Agent-Based Models of Financial Markets: A Comparison with Experimental Markets*

Nicholas T. Chan[†], Blake LeBaron[‡], Andrew W. Lo[††], and Tomaso Poggio[‡‡]

This Draft: September 5, 1999

Abstract

We construct a computer simulation of a repeated double-auction market, designed to match those in experimental-market settings with human subjects, to model complex interactions among artificially-intelligent traders endowed with varying degrees of learning capabilities. In the course of six different experimental designs, we investigate a number of features of our agent-based model: the price efficiency of the market, the speed at which prices converge to the rational expectations equilibrium price, the dynamics of the distribution of wealth among the different types of AI-agents, trading volume, bid/ask spreads, and other aspects of market dynamics. We are able to replicate several findings of human-based experimental markets, however, we also find intriguing differences between agent-based and human-based experiments.

Keywords: Agent-Based Models; Artificial Markets; Experimental Markets; Market Microstructure

JEL Classification: G12

---

*This research is part of the MIT Artificial Markets Project and is partially supported by the Artificial Intelligence Laboratory, the Center for Biological and Computational Learning, and the Laboratory for Financial Engineering, all at MIT, and the National Science Foundation (Grant Nos. ACS–9217041 and SBR–9709976). We thank seminar participants at the 1999 Computing in Economics and Finance Conference for valuable discussions and comments.

[†]MIT, Department of Electrical Engineering and Computer Science, and Center for Biological and Computational Learning, 45 Carleton Street, Cambridge, MA 02142.

[‡]Brandeis University, Graduate School of International Economics and Finance, Mail Stop 32, Waltham, MA 02454.

[††]MIT, Sloan School of Management, Laboratory for Financial Engineering, and Center for Biological and Computational Learning, 50 Memorial Drive, E52–432, Cambridge, MA 02142.

[‡‡]MIT, Department of Brain and Cognitive Sciences, and Center for Biological and Computational Learning, 45 Carleton Street, Cambridge, MA 02142.

Contents

1 Introduction

One of the most powerful ideas of modern economics is Adam Smith's (1776) Invisible Hand, the fact that agents acting in their own self-interest can reach an optimal allocation of scarce resources. This remarkable feature of perfectly competitive economies is due, of course, to the presence of markets, exchanges where buyers and sellers trade with each other and, in doing so, establish prices and quantities that equate supply and demand. Although these ideas were developed over two centuries ago, it is only within the past two or three decades that economists have begun to explore the specific mechanisms, i.e., the *market microstructure*, by which markets aggregate and disseminate information dynamically in a world of uncertainty and asymmetric information.

In many of these investigations, the theoretical analysis quickly becomes intractable for all but the simplest stylized models, and even the existence of an equilibrium cannot be guaranteed in many cases.[1] An alternative to this theoretical approach is an experimental one in which individuals are placed in a controlled market setting, given certain endowments of securities or cash or both, and allowed to trade with each other.[2] By varying the market structure, the design of the securities that can be traded, and the individuals' endowments, rewards, and information set, we can learn a great deal about the actual behavior of economic agents in a simple competitive environment and how markets perform their resource-allocation function so efficiently. Documenting and studying the interactions of optimizing individuals in an experimental setting is an important first step towards understanding their behavior in real markets.

However, the experimental-markets approach has its own limitations. In particular, although the market structure and economic environment are controlled by the experimenter, the motives and information-processing abilities of the economic agents are not. Therefore, it is often difficult to assess the impact of risk aversion, learning abilities, and the degree of individual rationality on prices and quantities in experimental markets. Moreover, there is no simple means to determine *how* agents process information and derive their trading rules in any given experiment, hence no assurance that any single experimental result is not an

[1]See Cohen, Maier, Schwartz & Whitcomb (1986), Schwartz (1993), O'Hara (1995), and Campbell, Lo & MacKinlay (1996, Chapter 3) for overviews of the theoretical and empirical market microstructure literature.

[2]Davis & Holt (1993) and Kagel & Roth (1995) are excellent surveys of this fast-growing literature.

artefact of the particular subjects in the experiment.

In this paper, we advocate a third approach—the use of *artificially intelligent* agents—to address some of the limitations of the theoretical and experimental alternatives. AI-agents are computer programs that contain certain heuristics and computational learning algorithms, with the intention of capturing particular aspects of human behavior. Although AI-agents are figments of our (and the computer's) imagination, their preferences and learning algorithms are transparent and, unlike experimental subjects, can be carefully controlled and modified. Using AI-agents, we can conduct a far broader set of experiments involving more complexities than with human agents. Moreover, the outcomes of such experiments are often more readily compared to theoretical models because we have eliminated the human "wildcard".

This approach, now commonly known as "agent-based models",[3] allows us to explore new areas of economic theory, especially in dynamic markets with asymmetric information, learning, and uncertainty—a combination that poses many insurmountable technical challenges from a theoretical perspective. However, agent-based models also bring with them new and untested algorithms, parameters that must be calibrated, and other ad hoc assumptions that are likely to be controversial. To address these concerns, we propose using data from human experimental markets to validate and calibrate our agent-based models. In particular, we have designed our market structure along the same lines as those in the experimental-markets literature and show that *simple* AI-agents—agents endowed with only rudimentary computational learning abilities—can replicate several features of human-based experimental markets.

Specifically, we construct a double-auction market for a single stock that pays one liquidating state-contingent dividend at the end of each trading period, and we allow several types of AI-agents—each endowed with its own preferences, information, and learning algorithm—to trade with each other during repeated trading periods. In the course of six different experimental designs, we investigate a number of features of our agent-based model: the price efficiency of the market (how close market prices are to the rational expectations equilibrum (REE), the speed at which prices converge to the REE, the dynamics of the distribution of

[3]Another term that has been proposed is "agent-based computational economics" or ACE. See `http://www.econ.iastate.edu/tesfatsi/ace.htm` for further discussion.

wealth among the different types of AI-agents, trading volume, bid/ask spreads, and other aspects of market dynamics. In these experiments, we are able to replicate several findings of human-based experimental markets, e.g., the dissemination of information from informed to uniformed traders, the aggregation of information from traders with private information, and convergence to the REE price after a number of trading sessions.

However, we also find significant differences between agent-based and human-based experiments. For example, in one of our experiments in which agents have heterogeneous preferences and heterogeneous information, prices never converge to the REE; the opposite result was reported by Plott and Sunder (1982) in an experimental market with human subjects. Such differences may point to key features of human learning and inference that we have not captured in the design of our AI-agents, and are just as important as for developing a better understanding of how human markets operate as the features that we are able to replicate.

In Section 2, we provide a brief review of both the experimental and computational literatures. We describe our market environment in Section 3 and provide the details of the particular experiments we conduct. The results of those experiments are summarized in Section 5, and we conclude in Section 6.

2 Review of the Literature

Our agent-based modeling approach draws on at least three distinct literatures: the market microstructure literature, the experimental markets literature, and the simulated markets literature, and we provide brief reviews of each in Section 2.1–2.3. However, we also wish to mention the recent paper by Farmer (1999) and the evolutionary view of financial markets espoused by LeBaron (1995) and Farmer & Lo (1999). In those papers, the emphasis is on modeling the dynamic interactions among agents of several types, where behavior is not always microfounded by optimizing expected utility and markets are not always in equilibrium. Instead, heuristics are often proposed as reasonable approximations to behavior, and the implications of those heuristics—in and out of equilibrium—are developed in some detail.

In this respect, Farmer (1999) and other agent-based models owe a great intellectual debt to the work of Herbert Simon (see, for example, Simon (1982)), whose notion of "bounded rationality" is the foundation on which much of the recent behavioral economics literature is built. Moreover, many of Simon's recent contributions have been at the intersection of economics, psychology, and computer science, providing enormous stimulation and inspiration for agent-based models and other related literatures.

2.1 Market Microstructure

The market microstructure area now has a well developed and mature literature which, some have argued, has its roots in Adam Smith's (1776) *An Inquiry into the Nature and Causes of the Wealth of Nations*. This literature provides important background and context for our experiments because many of the questions and issues that we focus on are those that the market microstructure literature has considered theoretically and empirically.

Although our approach takes a decidedly different tack from the recent market microstructure literature, nevertheless, there are several important papers that provide motivation and inspiration for the agent-based models. For example, Garman (1976) developed one of the earliest models of dealership and auction markets and went so far as to deduce the statistical properties of prices by simulating the order-arrival process. Cohen, Maier, Schwartz & Whitcomb (1983) and Hakansson, Beja & Kale (1990) propose more complex simulation models of market-making activities. And a number of papers develop optimal market-making behavior in certain theoretical contexts, e.g., Amihud & Mendelson (1980), Ho & Stoll (1981), and Kyle (1985). Cohen et al. (1986), Schwartz (1993), O'Hara (1995), and Campbell et al. (1996, Chapter 3) provide excellent overviews of the market microstructure literature.

The focus of this literature is primarily the structure of markets and market-making activities. In contrast, the focus of agent-based models of financial markets is broader. In this paper, we provide enough market structure to enable our agents to trade with each other, but we also specify the preferences and learning heuristics of all market participants, and it is the interaction of these two sets of specifications that yields the rich implications that we shall describe in Sections 4 and 5.

2.2 Experimental Markets

Davis & Holt (1993) and Kagel & Roth (1995) provide excellent coverage of the recent literature in experimental markets. In much of this literature, the rational expectations (RE) model has been the main benchmark, and has had mixed success in various studies. Studies of the informational efficiency of experimental markets relative to the RE benchmark generally fall into two categories: information dissemination between fully informed agents ("insiders") and uninformed agents, and information aggregation among many partially informed agents. The former experiments investigate the common intuition that market prices reflect insider information, hence uninformed traders should be able to infer the true price from the market. The latter experiments explore the aggregation of diverse information by partially informed agents, a more challenging objective because none of the agents possesses full information (traders identify the state of nature with certainly only by pooling their private information through the process of trading).

Plott & Sunder (1982) and Forsythe, Palfrey & Plott (1982) investigate markets with insiders and uninformed traders. They show that equilibrium prices do reveal insider information after several trials of the experiments and conclude that the markets disseminate information efficiently. Furthermore, Plott & Sunder (1982) show that even in markets in which traders are paid different dividends (the same security pays one trader a dividend of 3 in state A but pays another trader a dividend of 5 in the same state, proxying for differences in preferences between the two traders), prices still converge to the REE. They attribute the success of the RE model to the fact that traders learn about the equilibrium price and the state of nature simultaneously from market conditions.

On the other hand, results by Plott & Sunder (1988) and Forsythe & Lundholm (1990) show that a market aggregates diverse information efficiently only under certain conditions: identical preferences, common knowledge of the dividend structure, and complete contingent claims. These studies provide examples of the failure of the RE model and suggest that information aggregation is a more complicated situation. In a related study, O'Brien & Srivastava (1991) find that market efficiency—defined as full information aggregation—depends on "complexity" of the market, as measured by market parameters such as the number of stocks and the number of trading periods in the market.

2.3 Simulated Markets

Computer simulations of markets populated by software agents extend the experimental approach by allowing the experimenter to test various theories of learning behavior and market microstructure in a controlled environment. Unlike human-based experiments, in which the dynamics of the subjects' behavior over many trading periods are almost never modeled explicitly, agent-based models can easily accommodate complex learning behavior, asymmetric information, heterogeneous preferences, and ad hoc heuristics.

Garman (1976), Cohen et al. (1983), and Hakansson et al. (1990) were early pioneers of agent-based models of financial markets. More recently, Gode & Sunder (1993) uses this framework to demonstrate a remarkable property of competitive markets: even in the absence of any form of learning or intelligence, agents trading randomly eventually converged to the REE as long as budget constraints were continually satisfied.

Several other authors have added varying degrees of intelligence to Gode and Sunder's "zero-intelligence" (ZI) traders by restricting the range of bids and asks that they generate. Usually these restrictions involve some function of recently observed trades or quotes. Two examples are Jamal & Sunder (1996) and Cliff & Bruten (1997); both implement simple heuristics to try to limit and improve on simple random bidding.

Additional examples of trading algorithms for the simple double auction can be found in the report on the Santa Fe Institute Double Auction Tournament by Rust, Miller & Palmer (1992). This tournament focuses on the relative performance of various strategies played against each other. One of its key findings is that a very simple "parasite" strategy that feeds off the others performs best.

Finally, more complex computer-simulated asset markets that emphasize the evolution of trading behavior over time have also been created. LeBaron (forthcoming 1999) surveys many of these computational markets.[4] These simulations attempt to capture long-range market phenomenon as well as short-range trading dynamics, and share our emphasis of building behavioral theories starting at the individual level.

[4]Examples include Routledge (1994, 1999), Arifovic (1996), Arthur, Holland, LeBaron, Palmer & Tayler (1997), Lettau (1997), Youssefmir & Huberman (1997).

3 Experimental Design

Our experimental design consists of four components: the overall market structure and economic environment, the trading mechanism, the types of traders, and the learning algorithms that each type of trader employs. We describe each of these components in Sections 3.1–3.4, respectively.

3.1 Market Structure and Economic Environment

The general structure of our simulations is a double-auction market in which AI-agents trade a single security that pays a single liquidating state-contingent dividend at the end of a *trading period* by submitting orders for the security during the trading period. Each trading period consists of 40 *trading intervals*, and although the security pays no dividends until the last interval, trading occurs and information is revealed through prices and order flow in each interval. An *epoch* is defined to be a sequence of 75 consecutive trading periods, where an independently and identically distributed (IID) draw of the state of nature and private information is realized in each period. The state of nature is IID across periods and all the traders' endowments are reset at the start of each period, but the traders become more "experienced" as they learn from one period to the next.[5] Each of the six experiments we conduct (see Section 4) consists of 100 trials of an epoch, where each epoch begins with the same initial conditions (types of traders, wealth distribution, etc.). This experimental design is summarized in Table 1.

At the start of a period, three quantities are initialized (but not necesarily revealed): (1) the state of nature; (2) the agents' endowments of cash and stock, which is identical across all agents throughout our experiments; and (3) the private information of each agent. At the end of a period, the predetermined state of nature is revealed and dividends are distributed to the shareholders.

