.0)	
Market 4 Period (State)				5(Y)	6(X)	7(Y)	8(Y)	9(X)	10(Y)	11(X)	12(Y)	13(X)	Avg.(X)	Avg.(Y)	Avg. (All)	Avg. 6 Inf. Traders
5000 Iterations (Average of 50 Reps)	Trans (Eff)			7(79)	32(100)	6(80)	6(81)	27(99)	6(80)	23(96)	6(81)	22(96)	26.0(97.8)	6.2(80.2)	14.8(88.0)	18.0(94.1)
10,000 Iterations (Average of 50 Reps)	Trans (Eff)			8(78)	32(100)	6(80)	6(81)	26(100)	6(81)	25(100)	6(80)	25(100)	27.0(100)	6.4(80.0)	15.4(88.9)	

Table 3 shows the number of transactions and efficiency levels for simulated algorithmic traders with three informed traders and both 5000 iterations and 10,000 iterations. The final column shows the average of all sessions for the corresponding simulations with six informed traders from Table 2

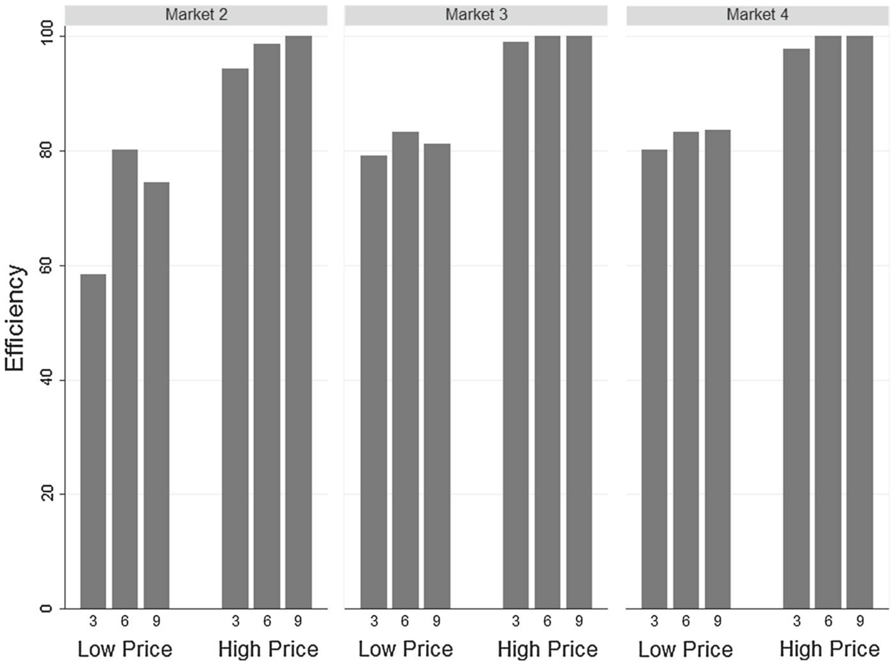

Fig. 8 Average allocative efficiency of markets with algorithmic traders by the number of informed traders in each market for high and low equilibrium price states

centered around the equilibrium levels derived from strong assumptions about optimization by individual agents. This occurs even though our algorithmic traders are unable to make even simple inferences and learn from experience to improve their current and future performance. This lack of learning preserves the spirit of Zero Intelligence (ZI) models of behavior (Gode and Sunder 1993) and makes it more difficult for our algorithmic traders to achieve the high levels of economic efficiency (and learning across periods) exhibited by human subjects in experiments.

Even if this key optimization assumption of theory were descriptively invalid, it does not necessarily undermine the validity and predictive value of the theory at the aggregate level. Our findings are consistent with Gigerenzer and Todd (1999) who built on Simon's bounded rationality paradigm by proposing that individuals use "fast and frugal" heuristics to successfully accomplish complex tasks.

The computational or other "cognitive" abilities of our algorithmic traders do not exceed, indeed are far weaker than, the documented faculties of human cognition. Yet, these simulated markets with insider trading based on asymmetric access to information converge to the close proximity of rational expectations equilibria and attain high allocative efficiency. Contrary to claims made in behavioral economics literature (Thaler 1986; Tversky and Kahneman 1974), we find that individuals using a simple means-end heuristic [analogous to Tversky and Kahneman's (1974) anchor–and–adjust heuristic] in a market setting generate outcomes close to the rational expectations equilibrium. We interpret the results to suggest that, even in these relatively more complex market environments [as compared to Gode and Sunder (1993,

1997) and Jamal and Sunder (1996)], allocative efficiency of markets remains largely a function of their structure, not intelligence, learning or optimizing behavior of agents. Attention to understanding the role of market structure, not just human cognition, may help advance our understanding of links between economic theory and market outcomes.

