MASSACHUSETTS INSTITUTE OF TECHNOLOGY

ARTIFICIAL INTELLIGENCE LABORATORY

and

CENTER FOR BIOLOGICAL AND COMPUTATIONAL LEARNING

DEPARTMENT OF BRAIN AND COGNITIVE SCIENCES

A.I. Memo No. 1646 September, 1998

C.B.C.L Paper No. 164

Information Dissemination and Aggregation in Asset Markets with Simple Intelligent Traders

Nicholas Chan, Blake LeBaron, Andrew Lo and Tomaso Poggio

This publication can be retrieved by anonymous ftp to publications.ai.mit.edu.

The pathname for this publication is: ai-publications/1500-1999/AIM-1646.ps

Abstract

Various studies of asset markets have shown that traders are capable of learning and transmitting information through prices in many situations. In this paper we replace human traders with intelligent software agents in a series of simulated markets. Using these simple learning agents, we are able to replicate several features of the experiments with human subjects, regarding (1) dissemination of information from informed to uninformed traders, and (2) aggregation of information spread over different traders.

This report describes research done within the Center for Biological and Computational Learning in the Department of Brain and Cognitive Sciences at the Massachusetts Institute of Technology. This research is sponsored by a grant from the National Science Foundation under contract ASC-9217041.

1 Introduction

Experimental asset markets have yielded many results on the properties of financial markets, and their abilities to disseminate and aggregate information. This understanding of the behavior of partially informed agents in experimental settings is a critical step toward understanding behavior in real markets. Various studies in experimental markets have shown that individuals are able to learn, and transmit information through prices in many different market situations. However, these studies are less specific about the actual mechanism that traders use to process information and learn from experience. This kind of generalization requires a deeper understanding of traders' trading strategies and the specification of the underlying learning processes.

We begin to address this question through the use of computational learning agents. These agents take the place of the experimental subjects and trade with each other in a simulated asset market. Unlike experimental subjects, the characteristics of the computer agents can be carefully controlled and modified to study the overall market behaviors with regard to different properties of the trader population. In the design of our trader agents, we strive to keep them as simple as possible in order to give us an idea of the lower bound of intelligence needed to replicate various market phenomenon. This simplicity also makes the agents more open to detailed analysis on how they are processing market information.

Computational models allow us to explore new areas of economic theory, especially in dynamic market situations with learning. However, computational models bring with them new untested algorithms, and parameters. Questions about where theory ends, and how simple ad hoc mechanisms begin are quite valid. We believe that experimental data provides one useful route for validation. For this reason we design our markets to follow those in used in the experimental literature. Using simple learning agents we are able to replicate several features of the actual experiments including dissemination of information from informed to uniformed traders, and aggregation of information spread over different traders.

Section 2 gives a brief review of both the experimental and computational literatures. Section 3 describes our market setup, and gives specifics about the experiments. Section 4 presents our results, and section 5 gives conclusions and ideas for the future.

2 Review of the Literature

2.1 Experimental Markets

The rational expectations (RE) model has received a considerable amount of attention in research on experimental markets. The RE model has had mixed success in various studies, depending on the complexity and structure of a market. The study of informational efficiency in the context of RE models can be categorized into two major areas. The first studies information dissemination from a group of insiders who have perfect information to a group of uninformed traders. The idea is that market prices reflect insider information so that uninformed traders can infer the true price from the market. The second examines information aggregation of diverse information in a market by a population of partially informed traders. Aggregation of diverse information is in general more difficult because no single agent possesses full information. Traders can identify the state of nature with certainly only by sharing their individual information in the process of trading.

Plott & Sunder (1982) and Forsythe, Palfrey & Plott (1982) study markets with insiders and uninformed traders. They show that the equilibrium prices do reveal insider information after repetition of experiments and conclude that the markets disseminate information efficiently. Plott & Sunder (1982) further show that convergence to the rational expectation equilibrium (REE) occurs in markets that pays diverse dividends to different traders. They attribute the success of the RE model to the fact that traders learn about the equilibrium price and the state simultaneously from market conditions. The results by Plott & Sunder (1988) and Forsythe & Lundholm (1990), on the other hand, show that a market aggregates diverse information efficiently only under certain conditions: identical preferences, common knowledge of the dividend structure, complete contingent claims. These studies provide examples of failure of the RE model and suggest that information aggregation is a more complicated situation. In another related study, O'Brien & Srivastava (1991) find that market efficiency in terms of full information aggregation depends on complexity of the market. In particular, complexity is induced by market parameters such as the number of stocks and the number of periods in the markets.