The state of nature is random and exogenously determined, and the underlying distribution of the state is common knowledge. For simplicity, we assume it is discrete and uniform. For example, in an economy with three states, each state has probability 1/3 of occurring.

[5]This, of course, applies only to those agents endowed with learning heuristics, e.g., empirical Bayesian and nearest-neighbor traders. See Section 3.4 for details.

We denote by $D = (0, 1, 2)$ a stock that pays a dividend of 0 in state 1, 1 in state 2, and 2 in state 3.

To model traders with homogeneous preferences, we assume that a security pays the same D regardless of who holds it. In contrast, to model traders with heterogeneous preferences, we assume that a security pays a *different* vector of dividends to different holders of the security. For example, in a market with two types of agents A and B, suppose the same security pays A a dividend of $D^a = (0, 1, 2)$, but pays B a dividend of $D^b = (2, 0, 1)$. This is a convenient device for capturing the fact that A may value a payoff in a particular state of nature more highly than B (in this example, A values a payoff in state 3 twice as highly as B). In economic terms, these agent-dependent payoffs may be viewed as marginal-utility-weighted payoffs (agents with different preferences will value identical dollar-payoffs differently). Heterogeneous preferences (payoffs) will be one motivation for trade in our market.

Differences in information about the likely state of nature is the other motive for trade. Information that is available to all market participants is public information, whereas information only known to some individuals is considered private information. The support of the distribution of dividends and their unconditional probabilities are public information, but some traders receive private information about the state of the nature. Specifically, traders are categorized into three groups according to their information: *insiders* know exactly which state will occur (for example, state 2 will occur, hence $D = (-, 1, -)$), *partially informed* traders who have imperfect information about the state (for example, state 3 will not occur, hence $D = (0, 1, -)$), and *uninformed* traders who have only public information (that is, $D = (0, 1, 2)$). Insiders and partially informed traders receive their private information at the beginning of each period. The distribution of private information is *not* common knowledge.

3.2 Trading Mechanism

The trading mechanism is a simplified double-auction market. Agents can either submit a bid or ask, or accept a posted bid or ask. If there is an existing bid for the stock, any subsequent bid must be higher than the current bid to be posted. Similarly, a subsequent ask following an existing ask must be lower than the current ask to be posted. A transaction occurs when an existing bid or ask is accepted (a market order matches with a limit order), or when the bid and ask cross (in which case the transaction price is set at the middle of the bid and ask).

For each trade, we restrict the quantity traded to be one share. There are two reasons for such a substantial simplification. First, allowing variable quantities complicates the analysis considerably, creating another strategic choice for which heuristics must be developed and then analyzed. Second, because one of the goals of our paper is to determine the minimal level of intelligence required to replicate certain features of more sophisticated human markets, we wish to keep our model as simple as possible while retaining the most essential features of a securities market, e.g., prices as a medium of information dissemination and aggregation. However, we recognize the importance of quantity as a choice variable—it is intimately associated with risk aversion, for example—and we hope to extend our analysis to incorporate variable shares traded in the near future.

No borrowing or short selling is permitted, and agents must satisfy their budget constraint at all times. Recall that each trading period consists of 40 trading intervals. At the beginning of each interval, a specific ordering of all the agents is drawn at random (uniformly). Following this randomly selected ordering, each agent submits one limit or market order. We fix the number of agents to be 20 for most of the experiments,[6] hence a maximum of $20 \times 40 = 800$ transactions can occur in any given period in such cases.

3.3 Agents

In designing our agents, we follow the spirit Gode and Sunder's (1993) "zero-intelligence" (ZI) traders by using the simplest heuristics to give us a sense of the lower bound of intelligence needed to replicate various human-market phenomena. This simplicity also allows us to analyze more easily the interactions among agents and how information is disseminated and aggregated.

Specifically, all traders are assumed to be risk neutral, and they maximize their end-of-

[6]In Experiment 4.6 we hold fixed the number of traders of one type while increasing the number of traders of another type.

period expected wealth by choosing between cash and stock. Agents maximize the end-of-period expected value of their portfolios by forecasting the liquidating dividend, and then buying when market prices are low relative to their forecast and selling when market prices are high. Although we do not explicitly model the utility functions of the agents, we do allow for some basic differences in preferences by allowing the dividend payments to differ across agents (see Section 3.1). All agents submit orders according to the procedure described in Table 2 but they differ in how they determine the expected value of the stock p^*, which we call the *base price*.[7] For example, if there exists only an ask (no outstanding bid) and the agent's base price is lower than the ask price, the agent posts a bid price that is uniformly distributed on the interval $(p^* - S, p^*)$, where p^* is the base price and S is a preset maximum spread.

Agents are of three possible types, depending on how they construct their forecasts: empirical Bayesian traders, momentum traders, and nearest-neighbor traders. Empirical Bayesian traders use market information to update their beliefs about the state of the economy.[8] They form their base price using these beliefs, and attempt to buy (sell) if the base price is higher (lower) than the market price, in which case the stock is under-valued (over-valued) from their perspective. Empirical Bayesian traders continuously observe market activities, update their beliefs, and adjust their positions accordingly. They stop trading when either the market price approaches their base price, or they run out of cash or stock.

Momentum traders are simple technical analysis traders whose forecast of tomorrow's return is today's return. Specifically, if at time t the two most recent transaction prices are p_t and p_{t-1}, then a momentum trader's forecast of the next transaction price is simply $p_t \times (p_t/p_{t-1})$. These traders reinforce and magnify the ups and downs of price movements, introducing extra volatility and irrational valuations of the security which make information aggregation and dissemination more difficult.

Nearest-neighbor traders attempt to exploit any patterns in historical prices to predict

[7] This procedure is inspired by the budget constrained ZI traders of Gode & Sunder (1993). It is also closely related to the heuristic trader mechanisms of Jamal & Sunder (1996) and Cliff & Bruten (1997), both of which suggest other methods for updating floor and ceiling levels which help to constrain bid and ask ranges.

[8] We use the term "empirical Bayesian" loosely—our traders will not actually be correctly updating their priors using all available time series data since this would be too complicated. They simplify past prices using a moving average and this is used as a proxy for the complete history of observed data, which is then used to update their priors.

market prices by using a nearest-neighbor learning heuristic (see Section 3.4). If the empirical Bayesian traders are the "fundamental investors" of the market, the nearest-neighbor traders can be viewed as sophisticated "technicians". Like the momentum traders, nearest-neighbor traders ignore any information regarding dividends and their associated probabilities. But instead of following a fixed strategy, they learn and adapt to changing market conditions.

3.4 Learning Mechanism

Empirical Bayesian traders condition their beliefs on market information. Specifically, the agents want to compute the expected dividend $\mathrm{E}[D|p_0, p_1, \ldots, p_t]$. For simplicity, we only consider transaction prices and ignore other market variables such as bid/ask prices and spreads and volume. We also assume that most of the relevant information is embedded in the transaction prices of the last k trades, hence a k-period moving average of prices m_t is used to summarize market information at time t,

$$m_t \quad = \quad \frac{1}{k} \sum_{\tau=t-k+1}^{t} p_\tau \ . \tag{3.1}$$

We set $k = 10$ in our simulations. Given the series of moving-average prices $m_k, m_{k+1}, \ldots, m_t$ and the realized dividend D_i, the conditional distribution $\mathrm{P}(m|D_i)$ can be estimated empirically, and using Bayes Theorem, $\mathrm{P}(D_i|m)$ can be determined:

$$\mathrm{P}(D_i|m) \quad = \quad \frac{\mathrm{P}(m|D_i)\mathrm{P}(D_i)}{\sum_{j=1}^{N} \mathrm{P}(m|D_j)\mathrm{P}(D_j)} \tag{3.2}$$

where $\mathrm{P}(D_i)$ is the prior probability of dividend state i given by a trader's private information set, and N is the number of possible states. Consequently, for $D = (D_0, D_1, \ldots, D_n)$ and given a moving-average price m, the conditional expectation of the dividend is

$$\mathrm{E}[D|m] \quad = \quad \sum_{i=1}^{N} \mathrm{P}(D_i|m) D_i \tag{3.3}$$

This conditional expectation is taken as the base price p^* for the empirical Bayesian traders. The order submission procedure, described in Table 2, is then followed.

In the actual implementation, the empirical Bayesian traders estimate the conditional density functions by constructing histograms with series of moving-average prices. Each histogram corresponds to a dividend state. A series is appended and the corresponding histogram is updated with the new moving-average prices after each period of an experiment. By participating in more periods, the empirical Bayesian traders attain more accurate estimates of the conditional probability. Intuitively, the empirical Bayesian traders learn the state by associating relevant market conditions with the realized state. They memorize these associations in form of histograms. These histograms give a picture of how well the agents discern different states given market data.

As for the nearest-neighbor traders, instead of observing the k-period moving-average prices, in each period i they form a sequence of n-tuples from the price series: $\mathbf{x}_n^i, \mathbf{x}_{n+1}^i, \ldots, \mathbf{x}_{T_i}^i$ where:

$$\mathbf{x}_t \quad = \quad (p_{t-n+1}, p_{t-n+2}, \ldots, p_t) \qquad , \qquad t \ = \ k, k+1, \ldots, T \ , \tag{3.4}$$

p_t is the market at time t, and T_i is the number of transactions in the period. Similar to the empirical Bayesians, the nearest-neighbor traders believe that all relevant information is embedded in the prices of the last n transactions. We set $n = 5$ in our experiments. Each of the n-tuples, $\mathbf{x}_t^i$, is associated with the end-of-period REE price, or dividend D_i, depending on the state of the economy. The pairs $(\mathbf{x}_n^i, D_i), (\mathbf{x}_{n+1}^i, D_i), \ldots, (\mathbf{x}_{T_i}^i, D_i), (\mathbf{x}_n^{i+1}, D_{i+1}), \ldots$ and so on represent the "memory" of a nearest-neighbor trader. The nearest-neighbor traders predict the dividend by first observing the most recent n-tuple in the current market, $\mathbf{x}_t^j$, then finding its r nearest neighbors in terms of Euclidean distance from memory. The forecast is defined to be the mean of the associated dividends of the r nearest neighbors.

The parameter r controls the robustness of the prediction by governing the trade-off between bias and variance of the estimate. If r is too large, the bias becomes large and the estimate is inaccurate. If r is too small, the variance is high and the estimate is noisy and sensitive to individual data points. Through simple trial-and-error, we settled on $r = 10$ as the best compromise between mean-squared-error and computational speed, but no formal optimization was performed.

4 Six Experiments

We conduct six distinct experiments, each consisting of 100 trials of an epoch (recall that an epoch is comprised of 75 consecutive trading periods). The market and information structures are identical across the six experiments, but we vary the composition of traders and the diversity of preferences. These differences are described in Sections 4.1–4.6 and summarized in Table 1. In all six experiments, we assume that there are three states of nature that occur with equal probability (unconditionally), and unless indicated otherwise, all agents begin each period with 10 units of cash and 5 shares of stock.

4.1 Information Aggregation and Identical Preferences

This experiment contains 20 agents with identical preferences (hence the dividend payoff D is the same for each agent) and all agents are partially informed that one of the three states is impossible. For example, if state 1 is the state that will be realized at the end of the period, at the beginning of the period one trader is informed that state 0 will not occur, and another trader is informed that state 2 will not occur. Although none of the traders knows in advance which state will occur, collectively, the market has perfect information about which state will occur. The REE price is simply the value of D in the realized state of nature, and the dividend payoff is $D = (0, 1, 2)$.

4.2 Information Dissemination and Identical Preferences

This experiment contains 20 agents with identical preferences, but there are 10 *insiders* who know what the state of nature is, and 10 uninformed traders who have only public information, i.e., the distribution of D. The REE price is D in the realized state, and the dividend payoff is $D = (0, 1, 3)$.

4.3 Information Aggregation and Heterogenous Preferences

This experiment contains 20 agents divided into two groups of 10 according to their preferences. In the three possible states of nature, Group A receives a dividend $D^a = (0, 1, 3)$ and Group B receives $D^b = (2, 0, 1)$. All traders have private information which rules out one of the two states that will not occur. Given the state of nature, the REE price is the higher of D^a and D^b in that particular state. For example, given that state 2 will occur, the REE price is 3.

We run this experiment twice. In the first run, we set agents' endowments at the usual levels: 10 units of cash and 5 shares of stock. In the second run, we increase each agent's cash endowment to 40 units, relaxing budget constraints considerably.

4.4 Information Dissemination and Heterogenous Preferences

There are two groups of traders with diverse preferences. Group A receives dividend $D^a = (0, 1, 3)$ and group B receives $D^b = (2, 0, 1)$. There are 5 insiders and 5 uninformed traders in groups A and B, respectively. The REE price is the higher of D^a and D^b given the state.

As in Experiment 4.3, we run this experiment twice. In the first run, we set agents' endowments at the usual levels: 10 units of cash and 5 shares of stock. In the second run, we increase each agent's cash endowment to 40 units, relaxing budget constraints considerably.

4.5 Empirical Bayesian and Momentum Traders

In this experiment we test the robustness of our market's price-discovery mechanism by varying the proportion of empirical Bayesian and momentum traders in the population. The empirical Bayesian traders provide the market with information and, by their trading activities, move market prices towards the REE. The momentum traders, on the other hand, introduce a substantial amount of noise and volatility into market prices. How much noise can the market "tolerate" before the price-discovery mechanism breaks down, i.e., prices no longer converge to the REE?

To answer this question, we fix the number of empirical Bayesian traders at 20 and perform a sequence of 14 experiments in which the number of momentum traders is increased incrementally from 0 in the first run to 150 in the 14th run.[9] By maintaining the same number of empirical Bayesian traders across these 14 runs, we keep constant the amount information in the market while successively increasing the amount of noise induced by momentum traders.

4.6 Empirical Bayesian and Nearest-Neighbor Traders

In this experiment we have 15 empirical Bayesian traders and 5 nearest-neighbor traders. The nearest-neighbor traders are designed to exploit any predictability in prices, hence their trading performance is a measure of the market's weak-form efficiency. If the market is weak-form efficient, then the empirical Bayesian traders should perform at least as well as the nearest-neighbor traders (because there is nothing for the "technicians" to pick up). On the other hand, if prices contain predictable components, the nearest-neighbor traders should outperform the empirical Bayesians.