Acknowledgments Jamal and Maier thank the Social Sciences and Humanities Research Council of Canada (SSHRC) for their financial support. Sunder thanks Yale University for financial support.

Appendix A

Outline of the Trading Algorithm

Variable Descriptions

BIDPRICE as a real number; stores the highest market bid price
ASKPRICE as a real number; stores the lowest market ask price
CURRENTBIDDERID as an integer; stores the index number of the highest bidder
CURRENTSELLERID as an integer; stores the index number of the lowest seller
STATE as an integer; stores the current state of the economy
GAMMA as a real number; stores the CAL adjustment parameter for the current replication
TRADER(n) as array of TRADER type
STATEPROB(s) as array of real numbers; an array containing the objective probabilities of each state

Type TRADER has the following properties:

CAL as real number; holds the Current Aspiration Level of the trader
INFORMED as a Boolean; determines whether the trader is informed or uninformed
TOKENS as an integer; records the number of tokens a trader has
CASH as a real number; records the amount of cash a trader has
DIV(s) as an array of real numbers; the dividend payable for STATE s

Step 1. Repeat Steps 2 to 12 50 times (number of replications)
Step 2. Generate a random number (R) between 0 and 100 to select a state
IF in 2 state version of program go to Step 2a
IF in 3 state version of program go to Step 2b
Step 2a. If R $<$ STATEPROB(1) then STATE $=$ 1 ELSE STATE $=$ 2 and go to Step 3
Step 2b. IF R $<$ STATEPROB(1) then STATE $=$ 1
IF R $>$ STATEPROB(1) and R $<$ STATEPROB(2) then STATE $=$ 2
IF R $>$ STATEPROB(2) then STATE $=$ 3
Go to Step 3
Step 3. For each TRADER(n)
Set TOKENS $=$ 2
If INFORMED $=$ TRUE then go to Step 3a
If INFORMED $=$ FALSE then go to Step 3b

Comment: Each trader is endowed with 2 new tokens to trade at the beginning of each period.

Step 3a. Set CAL = DIV(s) and go to Step 4
Comment: If the trader is informed then they know the true state of the world.

Step 3b. Set CAL = Σ (STATEPROB(s)*DIV(s)) if this is the first period and go to Step 4
Comment: If the trader is uninformed the starting point at the beginning of the first period is set to the expected value of the dividends. In subsequent periods the CAL carries over from period to period.

Step 4. Generate a random value GAMMA from uniform distribution (0.05,0.15) and Repeat Step 5–12 5000 times

Step 5. Select a trader at random from numbers 1 to 12 and go to Step 6

Step 6. Generate a random number (U) between 1 and 100
If the random number U is 50 or less then go to Step 7a
If the random number U is 51 or higher then go to Step 7b

Comment: This step generates whether a bid or ask is generated.

Step 7a. Generate a new random real number (T) between 0 and CAL of TRADER
If the new random number T is greater than BIDPRICE then
Set BIDPRICE = T
Set CURRENTBIDDERID = TRADERID
Go to Step 8

Comment: This step generates a random bid price for the selected trader between the lower bound of 0 and that traders CAL
If the new bid price is higher than the existing bid price in the market then the new bid price becomes the bid price in the market. The ID of the new high bidder is updated.

Step 7b. Generate a new random real number (U) between CAL of TRADER and 1
If the new random number U is less than ASKPRICE then
Set ASKPRICE = U
Set CURRENTSELLERID = TRADERID
Go to Step 8

Comment: This step generates a random ask price for the selected trader between the trader's CAL and the upper bound of 1. If the new ask price is lower than the existing ask price in the market then the new ask price becomes the ask price in the market. The ID of the new low asker is updated.

Step 8. If BIDPRICE > ASKPRICE then go to Step 9. Else return to Step 5
Comment: If the bid exceeds the ask, then a trade can occur.

Step 9. Set TRADEPRICE = (BIDPRICE – ASKPRICE) / 2 and go to Step 10
Comment: The trade price is determined as the halfway point between the bid and ask price.