2.2 Simulated Markets

Computer simulations and software agents extend the experimental approach by testing basic theories about learning behavior. Experiments use simple economic theories to test convergence properties, but the dynamics of the subjects behavior through the rounds is usually not modeled. The computer simulations performed here provide one possible method for testing the dynamics of learning in experimental settings, and developing theories in the form of agent algorithms which can be used to test further hypothesis on market designs and behavior.

Our agent design is based on the zero intelligence (ZI) traders used in Gode & Sunder (1993), where the generation of bids and offers contains a large random component. Gode & Sunder (1993) emphasize the impact of budget constraints alone on observed prices and market efficiency. Several other authors have begun adding some intelligence by further restricting the range of bids and asks that may be generated. Usually these restrictions involve some function of recently observed trades or quotes. Two examples of this are Jamal & Sunder (1996) and Cliff & Bruten (1997) which both implement simple heuristics to try to limit and improve on simple random bidding. Further examples of trading algorithms for the simple double auction can be found in the report on the Santa Fe Institute Double Auction Tournament, (Rust, Miller & Palmer 1992). This tournament focused on the relative performance of various strategies played against each other. One of its key findings was that a very simple parasite strategy that fed off the others performed the best.

Finally, more complex computer simulated asset markets, which emphasize the evolution of trading behavior over time have also been created. LeBaron (forthcoming 1998) surveys many of these other computational markets[1] These more complicated simulated markets are interested in long range market phenomenon, and less interested in the actual trading mechanisms which are at the center of our research. However, they share our emphasis of building behavioral theories starting at the individual level.

3 Design of the Simulation

Our simulations are conducted with computer traders, who possess private information of the economy, and participate in a double auction market. They trade a single stock that pays a liquidating and state-contingent dividend at the end of the period. A period of a simulation starts with the following: (1) the determination of the state of nature, (2) distribution of identical portfolios of cash and stock to all agents and (3) distribution of private information. A period is further divided into numerous trading rounds, at which agents submit orders. At the end of a period, the predetermined state of nature is revealed and dividends are distributed to the share holders. We conduct five different experiments with specific information and dividend structures. Each experiment has fifty periods, each consisting and independent draw of the state of nature and private

[1] A few early examples of these types of markets are, Arthur, Holland, LeBaron, Palmer & Tayler (1997), Arifovic (1996), Lettau (1997), Youssefmir & Huberman (1995).

information (see figure 1).

3.1 The Economy

There is a one-period stock that pays a state-contingent dividend at the end of the period. The state is unknown, random and exogenously determined. The underlying distribution of the state is common knowledge. In particular, it is discrete and uniform. In an economy with three possible states, for instance, if the stock pays a dividend $D = \{\$0, \$1, \$2\}$ to all traders, each share of it pays \$1 in state 1, \$2 in state 2 and so on. Each of the possible dividend states has a probability of 1/3. This is the case when traders have *homogeneous preferences*. When traders have *heterogeneous preferences*, the stock could pay different dividends to different stock holders, according to their preferences. In this sense, the dividend is considered as a measure of utilities to the stock holders. For example, a market could have two types of agents. Type I and type II agents have dividend profiles $D^I = \{\$0, \$1, \$2\}$ and $D^{II} = \{\$2, \$0, \$1\}$ respectively. This means that in state 2, type I agents receive \$1 and type II get \$0 for owning the stock. The differences in traders' preferences is one of the incentives for trading in the market.