5 Results and Discussion

In all six experiments, we focus on the informational efficiency of the market, i.e., do "prices fully reflect all available information"? Specifically, we compare market prices to their REE counterpart by measuring their average absolute price-deviation:

$$\Delta_p \quad = \quad \frac{1}{T}\sum_{t=1}^{T}|p_t \ - \ D| \tag{5.1}$$

where p_t is the transaction price and D is the REE price, and by considering the rate of convergence of p_t to D over the epoch.

In addition, we investigate bid-ask spreads, trading volume, and the wealth distribution across the different types of traders. Narrowing bid-ask spreads show that prices are con-

[9] Specifically, the 14 runs correspond to experiments with 0, 5, 10, 15, 20, 25, 30, 40, 50, 60, 75, 100, 125, and 150 momentum traders, respectively.

verging, implying that buyers and sellers are reaching a common price. Diminishing volume, on the other hand, suggests that the market is approaching its equilibrium. This is either because all traders come to the same expected price and therefore have no incentives to trade, or they simply run out of cash or stock to transact further. And the difference in wealth between two types of traders provides an indication of the economic impact of the differences among the traders. For example, in the case of insiders versus uninformed traders (Experiment 4.2), the differences in wealth between the two groups provide a measure of the value of insider information. We measure this difference as $\Delta_w(i,j)$ where

$$\Delta_w(i,j) \quad \equiv \quad \frac{W_i \; - \; W_j}{W_j} \; \times \; 100 \tag{5.2}$$

and W_i and W_j are the total wealth levels of the two types of traders.

We also investigate the expectations formed by the agents by examining their empirical conditional density functions of the moving-average price given the states. This collection of conditional density functions represents the agents' beliefs formed with their prior information and updated continuously with market prices. The agents use these density functions to distinguish one state from another, hence these functions are central to understanding how the agents learn.

In experiments that have a diverse dividend structure, we define *allocative efficiency*, following Smith (1962), as the ratio between total dividends earned by all traders and the total maximum dividends that can possibly be extracted from the market. For example, 100% allocative efficiency implies that all shares are held by traders in the group that receives the highest dividend in the realized state. The REE predicts 100% allocative efficiency in that all shares will be allocated to the traders valuing them most highly.

Recall that each experiment consists of 100 trials of an epoch consisting of 75 consecutive trading periods, and a trading period contains 40 trading intervals. Because of the enormous quantity of data generated from these simulations, it is difficult to provide numerical summaries of the results. Therefore, we summarize our findings in a series of graphs (Figures 2a–7b) and discuss them in Sections 5.1–5.4.

5.1 Homogeneous Preferences

With identical preferences, the results from our simulation are similar to those in the human-based experimental markets literature, and are captured in Figures 2a and 2b which plot the transaction, bid, and ask prices and volume in selected periods of a typical epoch of Experiment 4.1.

In this experiment the convergence to the REE is apparent. Figure 2a shows market activities in the earlier periods of the market in a typical trial of the experiment. In this stage, agents are actively learning and observing, with little evidence of convergence. However, in the later periods (see Figure 2b), after agents have accumulated sufficient knowledge regarding how states and prices are related, convergence becomes more apparent.

Figure 2c plots the average price-deviations Δ_p (see (5.1)) for each of the 75 periods of the epoch, averaged over the 100 trials of the experiment. Market efficiency clearly improves substantially over the epoch. Figure 2d plots the conditional distribution of the moving-average price (see (3.1)) for each of the three states, obtained by summing up the frequency counts for m_t across the 100 trials and for each state. These histograms show that the three states are clearly distinguishable by the agents.

In Experiment 4.2, the evidence of convergence is even more compelling (see Figures 3a and 3b). In contrast to Experiment 4.1, prices converge faster in this experiment and are closer to the REE price (Figure 3c), and bid-ask spreads are smaller. There are two reasons for such a difference in the two experiments, despite the fact that both markets have approximately the same amount of information. First, in Experiment 4.1 traders must trade with each other to "pool" their information to determine the correct price, whereas in Experiment 4.2 the insiders know the correct price. Second, in the former case the distribution of information to the traders is random. For example, there may be many more traders given the information $D = (0, 1, -)$ than those given $D = (-, 1, 3)$, biasing the consensus in one direction or another.

Figure 3e plots the cross-sectional distribution of percentage wealth differences Δ_w (see (5.2)) between the informed and uninformed traders for the 75 trading periods. For each period, we compute the average wealth within the two groups, take the percentage difference, and plot the deciles of these differences over the 100 trials. Not surprisingly, insiders have a substantially higher wealth than the uninformed. The difference in their wealth represents the value of the insider information and may be an estimate of the price traders would be willing to pay if information signals were sold. Observe that the value of insider information is diminishing over the epoch as uninformed traders learn. This is consistent with results from human-based experimental markets such as Sunder (1992) in which information is sold in a sealed bid auction. In such experimental markets, traders lower their bids for information once they learn to infer the states after a few periods of experience.

5.2 Heterogeneous Preferences

In contrast to the identical-preference cases, the prices in experiments involving diverse preferences do not seem to converge to the REE price. This can be explained by the fact that our agents attempt to recover the state of nature from market information alone, and not from the preferences of other agents (which is not common knowledge), despite the fact that heterogeneity is an important feature of their world. In fact, they are not even "aware" of the possibility of differences in dividend payoffs across traders.

Figures 4a–5e summarize the results from Experiments 4.3 and 4.4. Because our agents must infer the state of nature from market prices alone, we expect the REE model to fail in both experiments. The intuition for this conjecture comes from the fact that market prices are less useful for discriminating among states of nature in the presence of heterogeneity. For example, Figure 5e plots the conditional distribution of the moving-average price in Experiment 4.4; the probability of such a realization is almost identical in states 1 and 2, making the two virtually indistinguishable. Even if agents were told which state will occur, they would still have trouble reaching a unanimous price because of the heterogeneity in their preferences.

However, the degree of market efficiency—as measured by the average absolute price-deviation and allocative efficiency—is influenced by the traders' initial cash endowments. The outcomes of two experiments, a low-cash (10 units) and a high-cash (40 units) experiment, are summarized in Figures 4c and 4d. These figures plot average absolute price-deviations and allocative efficiency, respectively, for the two experiments over the 75 periods of an epoch and averaged over 100 trials. Figure 4c shows that the standard cash endowment of 10 units does not lead to convergence; average absolute price-deviations and allocative efficiencies do not improve much over the 75 periods. However, an initial cash endowment of 40 units does yield some convergence in Experiments 4.3 and 4.4.

A concrete example will help to illustrate how the market reaches equilibrium in the high-cash case. In Experiment 4.3, type A and type B insiders will receive 3 and 1, respectively, for one share of the stock in state 3. These are their reservation prices. Agents will not buy above or sell below these prices. Between the two groups of insiders, it is only possible for type B to buy from type A. The uninformed agents, without any private information, will have a reservation price approximately equal to 1 regardless of their dividend profile.[10] Hence, we can conjecture that the transaction prices will range from 1 to 3. Note that type B insiders will bid the highest price—close to 3—and they will never sell the shares. The rest will attempt to buy or sell at roughly 1 but type B insiders will be responsible for most of the buying. Consequently supply diminishes and the price converges gradually to 2.

Not surprisingly, we also observe close to 100% allocative efficiency in the high-cash experiment as Figure 4d shows. However, the large bid-ask spreads displayed in Figure 4b imply that many traders are still interested in trading at prices far from the REE price, and there is little improvement in this spread across the periods.

Information dissemination in a market with diverse dividends (Experiment 4.4) is unsuccessful by our learning agents. This contrasts sharply with the human-based experimental markets studied by Plott and Sunder (1982), where after a few trials, insiders begin to realize that the equilibrium price can be different from what their dividend profiles imply, and they adjust their trading strategy accordingly. Uninformed human traders are also able infer the equilibrium price from market conditions. The key distinctions between these experimental markets and our simulations are human traders' knowledge of the existence of heterogeneous preferences (diverse dividend payoffs), and their ability to learn the relation between the equilibrium price and the state of nature.

In the market of diverse information and heterogeneous preferences (Experiment 4.3), the end-of-period price does not come close to the REE price. We recognize that a market with diverse information is a more difficult scenario than one with insider information. In

[10] This is approximate because their beliefs, conditioned on the market prices, can affect their estimates of the price.

similar experiments with human subjects, Plott and Sunder (1988) show that information aggregation was unsuccessful in a market with heterogeneous preferences, and they attributed the failure to the complexity involved to inferring the state from market information. In two other sets of experiments, they found that the market aggregates information efficiently by having identical dividends across all traders (as in Experiment 4.1), or by replacing the single three-state security with three state-contingent claims. In a separate study, Forsythe & Lundholm (1990) confirmed similar results and added that information aggregation can be successful if the information about the heterogeneity in dividend payoffs is made available to all traders. Nevertheless, here our empirical Bayesian traders fail to aggregate information for the same reasons as they fail to disseminate information in Experiment 4.4.

5.3 Momentum Traders

In Experiment 4.5, we add momentum traders to the market to introduce extra noise and volatility to the "signal" perceived by the partially informed empirical Bayesian traders. To quantify the effect that momentum traders have on the market, we plot in Figure 6a the average absolute price-deviations in periods 30, 40, 50, and 75, each averaged over 100 trials for each of 14 different runs of this experiment, each run corresponding to a different number of momentum traders, from 0 in run 1 to 150 in run 14 (the number of empirical Bayesian traders is fixed at 20 for all runs).[11] As expected, the absolute price-deviation curve is highest for the period-30 plot and lowest for the period-75 plot—the market becomes more efficient over time as agents learn.

Figure 6a also shows that in all four periods, the absolute price-deviations decrease initially as momentum traders are introduced, but generally increase after the number of momentum traders exceeds 5. Momentum traders add not only noise but also liquidity to the market, and with a small population of these irrational agents in the market, the empirical Bayesians manage to take advantage of the additional liquidity in making the market more efficient. However, when the number of momentum traders reaches 25 or more, the average price-deviation exceeds that of the benchmark case where no momentum traders are present.

[11] See footnote 9.

Figure 6b provides a more detailed look at the impact of momentum traders on market prices through plots of the average absolute price-deviation over three different runs of Experiment 4.5: runs with 0, 25, and 50 momentum traders (each plot is the average over 100 trials). Not surprisingly, the average absolute price-deviations increase with the number of momentum traders. The irrational trading of the momentum traders adversely affects the price convergence at early stage of the markets (roughly from periods 1 to 40). However, as the empirical Bayesian traders learn from and adapt to the strategies of the momentum traders, they are eventually able to overcome the noise from the irrational trading. From periods 65 to 75, the three markets are about equally efficient as measured by price deviation. Both the learning of the empirical Bayesian traders and the liquidity provided by the momentum traders contribute to efficiency of the markets.

Figure 6c plots the empirical conditional distributions of the moving-average prices in an experiment with 20 empirical Bayesians and 20 momentum traders. Despite the fact that these distributions have high dispersion, the three states are still distinguishable. In such circumstances, we expect the empirical Bayesians to exploit the irrational momentum traders and end up with a much higher level of end-of-period wealth. After all, the farther the price deviates from the RE price, the higher is the gain of the empirical Bayesians. However, Figure 6d shows that this intuition is not complete. Although median wealth differences between empirical Bayesian and momentum traders increase initially (from periods 1 to 5), they generally decline afterwards. The initial increase can be attributed to the empirical Bayesians' learning about the influence of momentum traders. But after some point, the market becomes more efficient, i.e., prices become more informative and closer to the REE. This is consistent with the patterns documented in Figure 6b—the initial advantages of the empirical Bayesians diminish through time as profit opportunities are bid away.

5.4 Nearest-Neighbor Traders

In the previous experiments, we have shown that the empirical Bayesian traders are successful in disseminating and aggregating information in homogeneous-preferences cases. However, we have not investigated the weak-form efficiency of these markets, i.e., how predictable are price changes? In Experiment 4.6, the empirical Bayesian traders are combined with nearest-neighbor traders, traders that attempt to uncover and exploit predictabilities in past prices. Our hypothesis is that if market prices are informationally efficient and do fully reveal all available information, then nearest-neighbor traders will perform poorly against empirical Bayesians.

Figure 7a shows the price convergence of this market. The price deviations reach the same levels as those in Experiment 4.1, converging rapidly after 50 periods.[12] Unlike momentum traders, nearest-neighbor traders do not appear to hinder the process of information aggregation.

With respect to the relative performance of the two types of traders, Figure 7b shows that in terms of median percentage wealth differences, nearest-neighbor traders outperform empirical Bayesian traders in all but the first 4 periods, implying that market prices do have some predictability to be exploited. In fact, the nearest-neighbor traders significantly outperform the empirical Bayesians roughly from period 5 to 40, after which the median wealth difference between the two groups becomes less significant.

This suggests that the predictability in prices is temporary (but more than just a few periods), and that nearest-neighbor traders learn faster than the empirical Bayesians. The first implication is consistent with our observation of the decreasing price deviations (in Figure 7a), or equivalently increasing price efficiency, from periods 1 to 40. Nearest-neighbor traders help make the market more efficient.

With respect to the second implication, the two types of traders start learning at the same time and compete with each other to discover the REE price. Evidently, the nearest-neighbor traders are able to exploit predictabilities more quickly hence they outperform empirical Bayesians initially. Eventually, empirical Bayesians are able to adapt to the strategy of the nearest-neighbor traders and more accurately infer the state of nature from market data. Consequently, as price becomes more efficient, the advantage enjoyed by the nearest-neighbor traders diminishes. However, the distribution of wealth differences (Figure 7b) show that even in the later periods, there are some realizations in which nearest-neighbor traders exhibit small gains over empirical Bayesians. These gains are not caused by price inefficiency, but are due to the fact that empirical Bayesians trade on rather inaccurate unconditional expected

[12]For this experiment, we extend the epoch to include 100 periods to ensure that prices were converging to the REE instead of cycling.

dividend at the beginning of each period.[13]

6 Conclusions

The rich implications of our agent-based model of financial markets underscore the potential for this new approach to shed light on challenging financial issues that currently cannot be addressed in any other way. Our simulation results accord well with human-based experimental market studies in many cases. Our simple AI-agents can accurately infer and aggregate diverse pieces of information in many circumstances, and they have difficulties in cases where human traders are also unable to determine the rational expectations equilibrium.