Step 10. For TRADER(CURRENTBIDDERID)
Set TOKENS = TOKENS + 1
Set CASH = CASH – TRADEPRICE
Go to Step 11

Comment: The bidder has bought and this step records the purchase of the token

Step 11. For TRADER(CURRENTSELLERID)
Set TOKENS = TOKENS −1
Set CASH = CASH + TRADEPRICE
Go to Step 12
Comment: This step records the sale of the token

Step 12. For each trader with INFORMED = False
CAL = (1 − GAMMA) * CAL + GAMMA * TRADEPRICE
Comment: Carries out the step to update the players CAL in response to observed trade. Only uninformed traders update their CAL.

References

Angerer, M., Huber, J., Shubik, M., & Sunder, S. (2010). An economy with personal currency: Theory and experimental evidence. *Annals of Finance*, *6*(4), 475–509.
Becker, G. (1962). Irrational behavior and economic theory. *Journal of Political Economy*, *70*, 1–13.
Carlson, J. A., & Okeefe, T. B. (1969). Buffer stocks and reaction coefficients–Experiment with decision making under risk. *Review of Economic Studies*, *36*(4), 467–484.
Carlson, John. (1967). The stability of an experimental market with a supply-response lag. *Southern Economic Journal*, *33*(3), 305–321.
Dickhaut, J., Lin, S., Porter, D., & Smith, V. L. (2012). Commodity durability, trader specialization, and market performance. *Proceedings of the National Academy of Sciences, 109*(5), 1425–1430.
Forsythe, R., & Lundholm, R. (1990). Information aggregation in an experimental market. *Econometrica*, *58*(2), 309–347.
Forsythe, R., Palfrey, T. R., & Plott, C. R. (1982). Asset valuation in an experimental market. *Econometrica*, *50*(3), 537–567.
Gigerenzer, G., & Todd, P. (1999). *Simple heuristics that make us smart*. New York: Oxford University Press.
Gode, D. K., & Sunder, S. (1993). Allocative efficiency of markets with zero-intelligence traders: Market as a partial substitute for individual rationality. *Journal of Political Economy*, *101*(1), 119–137.
Gode, D. K., & Sunder, S. (1997). What makes markets allocationally efficient? *Quarterly Journal of Economics*, *112*(2), 603–630.
Huber, J., Shubik, M., & Sunder, S. (2010). Three minimal market institutions with human and algorithmic agents: Theory and experimental evidence. *Games and Economic Behavior*, *70*(2), 403–424.
Jamal, K., & Sunder, S. (1996). Bayesian equilibrium in double auctions populated by biased heuristic traders. *Journal of Economic Behavior & Organization*, *31*(2), 273–291.
Muth, J. F. (1961). Rational-expectations and the theory of price movements. *Econometrica*, *29*(3), 315–335.
Newell, A., & Simon, H. A. (1972). *Human Problem Solving*. Englewood Cliffs: Prentice Hall.
North, D. (1990). *Institutions, institutional change and economic performance*. New York: Cambridge University Press.
Plott, C. R., & Sunder, S. (1982). Efficiency of experimental security markets with insider information: An application of rational-expectations models. *Journal of Political Economy*, *90*(4), 663–698.
Plott, C. R., & Sunder, S. (1988). Rational-expectations and the aggregation of diverse information in laboratory security markets. *Econometrica*, *56*(5), 1085–1118.
Simon, H. A. (1955). A behavioral model of rational choice. *Quarterly Journal of Economics*, *69*(1), 99–118.
Simon, H. A. (1956). Rational choice and the structure of the environment. *Psychological Review*, *63*(2), 129–138.
Simon, H. A. (1969). *The sciences of the artifical*. Cambridge: MIT Press.
Smith, V. L. (1962). An experimental-study of competitive market behavior. *Journal of Political Economy*, *70*(2), 111–137.
Smith, V. L. (2008). *Rationality in economics: constructivist and ecological forms*. New York: Cambridge University Press.

Thaler, R. H. (1986). The psychology and economics conference handbook–Comment. *Journal of Business*, *59*(4), S279–S284.
Tversky, A., & Kahneman, D. (1974). Judgment under uncertainty–heuristics and biases. *Science*, *185*(4157), 1124–1131.
Tversky, A., & Kahneman, D. (1986). Rational choice and the framing of decisions. *Journal of Business*, *59*(4), S251–S278.