The difference in information or knowledge about the economy is the other reason why agents trade with each other. Information that is available to all market participants is public information, whereas information only known to some individuals is considered private information. The possible dividend payoffs and their associated probabilities is public information. Some traders receive private information about the state of the economy. Traders are categorized into three groups according to what they know about the economy. There are insiders who know exactly what the states will be (for instance, $D = \{\$1\}$), partially informed traders who have imperfect information about the states (such as $D = \{\$0, \$1\}$ or $D = \{\$1, \$2\}$), and uninformed traders who have only the public information (that is $D = \{\$0, \$1, \$2\}$). Insiders and partially informed traders receive their private information at the beginning of each period of the experiment. The distribution of private information is not common knowledge.

3.2 Trading Mechanism

The trading mechanism is a simplified double auction market. Traders can either submit a bid or ask, or accept a bid or ask. If there is an existing bid for the stock, any subsequent bid must be higher than the current bid. Similarly, a subsequent ask following an existing ask must be lower than the current ask. A transaction occurs when an existing bid or ask is accepted (a market order matches with a limit order), or

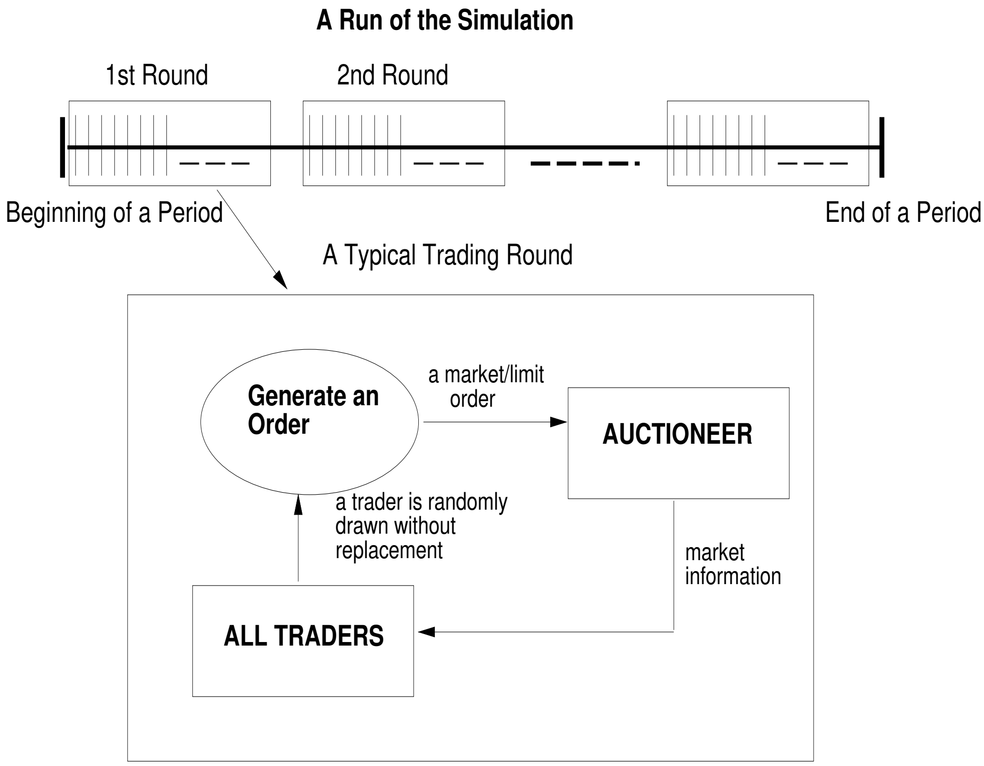

Figure 1: The design of the simulation.

when the bid and ask cross (in which case, the transaction price is the middle of the bid and ask).

The quantity of each trade is restricted to one share. There are two reasons for such a substantial simplification. First of all, we want to keep the modeling of agents unsophisticated and intuitive. The main idea of this paper is to show how agents with simple heuristics can replicate what sophisticated humans do in a market. It is also our interest to focus on the prices of the stock as a media of information transmission among different agents.[2]

No borrowing or short selling is permitted. Traders have to trade subject to their budget constraints. Each period of an experiment consists of 40 trading rounds for all traders. At the beginning of each round, a random permutation of the traders is determined. Following the sequence in the permutation, each trader submits one limit or market order. There are 20 agents and therefore a maximum of 800 transactions can be recorded in one period.