In a small number of cases our markets behave differently from human-based experimental markets. In our view, these discrepancies are just as significant as the concordances. For example, the sharp contradiction between Plott and Sunder (1982) and our experimental results in the case of information dissemination under heterogeneous preferences points to several important issues that warrant further investigation (more sophisticated learning algorithms for our agents, non-price learning and communication by human subjects, the dynamics created by heterogeneous preferences, etc.).

The use of AI-agents with simple heuristic trading rules and learning algorithms allows us to perform many new experiments that are well beyond the capabilities of experimental markets with human subjects. For example, we show that adding momentum traders to a population of empirical Bayesians has an adverse impact on market performance and the momentum traders do poorly overall. However, this effect diminishes over time as the market becomes more efficient. But in our final experiments in which nearest-neighbor traders—traders that simply trade on patterns in past prices—are added to a population of empirical Bayesians, they are relatively successful free riders, not only matching the performance of empirical Bayesians in the long run, but outperforming the Bayesians in the short run. We conjecture that this advantage comes from the nearest-neighbor traders' ability to exploit

[13] Recall from Section 3.4, empirical Bayesian traders compute their expected dividend condition on a k-period moving average price m_t. At the beginning of each period, before k prices are available, they simply trade on unconditional expected dividend.

short-term predictabilities more efficiently (that is what they are designed to do), and such predictabilities are more readily available in the early periods of trading.[14] These findings raise interesting possibilities when viewed from an evolutionary point of view. In the early periods, selective pressures favor the nearest-neighbor traders, not the empirical Bayesian. If enough free riders enter the market, then prices might fail to converge to the rational expectations price because the market will contain too many free riders hoping to learn from price patterns alone.

Agent-based models are also ideally suited to address some of the most challenging issues in market microstructure: What are the relative merits of a monopolistic marketmaker versus multiple dealers? What are the likely effects of decimalization on the bid-ask spreads and volume? Do "circuit breakers" ameliorate or exacerbate market volatility? How do we define "liquidity"? And what are the social-welfare implications of the growing number of ECN's and the corresponding fragmentation that they create? Although each of these issues has been subjected to theoretical and empirical scrutiny, the complexities of the interactions among market participants and institutional structure are so great that very few practical implications can be expected from such studies. Agent-based models provide a natural alternative, and we plan to explore these issues more fully in future research.

Our paper is a growing research program in which computer-simulated market interactions of AI-agents are yielding many insights into complex issues such as learning dynamics, the evolution of market structure, and the nature of human intelligence in an economic context. We hope to have provided a bridge between ad hoc learning models and market experiments with human subjects.[15] Placing AI-agents into a well-defined market environment imposes a certain discipline upon the experimenter to be precise about learning algorithms and behavioral heuristics. Moreover, the results of such simulations may suggest hypotheses for human subjects that can then be taken into the laboratory for empirical validation.

Future agent-based simulations need not be restricted to AI-agents. We believe that there are many interesting experiments to be performed with human and software agents combined, and these "mixed" experiments may provide new methods for exploring the nature of human cognition in economic settings. Indeed, one can imagine a financial "Turing

[14] These results are closely related to parasite strategies documented in Rust et al. (1992).

[15] Other papers that have already blazed this trail are Andreoni and Miller (1995) and Arifovic (1996).

Test" in which human traders are asked to distinguish between human and electronic counterparties based solely on their trading patterns. With the recent plethora of electronic day-trading companies and corresponding technologies, we may soon see AI-agents acting as broker/dealers for human clients, hence an agent-based modeling approach to financial markets may have practical implications as well.

References

Amihud, Y. & Mendelson, H. (1980), 'Dealership markets: Market making with uncertainty', *Journal of Financial Economics* **8**, 31–54.

Arifovic, J. (1996), 'The behavior of the exchange rate in the genetic algorithm and experimental economies', *Journal of Political Economy* **104**, 510–541.

Arthur, W. B., Holland, J., LeBaron, B., Palmer, R. & Tayler, P. (1997), Asset pricing under endogenous expectations in an artificial stock market, *in* W. B. Arthur, S. Durlauf & D. Lane, eds, 'The Economy as an Evolving Complex System II', Addison-Wesley, Reading, MA, pp. 15–44.

Campbell, J. Y., Lo, A. W. & MacKinlay, A. C. (1996), *The Econometrics of Financial Markets*, Princeton University Press, Princeton, NJ.

Cliff, D. & Bruten, J. (1997), Zero is not enough: On the lower limit of agent intelligence for continuous double auction markets, Technical Report HPL-97-141, Hewelett-Packard Laboratories, Bristol, UK.

Cohen, K. J., Maier, S. F., Schwartz, R. A. & Whitcomb, D. K. (1986), *The Microstructure of Securities Markets*, Prentice-Hall, Englewood Cliffs, New Jersey.

Cohen, K., Maier, S., Schwartz, R. & Whitcomb, D. (1983), 'A simulation model of stock exchange trading', *Simulation* **41**, 181–191.

Davis, D. & Holt, C. (1993), *Experimental Economics*, Princeton University Press, Princeton, NJ.

Farmer, J. D. (1999), Market force, ecology, and evolution, Technical report, Santa Fe Institute, Santa Fe, NM.

Farmer, J. D. & Lo, A. (1999), 'Frontiers of finance: Evolution and efficient markets', *Proceedings of the National Academy of Sciences* **96**, 9991–9992.

Forsythe, R. & Lundholm, R. (1990), 'Information aggregation in an experimental market', *Econometrica* **58**, 309–47.

Forsythe, R., Palfrey, T. R. & Plott, C. R. (1982), 'Asset valuation in an experimental market', *Econometrica* **50**, 537–567.

Garman, M. (1976), 'Market microstructure', *Journal of Financial Economics* **3**, 257–275.

Gode, D. K. & Sunder, S. (1993), 'Allocative efficiency of markets with zero intelligence traders', *Journal of Political Economy* **101**, 119–37.

Hakansson, N., Beja, A. & Kale, J. (1990), 'On the feasibility of automated market making by a programmed specialist', *Journal of Finance* **40**, 1–20.

Ho, T. & Stoll, H. (1981), 'Optimal dealer pricing under transactions and return uncertainty', *Journal of Financial Economics* **9**, 47–73.

Jamal, K. & Sunder, S. (1996), 'Bayesian equilibrium in double auctions populated by biased heuristic traders', *Journal of Economic Behavior and Organization* **31**, 273–91.

Kagel, J. & Roth, A., eds (1995), *Handbook of Experimental Economics*, Princeton University Press, Princeton, NJ.

Kyle, A. S. (1985), 'Continuous auctions and insider trading', *Econometrica* **53**, 1315–1335.

LeBaron, B. (1995), Experiments in evolutionary finance, Technical report, The University of Wisconsin - Madison, Madison, Wisconsin.

LeBaron, B. (forthcoming 1999), 'Agent based computational finance: Suggested readings and early research', *Journal of Economic Dynamics and Control* .

Lettau, M. (1997), 'Explaining the facts with adaptive agents: The case of mutual fund flows', *Journal of Economic Dynamics and Control* **21**, 1117–1148.

O'Brien, J. & Srivastava, S. (1991), 'Dynamic stock markets with multiple assets', *Journal of Finance* **46**, 1811–38.

O'Hara, M. (1995), *Market Microstructure Theory*, Blackwell Publishers, Cambridge, MA.

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**, 663–698.

Plott, C. R. & Sunder, S. (1988), 'Rational expectations and the aggregation of diverse information in laboratory settings', *Econometrica* **56**, 1085–1118.

Routledge, B. R. (1994), Artificial selection: Genetic algorithms and learning in a rational expectations model, Technical report, GSIA, Carnegie Mellon, Pittsburgh, PA.

Routledge, B. R. (forthcoming 1999), 'Adaptive learning in financial markets', *Review of Financial Studies* .

Rust, J., Miller, J. & Palmer, R. (1992), Behavior of trading automota in a computerized double auction market, *in* D. Friedman & J. Rust, eds, 'The double auction market: Institutions, Theories, and Evidence', Addison-Wesley, Reading, MA, pp. 155–198.

Schwartz, R. (1993), *Reshaping the Equity Markets*, Business One Irwin, Homewood, IL.

Simon, H. (1982), *Models of Bounded Rationality, Volumes 1 and 2*, MIT Press, Cambridge, MA.

Sunder, S. (1992), 'Market for information: Experimental edvidence', *Econometrica* **60**, 667–95.

Youssefmir, M. & Huberman, B. A. (1997), 'Clustered volatility in multiagent dynamics', *Journal of Economic Behavior and Organization* **32**, 101–118.

Experiment Number	Information	Preferences	Dividend Payoff	Endowment		Number of Traders		
				Cash	Stock	E.B.	Mom.	N.N.
1	Aggregation	Homogeneous	$(0,1,2)$	10	5	20	0	0
2	Dissemination	Homogeneous	$(0,1,2)$	10	5	20	0	0
3a	Aggregation	Heterogeneous	$(0,1,3)^a$ $(2,0,1)^b$	10	5	10 10	0	0
3b	Aggregation	Heterogeneous	$(0,1,3)^a$ $(2,0,1)^b$	40	5	10 10	0	0
4a	Dissemination	Heterogeneous	$(0,1,3)^a$ $(2,0,1)^b$	10	5	10 10	0	0
4b	Dissemination	Heterogeneous	$(0,1,3)^a$ $(2,0,1)^b$	40	5	10 10	0	0
5	Aggregation	Homogeneous	$(0,1,2)$	10	5	20	$0,\ldots,150$	0
6	Aggregation	Homogeneous	$(0,1,2)$	10	5	15	0	5

Table 1: Summary of the six experiments conducted in the Artificial Markets simulations. 'E.B.', 'Mom.', and 'N.N.' denote the number of Empirical Bayesian, momentum, and nearest-neighbor traders in each experiment. Each experiment consists of 100 statistically independent repetitions of 75 trading periods.

Scenario	Action
existing bid, existing ask	
$p^* > a$	buy at market
$p^* < b$	sell at market
$b < p^* < a$ and $a - p^* > p^* - b$	post an ask distributed $\mathrm{U}(p^*, p^* + S)$
$b < p^* < a$ and $a - p^* \leq p^* - b$	post a bid distributed $\mathrm{U}(p^* - S, p^*)$
no bid, existing ask	
$p^* > a$	buy at market
$p^* \leq a$	post a bid distributed $\mathrm{U}(p^* - S, p^*)$
existing bid, no ask	
$p^* < b$	sell at market
$p^* \geq b$	post an ask distributed $\mathrm{U}(p^*, p^* + S)$
no bid, no ask	
with probability 1/2	post an ask distributed $\mathrm{U}(p^*, p^* + S)$
with probability 1/2	post a bid distributed $\mathrm{U}(p^* - S, p^*)$

Table 2: The order-submission algorithm of AI-agents in the Artificial Markets simulations, where a denotes the best ask price, b the best bid price, p^* the agent's base price, S the maximum spread from the base price, and $\mathrm{U}(x_1, x_2)$ the uniform distribution on the open interval from x_1 to x_2.

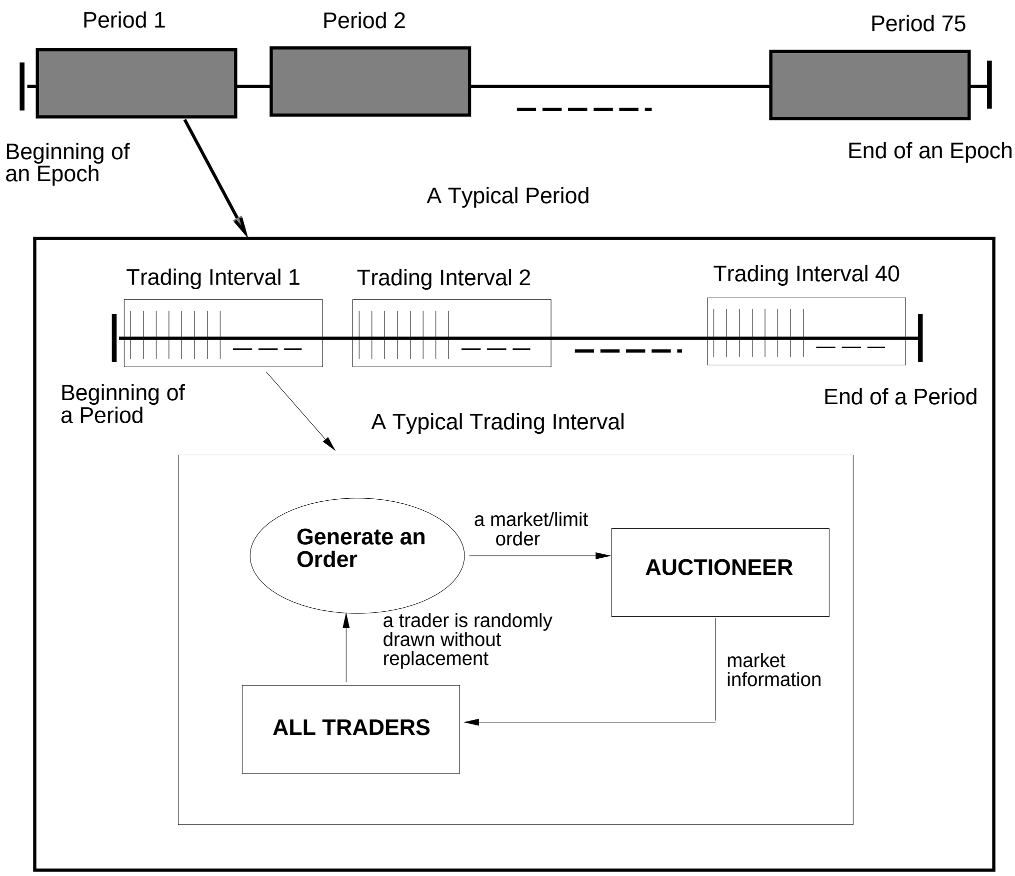

Figure 1: The experimental design of the Artificial Markets simulations. An epoch consists of 75 trading periods, and each period contains 40 trading intervals.

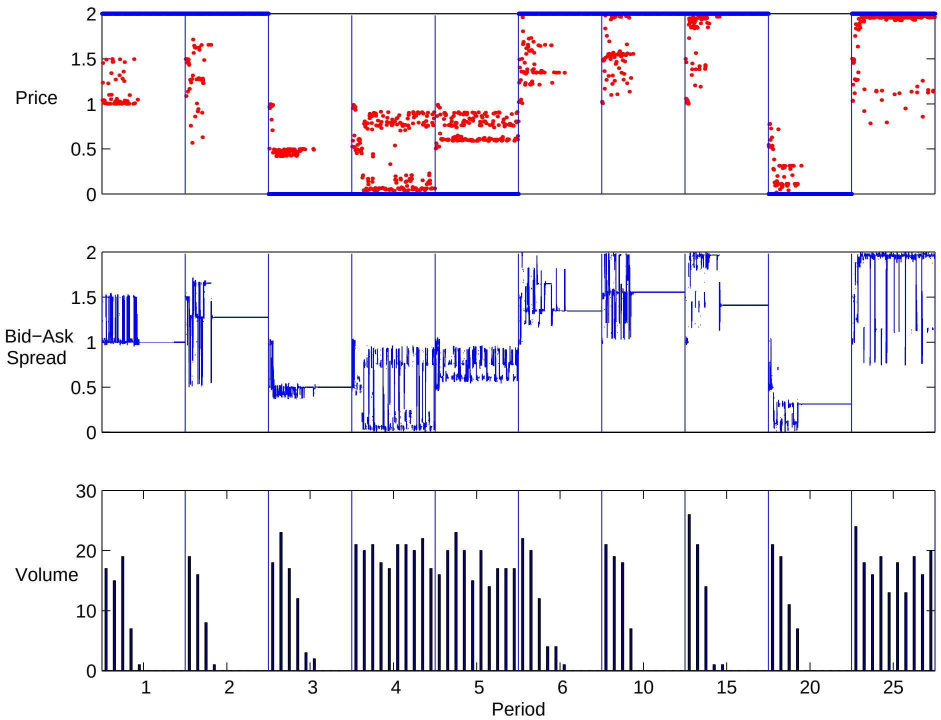

Figure 2a: Prices, bid-ask spreads, and volume in the early periods of a typical realization of Artificial Markets Experiment 4.1 (information aggregation with identical preferences).

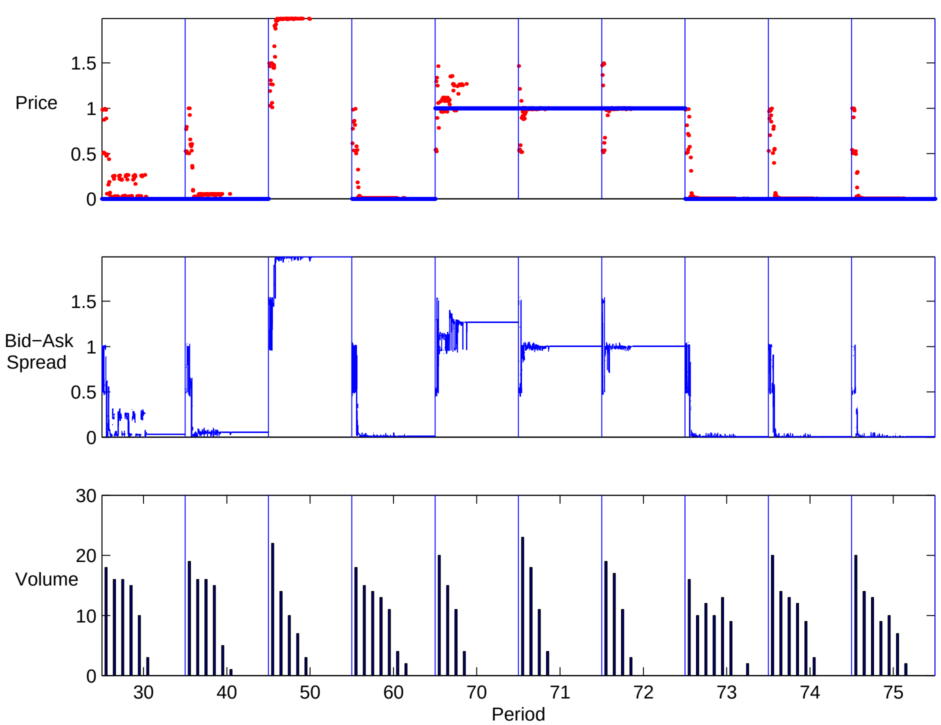

Figure 2b: Prices, bid-ask spreads, and volume in the later periods of a typical realization of Artificial Markets Experiment 4.1 (information aggregation with identical preferences).

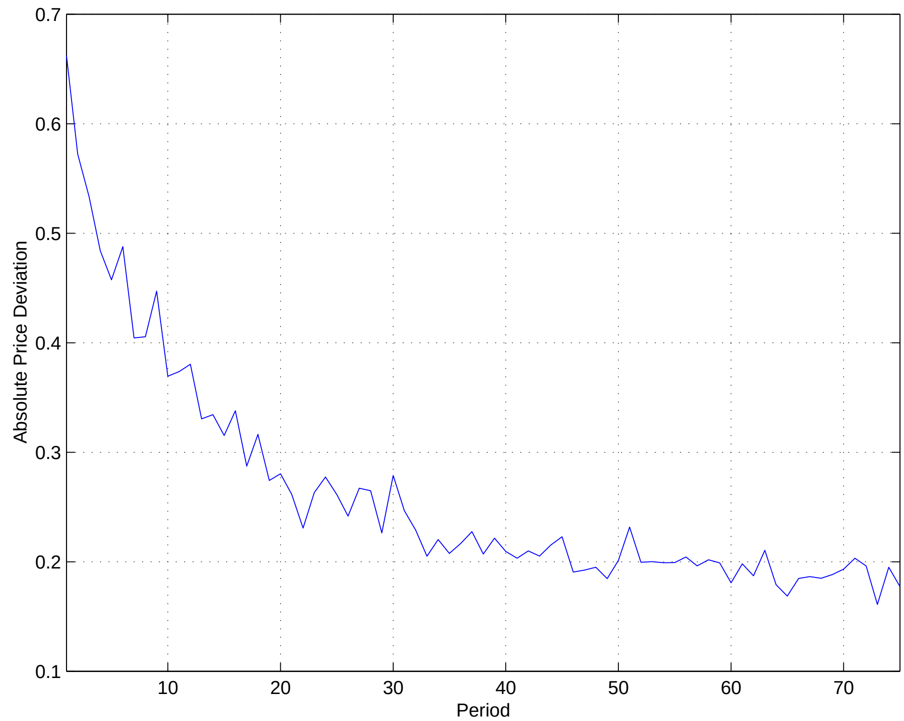

Figure 2c: Absolute price-deviations of market prices from the rational expectations equilibrium price, averaged over 100 repetitions of Artificial Markets Experiment 4.1 (information aggregation with identical preferences).

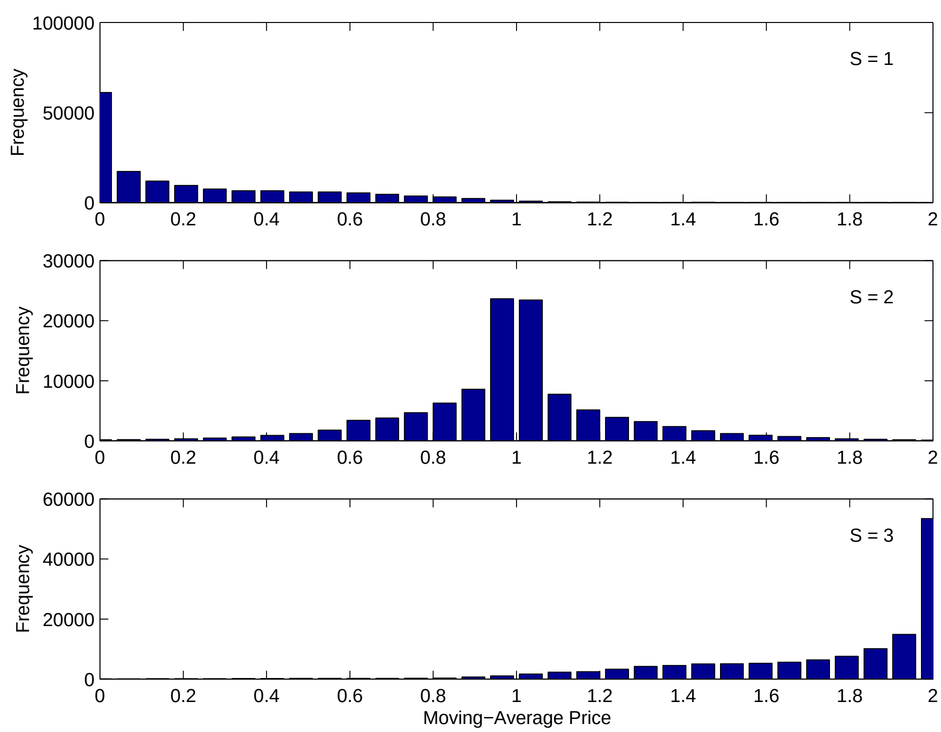

Figure 2d: Empirical distribution of moving-average prices, conditioned on the state of nature S, in Artificial Markets Experiment 4.1 (information aggregation with identical preferences).

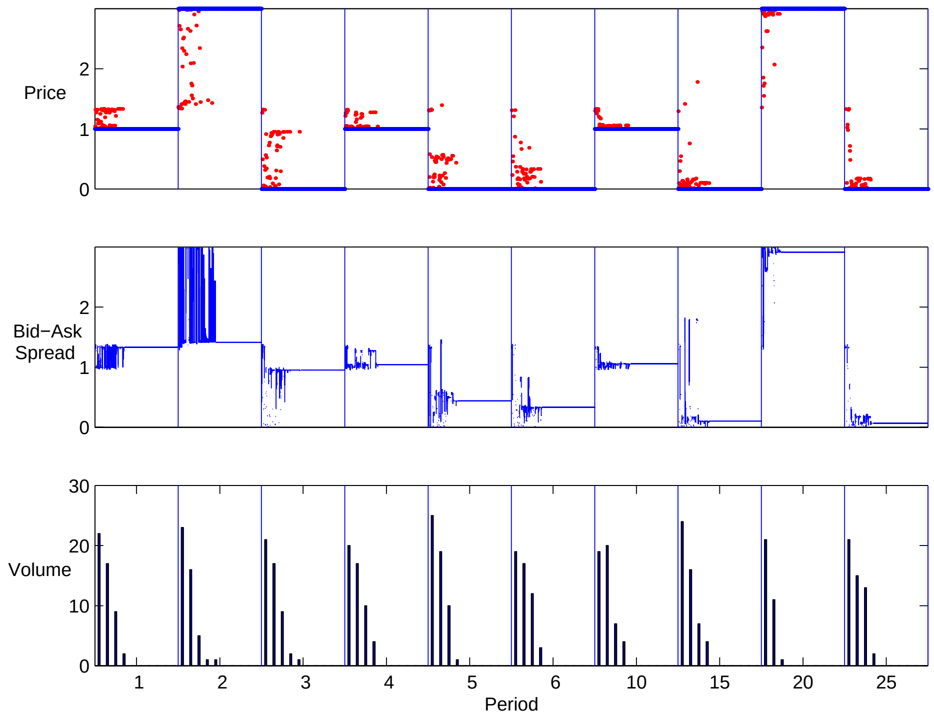

Figure 3a: Prices, bid-ask spreads, and volume in the early periods of a typical realization of Artificial Markets Experiment 4.2 (information dissemination with identical preferences).

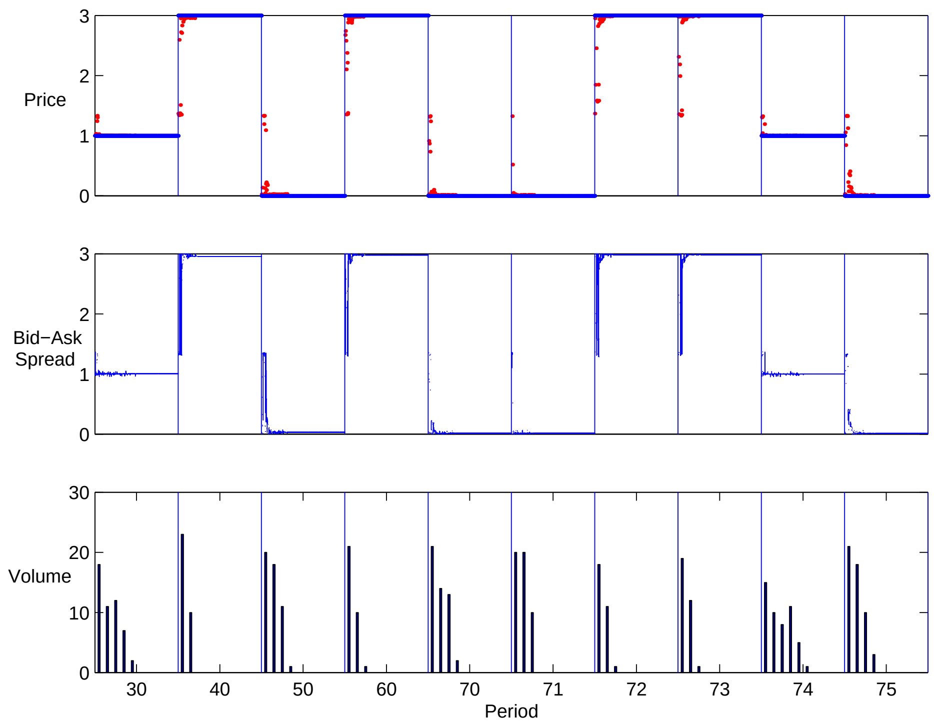

Figure 3b: Prices, bid-ask spreads, and volume in the later periods of a typical realization of Artificial Markets Experiment 4.2 (information dissemination with identical preferences).