3.3 Traders

Agents are supposed to maximize the the end-of-period value of their portfolios by forecasting the dividend, buying low and selling high at the market. The forecast is done by utilizing the public, private and market information.

The market is populated with three types of traders according to their trading objectives and trading strategies: noise traders, empirical Bayesian traders[3] and momentum traders. The former ignore all the available information and trade randomly for liquidity needs but not for profit making. They have no memory or learning. They only submit market orders randomly according to a preset probability to trade.

Empirical Bayesian traders, on the other hand, act more intelligently. They utilize market information to update their beliefs regarding the states of the economy. According to their beliefs, they form expectations of the price, which we call the base price. They 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). Orders are submitted according to the procedure described in table (1).[4] For instance, if there exists only an ask (no outstanding bid) and the agent's base price is lower than the ask price, it posts a bid uniformly distributed from $(bp - S, bp)$, where

[2] It is important to keep in mind that quantities of transactions is an important aspect of the market. It could be associated with, for example, agents' risk aversion. This aspect of the market will be studied in the future.

[3] We use the term empirical Bayesian, but 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 summary of observed data which is used to update the priors.

[4] 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.

bp is the base price and S is a preset maximum spread. Empirical Bayesian traders are assumed to be risk neutral and maximize the end-of-period wealth by choosing between cash and stock.[5] They continuously observe the market activities, update their beliefs and adjust their positions. They stop trading when either the market price approaches their expected price, or they run out of cash or stock.

Momentum traders are simple technical analysis traders who believe that tomorrow's return equals today's return. In our simulation, suppose at time t the previous two transaction prices are p_t and p_{t-1}, a momentum trader predicts that the next transaction will occur at $p_{t+1} = p_t \left(\frac{p_t}{p_{t-1}}\right)$. In effect, these traders reinforce and magnify the ups and downs of price movements. The presence of momentum traders introduce extra randomness and irrational valuations of the security, which make information aggregation and dissemination more difficult.

3.4 Learning Mechanism

The empirical Bayesian traders condition their beliefs on market information. Specifically, the agents want to compute the expected dividend $E(D|p_0, p_1, ..., p_t)$. For simplicity, we only consider transaction prices and ignore other market variables such as bid, ask, or volume. Empirically, we further assume that most of the relevant information is embedded in the transaction prices of the last k trades. The k-period moving average of prices m_t is used to summarize market information at time t,

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

Given the series of moving average price $m_k, m_{k+1}, ..., m_t$ and the realized dividend D_i, one can empirically estimate the conditional distribution $P(m|D_i)$. Using Bayes Theorem, $P(D_i|m)$ can be determined,

$$P(D_i|m) = \frac{P(m|D_i)P(D_i)}{\sum_{j=1}^{N} P(m|D_j)P(D_j)} \ ,$$

where $P(D_i)$ 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, in the case when $D = \{D_0, D_1, ..., D_n\}$, given a price statistic m, the conditional expectation of the dividend is

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

[5] Note that we do not explicitly model the utility functions of the traders. Their preferences are reflected in the dividends they receive. For instance, a trader who gets $D_1 = \$2$ and $D_2 = \$0$ prefers state 1 to state 2.

scenario		action
there exists a bid and an ask		
	$bp > a$	buy at market
	$bp < b$	sell at market
	$b < bp < a$ and $a - bp > bp - b$	post an ask distributed $(bp, bp + S)$
	$b < bp < a$ and $a - bp \leq bp - b$	post an bid distributed $(bp - S, bp)$
there exists only an ask		
	$bp > a$	buy at market
	$bp < a$	post a bid distributed $(bp - S, bp)$
there exists only an bid		
	$bp > a$	sell at market
	$bp < a$	post an ask distributed $(bp, bp + S)$
there exists no bid or ask		
	flip a coin	
	HEAD	post an ask distributed $(bp, bp + S)$
	TAIL	post a bid distributed $(bp - S, bp)$

Table 1: This table describes the procedure followed by computer traders to submit an order. The variables a denotes the best ask, b the best bid, bp a trader's base price and S the maximum spread from the base price.