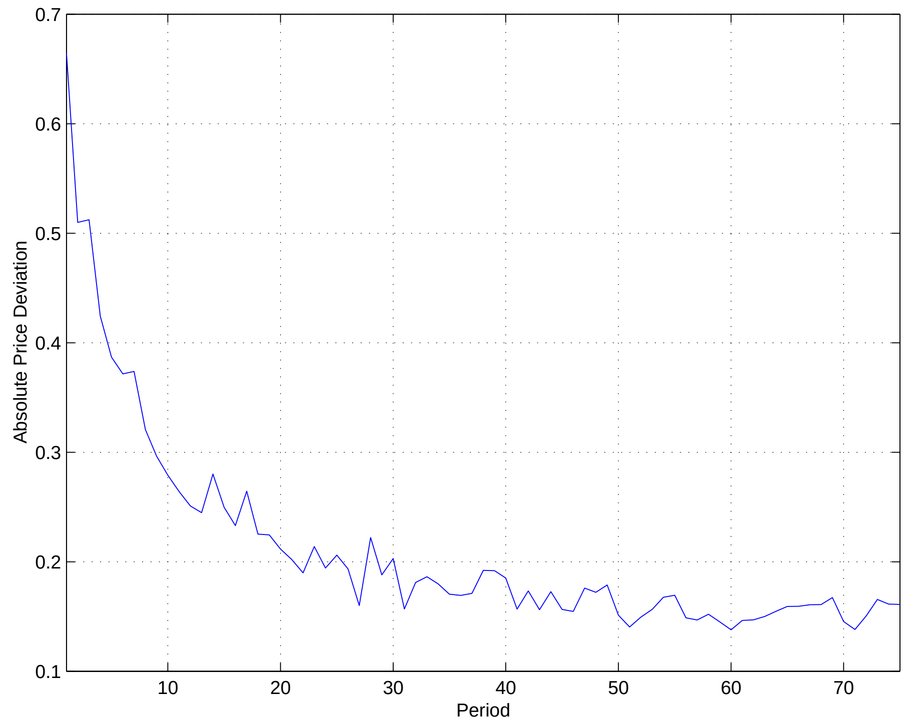

Figure 3c: Absolute price-deviations of market prices from the rational expectations equilibrium price, averaged over 100 repetitions of Artificial Markets Experiment 4.2 (information dissemination with identical preferences).

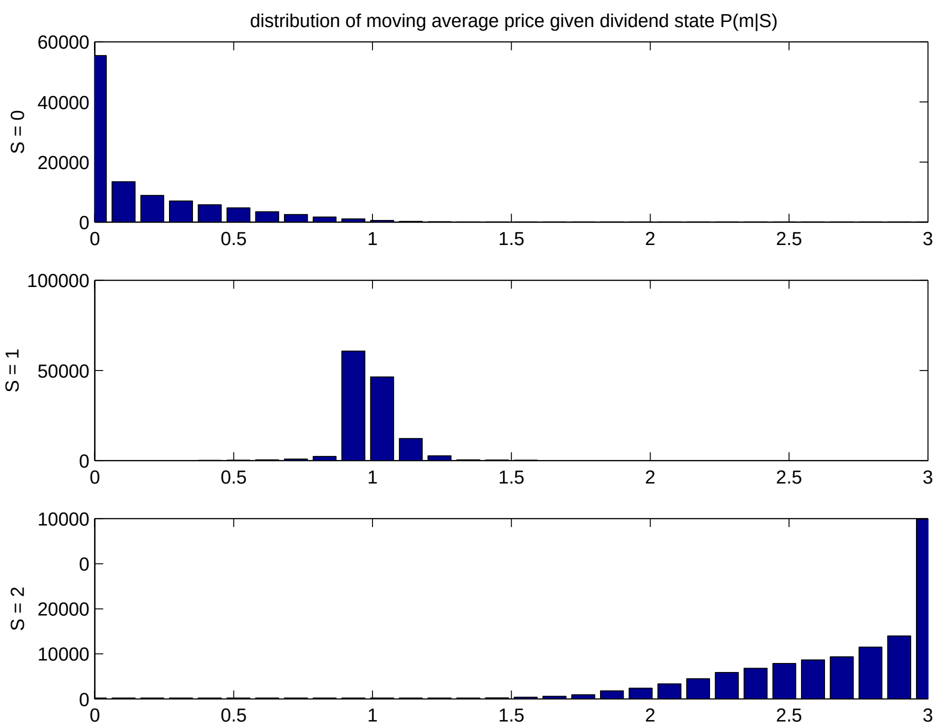

Figure 3d: Empirical distribution of moving-average prices, conditioned on the state of nature S, in Artificial Markets Experiment 4.2 (information dissemination with identical preferences).

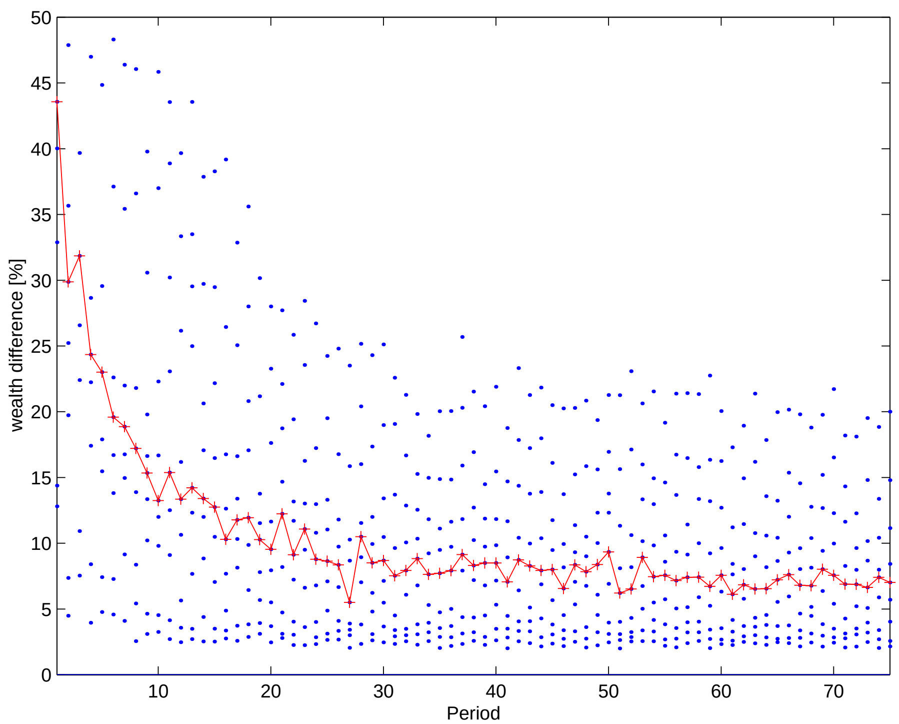

Figure 3e: Deciles of percentage wealth differences between insiders and uninformed traders in 100 repetitions of Artificial Markets Experiment 4.2 (information dissemination with identical preferences). Medians are indicated by the symbol '+'.

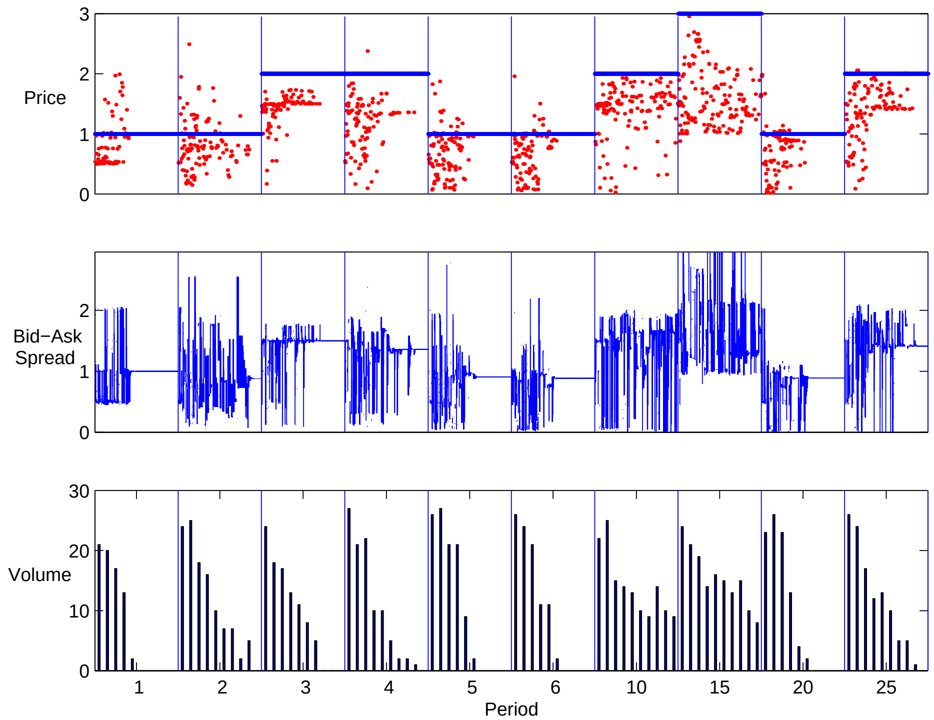

Figure 4a: Prices, bid-ask spreads, and volume in the early periods of a typical realization of Artificial Markets Experiment 4.3 (information aggregation with heterogeneous preferences).

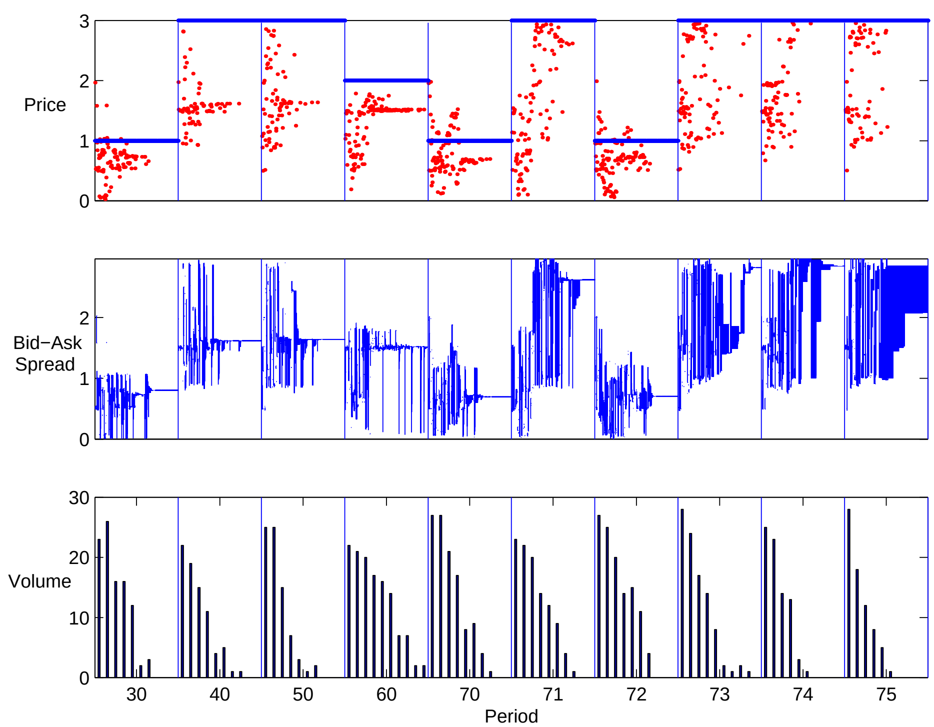

Figure 4b: Prices, bid-ask spreads, and volume in the later periods of a typical realization of Artificial Markets Experiment 4.3 (information aggregation with heterogeneous preferences).

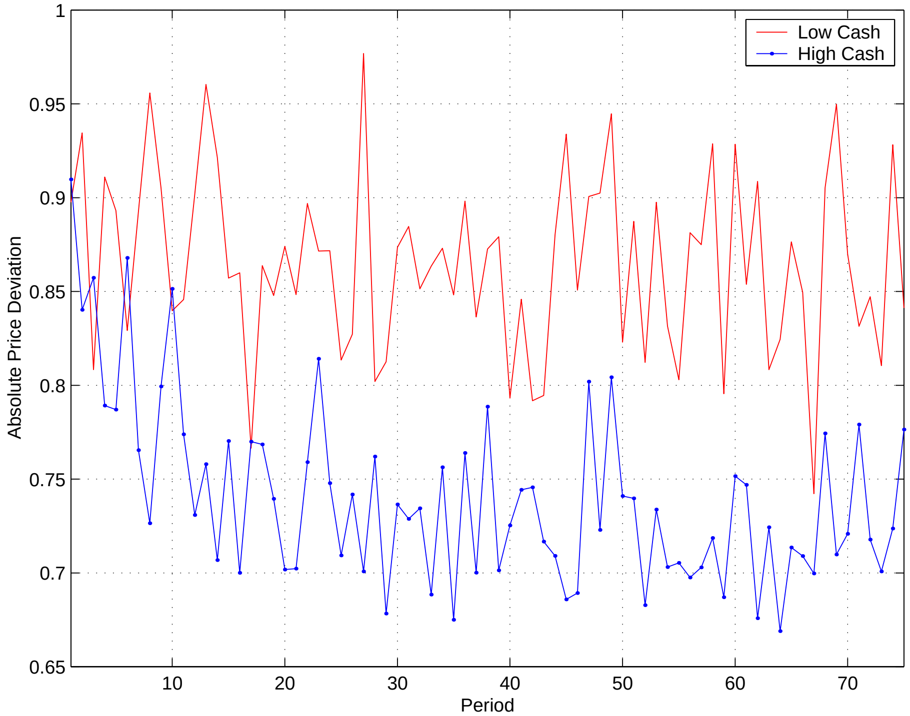

Figure 4c: Absolute price-deviations of market prices from the rational expectations equilibrium price, averaged over 100 repetitions of each of two runs of Artificial Markets Experiment 4.3 (information aggregation with heterogeneous preferences), the 'low-cash' and 'high-cash' experiments.

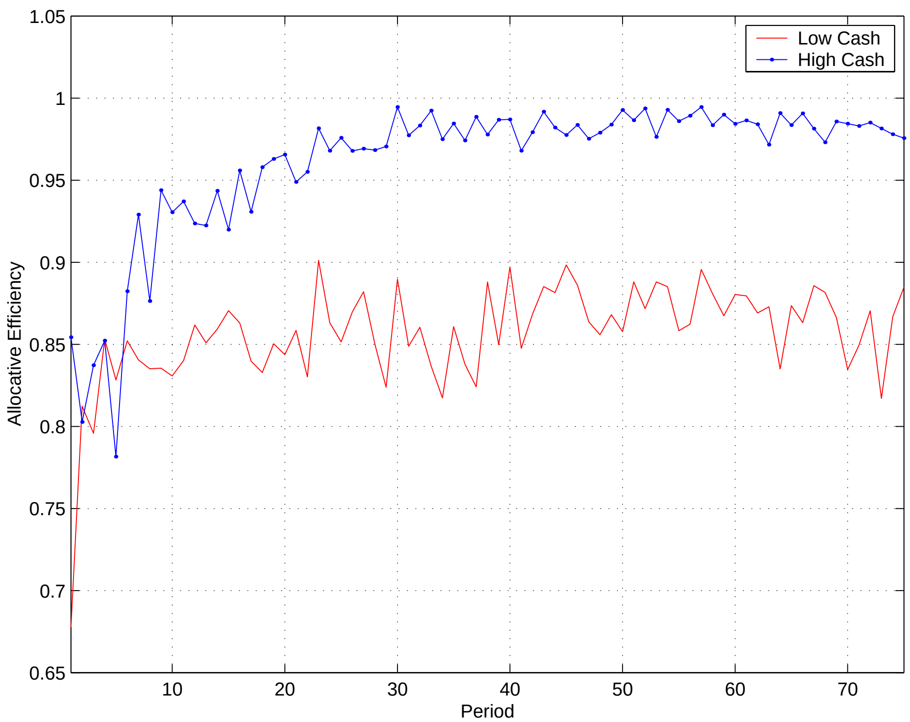

Figure 4d: Allocative efficiency, averaged over 100 repetitions of each of two runs of Artificial Markets Experiment 4.3 (information aggregation with heterogeneous preferences), the 'low-cash' and 'high-cash' experiments.

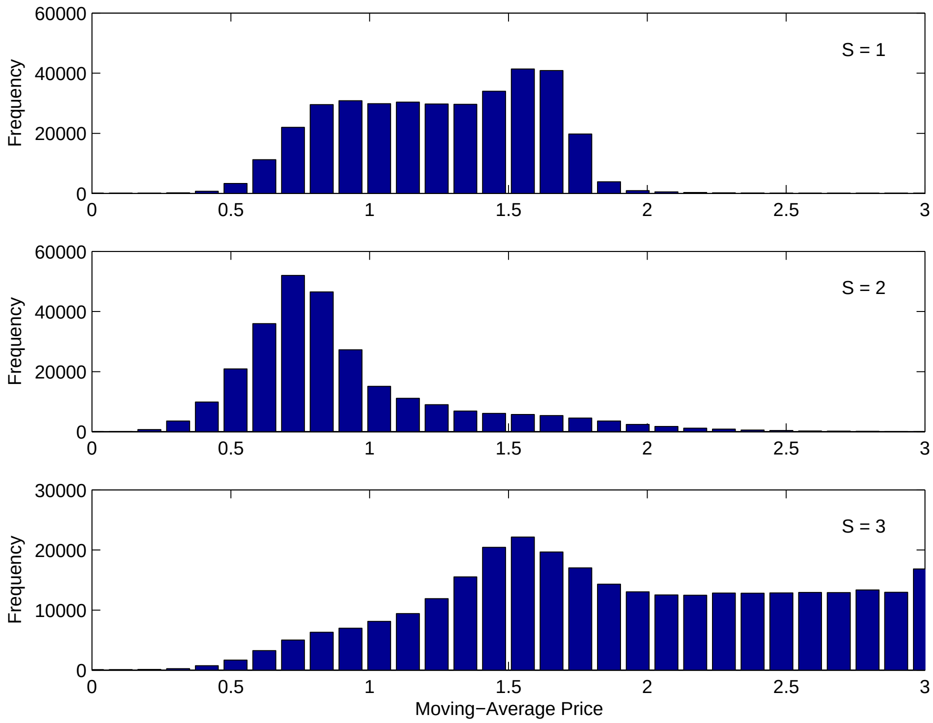

Figure 4e: Empirical distribution of moving-average prices, conditioned on the state of nature S, in Artificial Markets Experiment 4.3 (information aggregation with heterogeneous preferences).

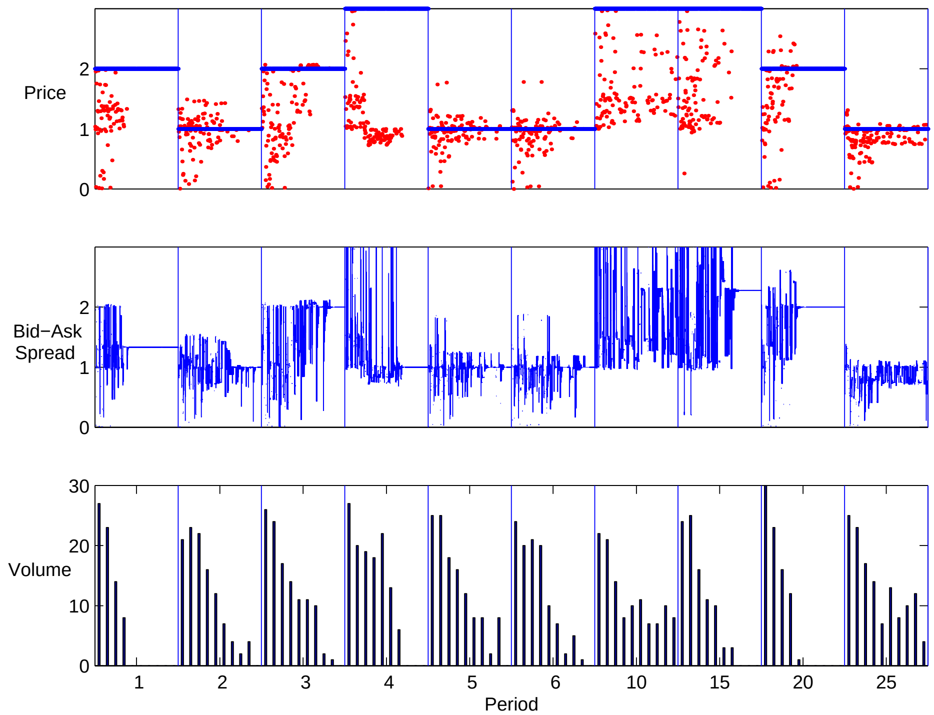

Figure 5a: Prices, bid-ask spreads, and volume in the early periods of a typical realization of Artificial Markets Experiment 4.4 (information dissemination with heterogeneous preferences).

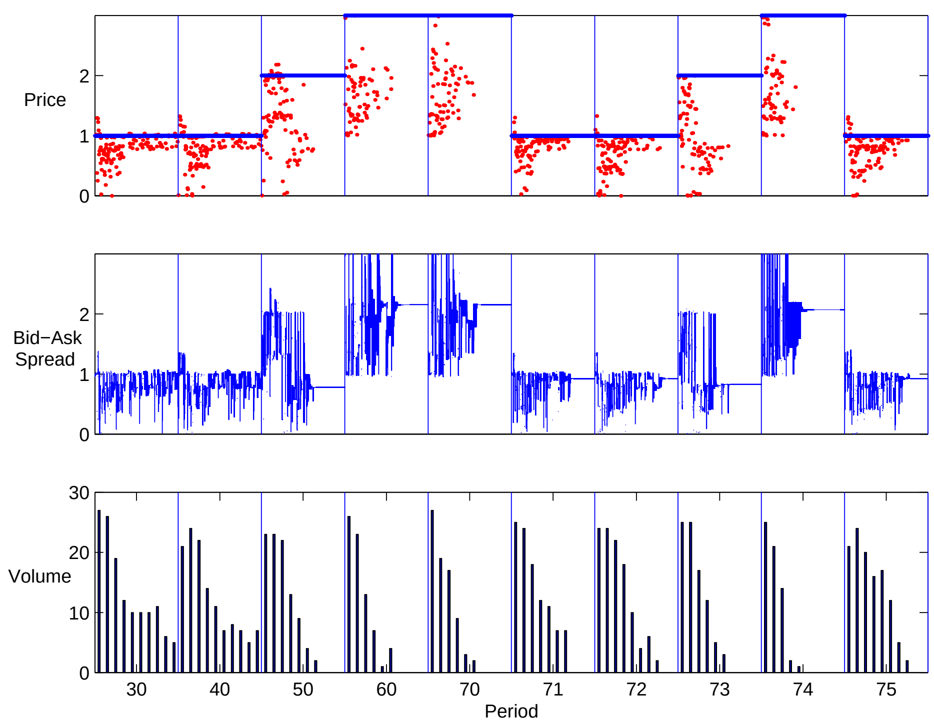

Figure 5b: Prices, bid-ask spreads, and volume in the later periods of a typical realization of Artificial Markets Experiment 4.4 (information dissemination with heterogeneous preferences).

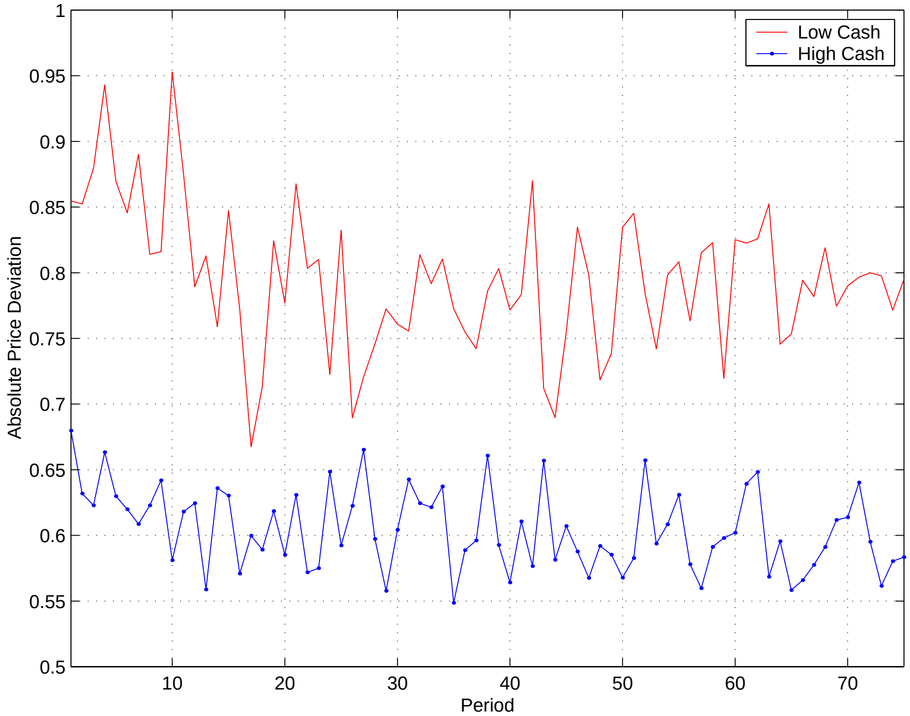

Figure 5c: Absolute price-deviations of market prices from the rational expectations equilibrium price, averaged over 100 repetitions of each of two runs of Artificial Markets Experiment 4.4 (information dissemination with heterogeneous preferences), the 'low-cash' and 'high-cash' experiments.

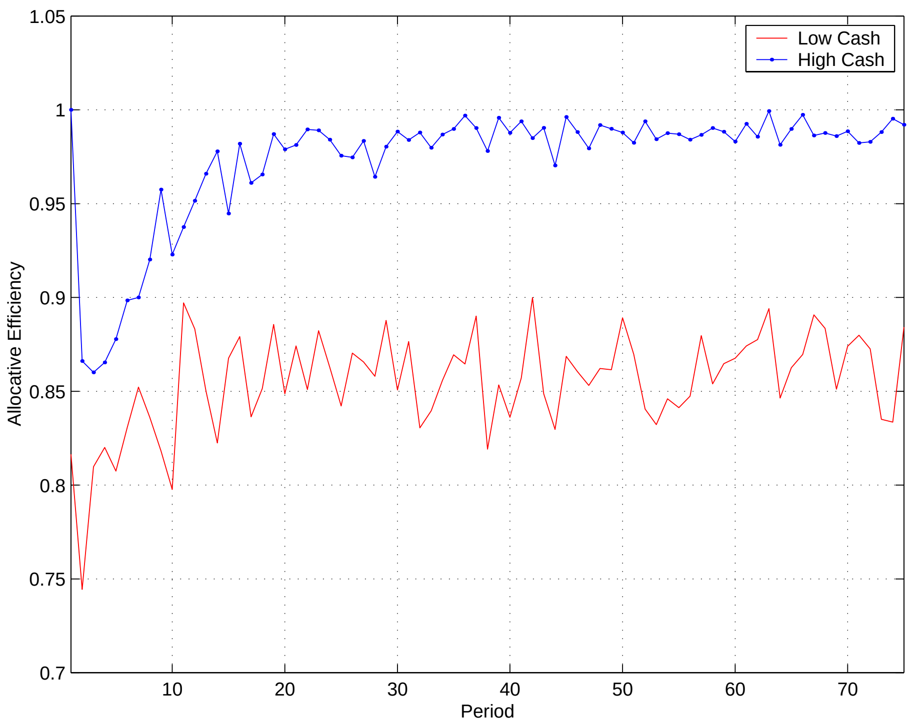

Figure 5d: Allocative efficiency, averaged over 100 repetitions of each of two runs of Artificial Markets Experiment 4.4 (information dissemination with heterogeneous preferences), the 'low-cash' and 'high-cash' experiments.