The conditional expected price is taken as the base price for the empirical Bayesian traders. The order submission procedure, as described in table(1), is 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.

3.5 Experiments

We conducted five computer experiments, each of which has the same market and information structure but differs in two aspects. First, the composition of traders could be different regarding their information sets. In particular, we considered markets populated with partially informed traders, and markets with insiders and uninformed traders. In the former case, we studied how diverse information is aggregated in the market.

The latter is for the study of information dissemination. Second, traders in the markets could have either identical or diverse preferences, which depend on the state-contingent dividends they get.

In all experiments, we are interested in the informational efficiency of the markets. In particular we focus on how efficient the markets aggregate and disseminate information by measuring to what extent the prices reflect all the available information. Specifically, we focus on the convergence of prices to what the REE predicts. In the cases when traders have diverse preferences, allocative efficiency is also studied.

- **Experiment 1: Information aggregation by traders with identical preferences**

  The economy has three possible states and respectively the stock pays a dividend $D = \{\$0, \$1, \$2\}$ with equal probability. There are 20 traders. All are partially informed that one of the three states is impossible. If state 1 is the true state, for instance, a trader will be told that state 0 or 2 is impossible with probability of 1/2. This is a situation where none of the traders knows the state of nature, but all traders, or the market as a whole has perfect information about true state. The REE price is simply the value of D given the state of nature.

- **Experiment 2: Information dissemination by traders with identical preferences**

  The stock pays a dividend $D = \{\$0, \$1, \$3\}$ to all traders. There are 10 insiders who know what the state of nature is, and 10 uninformed traders who have only public information: the distribution of D. The REE price is D given the state.

- **Experiment 3: Information aggregation by traders with diverse preferences**

  Traders are divided into two groups of 10 according to their preferences, 10. In the three possible states of the economy, Group I gets dividend $D^I = \{\$0, \$1, \$3\}$ and Group II gets $D^{II} = \{2, 0, 1\}$. All traders have private information that eliminates one of impossible states. Given the state of nature, the REE price is the higher of D^I and D^{II} in that particular state. For example, given that state 2 will occur, the REE price equals \$3. In this experiment and experiment 4, we also vary the amount of cash assigned to the traders at the beginning of each period. The amount of cash determines the budget constraints faced by the traders. There are two levels of cash endowment. The low cash endowment implies a stringent budget constraint and high cash endowment corresponds to a relax budget constraint.

- **Experiment 4: Information dissemination by traders with diverse preferences**

  There are two groups of traders with diverse preferences. Group I gets dividend $D^I = \{\$0, \$1, \$3\}$ and group II gets $D^{II} = \{\$2, \$0, \$1\}$. There are 5 insiders and 5 uninformed traders in group I and II respectively. The REE price is the higher of D^I and D^{II} given the state.

- **Experiment 5: Information aggregation by empirical Bayesian and momentum traders**

  The purpose of this experiment is to test the robustness of the previous results. The empirical Bayesian traders not only serve as a source of information but an important mechanism for the market to reach the REE. The momentum traders, on the other hand, introduce a substantial amount of noise and false information to the markets. It is interesting to observe whether the empirical Bayesian traders can maintain an efficient market, in the presence of erroneous signals.

  In this experiment, there are three markets each of which has a different composition of traders. In term of the number of empirical Bayesian traders versus momentum traders, we have the following compositions: (1) 5 versus 5, (2) 5 versus 25 and (3) 10 versus 10.

4 Results and Discussion

We study the price efficiency of a market by observing how close and how fast prices converge to the REE prices. Since the REE is a full information revealing situation, convergence to the REE price indicates how efficient information is aggregated or disseminated in a market. In particular, we consider the absolute price deviation of the last 20 transactions[6] from the theoretical price as a measure of closeness to the REE. Other variables that we consider are the bid-ask spread, trading volume and the wealth of the traders. Narrowing bid-ask spreads shows that the prices are converging. It implies that buyers and sellers have reached a common price. A 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 further transact. Lastly, in the case of insiders versus uninformed traders, the differences in wealth provides a measure of how much insider information is worth. Specifically, the difference is presented as a percentage difference in total wealth between two groups of traders[7].