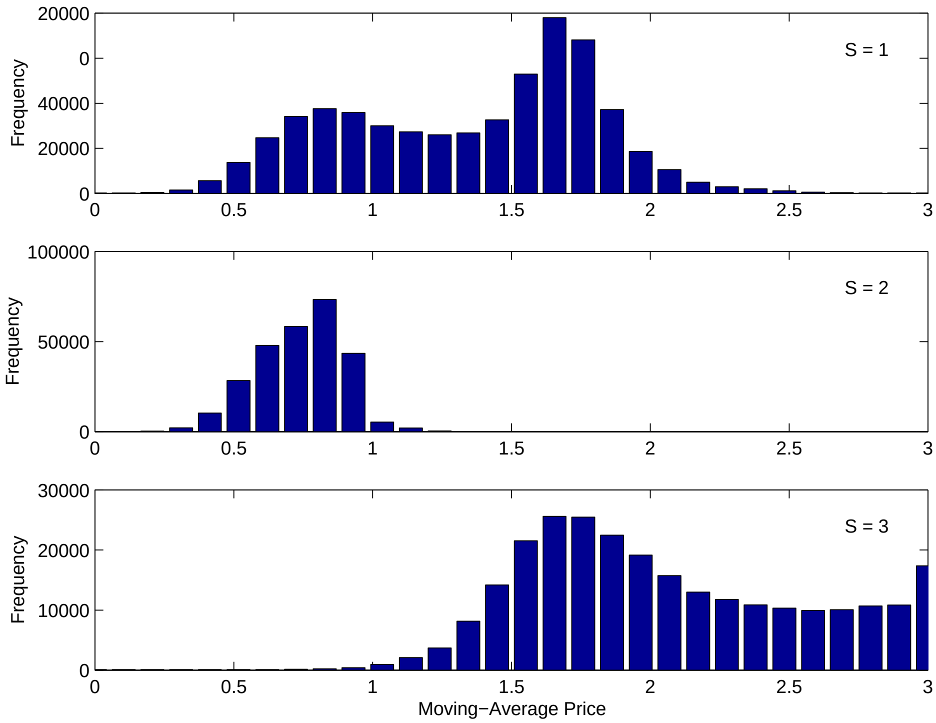

Figure 5e: Empirical distribution of moving-average prices, conditioned on the state of nature S, in Artificial Markets Experiment 4.4 (information dissemination with heterogeneous preferences).

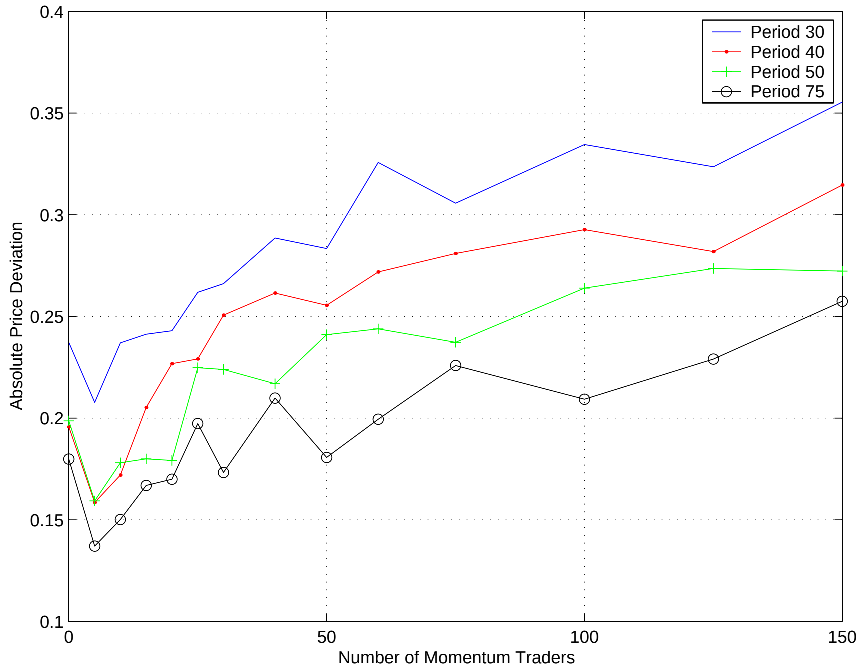

Figure 6a: Absolute price-deviations of market prices from the rational expectations equilibrium price in periods 30, 40, 50 and 75, averaged over 100 repetitions, as a function of the number of momentum traders present in Artificial Markets Experiment 4.5 (information aggregation with empirical Bayesian and momentum traders).

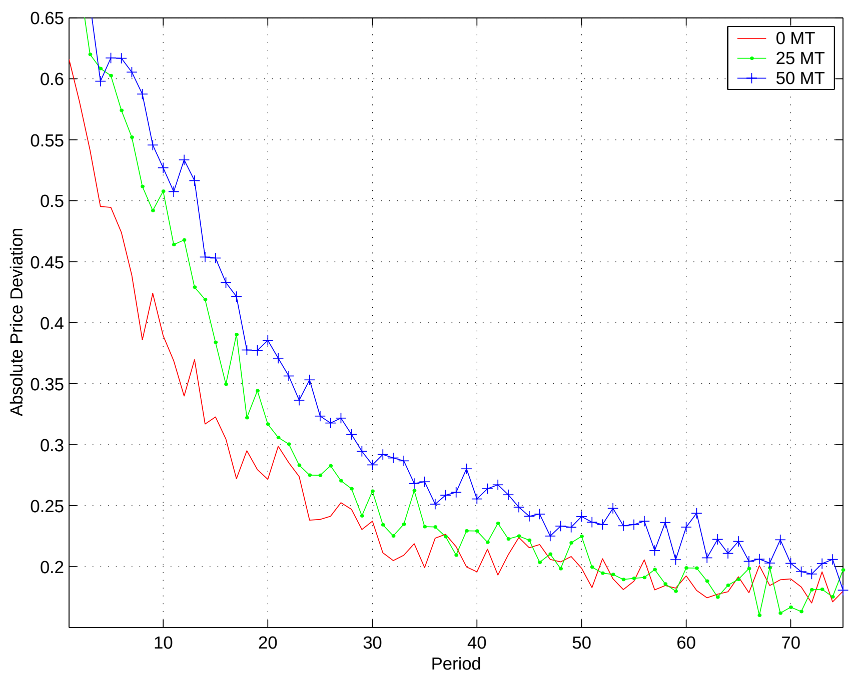

Figure 6b: Absolute price-deviations of market prices from the rational expectations equilibrium price, averaged over 100 repetitions, over the epoch for 0, 25, and 50 momentum traders in Artificial Markets Experiment 4.5 (information aggregation with empirical Bayesian and momentum traders).

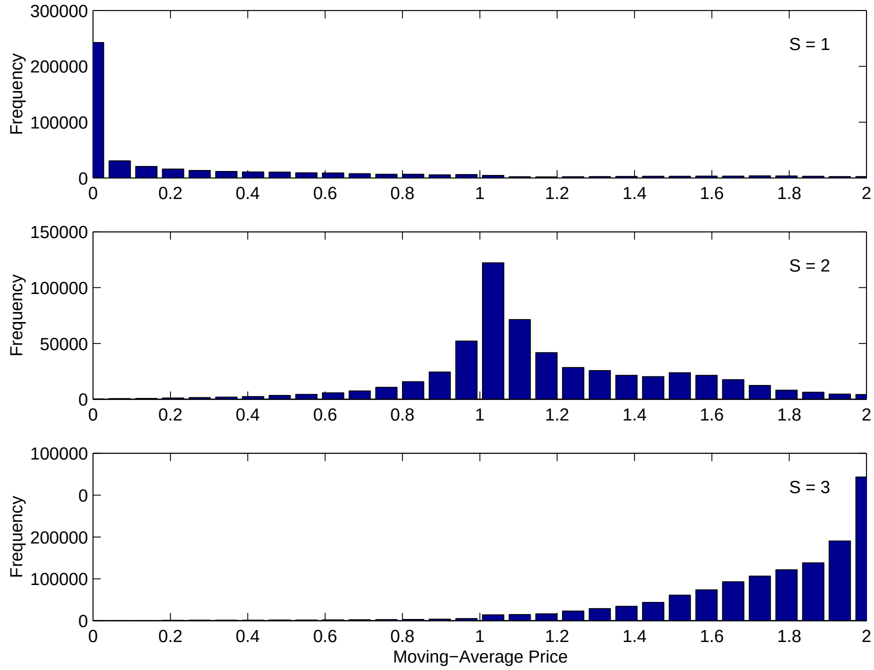

Figure 6c: Empirical distribution of moving-average prices, conditioned on the state of nature S, in Artificial Markets Experiment 4.5 (information aggregation with 20 empirical Bayesian and 20 momentum traders).

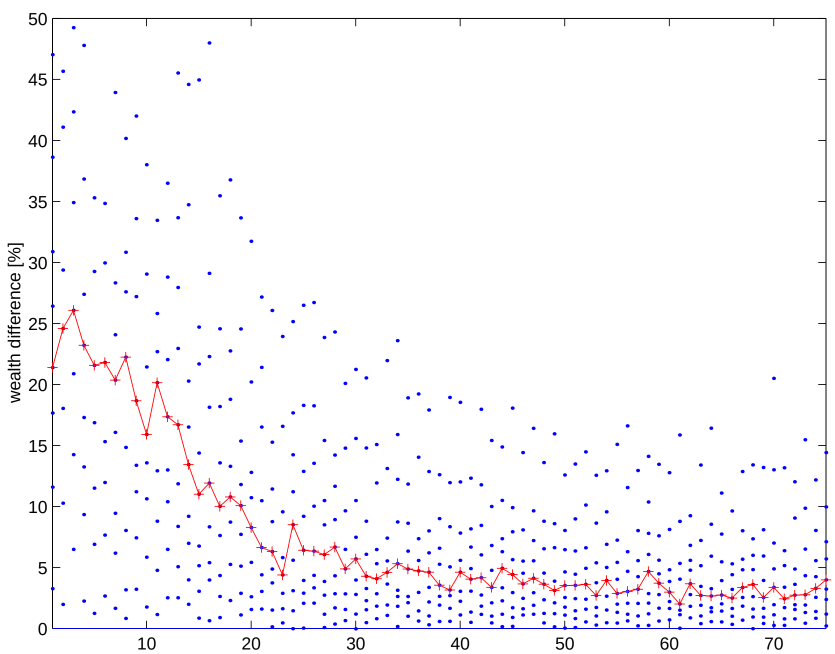

Figure 6d: Deciles of percentage wealth differences between empirical Bayesian and momentum traders in 100 repetitions of Artificial Markets Experiment 4.5 (information aggregation with 20 empirical Bayesian and 20 momentum traders). Medians are indicated by the symbol '+'.

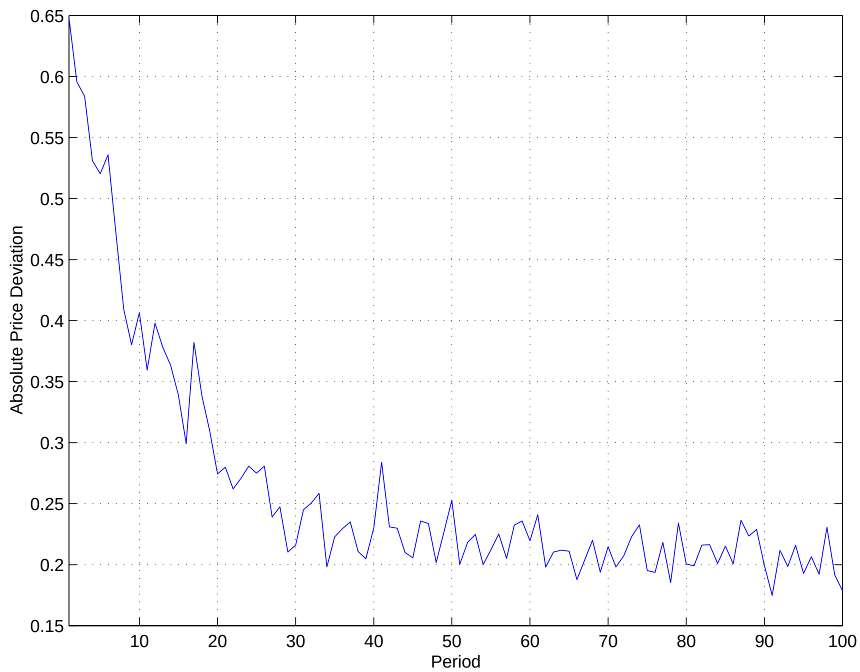

Figure 7a: Absolute price-deviations of market prices from the rational expectations equilibrium price, averaged over 100 repetitions, in Artificial Markets Experiment 4.6 (information aggregation with 15 empirical Bayesian and 5 nearest-neighbor traders).

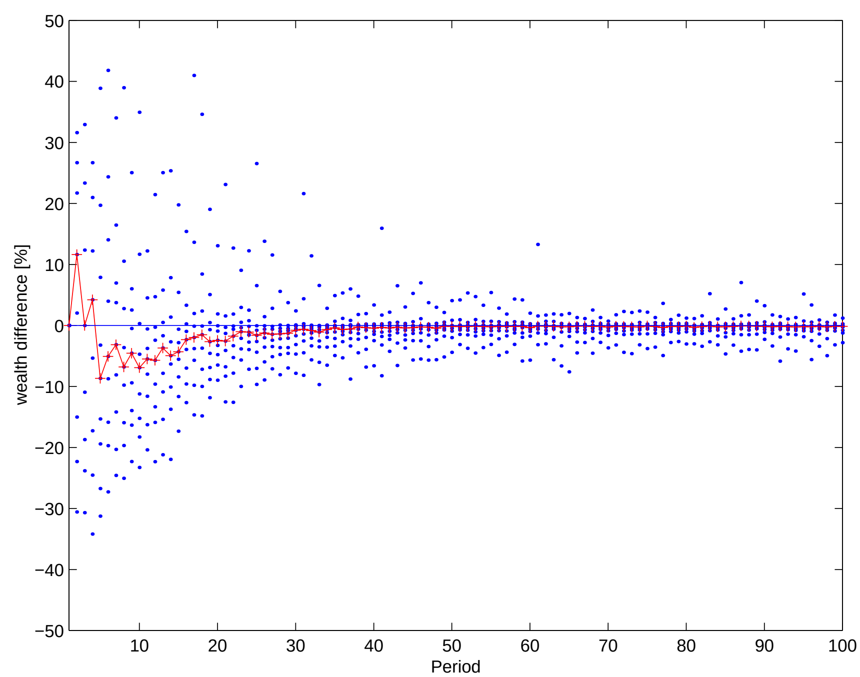

Figure 7b: Deciles of percentage wealth differences between empirical Bayesian and nearest-neighbor traders in 100 repetitions of Artificial Markets Experiment 4.6 (information aggregation with 15 empirical Bayesian and 5 momentum traders). Medians are indicated by the symbol '+'.