We also examine the expectations of the agents by noting their empirical conditional density functions of the moving average price given the states. This collection of conditional density functions is the agents' belief formed using their prior information and updated continuously by observing the market prices. The

[6] Only the last 20 transactions are considered because prices usually fluctuate immensely before the market reaches equilibrium. Traders need time to observe the market and infer the prices.

[7] Given the total wealth of group 1 (w_1) and group 2 (w_2), the percentage difference is defined as $\frac{w_1 - w_2}{w_2}$.

agents use these density functions to distinguish one state from another. These functions are crucial to understanding how the agents are learning.

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 can possibly extracted from the market. For example, 100% allocative efficiency implies that all shares are held by traders from the group that gets the highest dividend in the realized state. The REE predicts a 100% market efficiency in that all shares will be allocated to the highest valuing traders.

In general, we find convergence to the REE in the identical preference cases, but difficult or unattainable in experiments that involve diverse preferences. This can be explained by the fact that our software agents only attempt to recover the state of nature from market information, but not preferences of other agents, which is not common knowledge. In fact, they do not even realize or consider the possibility of different types of dividend payoffs.

Most of figures and statistics presented in this section are either total or average measures taken from 100 trials of the same experiment. For instance, price deviations reported are averaged over 100 trials of the same experiment and the frequencies presented in the empirical distributions are summed values.

4.1 Homogeneous Preferences

Under uniform preferences, the results from our simulation are similar to what were found in the experimental markets literature. In both cases of information aggregation (**Experiment 1**) and dissemination (**Experiment 2**), the convergence to the REE is apparent. Figure 2 shows the market activities in the some earlier periods of the market in a typical trial of the simulation. This is the stage when agents are actively learning and observing with little evidence for convergence. In the later periods (see figure 3), when the agents accumulate enough knowledge on how states and prices are related, convergence becomes more observable. In terms of the absolute price deviation, market efficiency improves substantially over the periods (figure 4). The distributions of moving average prices given the state (figure 5) shows that the three states are clearly distinguishable by the agents.

In Experiment 2 the evidence of convergence is even more compelling (see figures 6 and 7). Comparing to the market in Experiment 1, the prices converge faster and closer to the theoretical price, and the bid-ask spreads are smaller. Further, the price deviations are lower than that in Experiment 1 (figure 8). There are two reasons for such a difference in the two experiments. Arguably both markets have the same amount of information. However in Experiment 1 traders need to trade with each other to "pool" their information to come up with true price, whereas in Experiment 2 the insiders simply know what the true price is. Secondly, in the former case the distribution of imperfect information to the traders is random. For instance, there could be much more traders who receive the information that $D = \{\$0, \$1\}$ than those who receive $D = \{\$1, \$3\}$. Therefore, their consensus could be biased in one way or another. Figure 10 shows the accumulated wealth of the insiders and uniformed traders. The insiders have a substantially higher accumulated wealth than the uninformed. The difference in their wealth represents the value of the insider information and might be an estimate of the price traders would be willing to pay if the information signals were sold. Note that the value of insider information is diminishing across the periods as learning occurs. This result is consistent with that from markets in Sunder (1992) where information is sold in a sealed bid auction. Traders in these markets lower their bids for information after learning to infer the states with a few periods of experience.

4.2 Heterogeneous Preferences

This is a more difficult case. Agents have to learn not only the states of nature, but also the others' preferences (or the dividend payoff schedule). Figures 11 to 20 show the results from **Experiment 3** and **4**. Since our agents are only capable of learning the state from market prices, we expect the REE model to fail in both cases. The conjecture can be explained by the empirical conditional distributions or the histograms. In figure 20, for example, given the moving average price $m = 0.9$, it is almost equally likely for state 0 or 1 to be true. Hence, agents cannot effectively identify one state from another. Even if they were told what the state is, they would still have trouble reaching a unanimous price, because that depends on their individual dividend payoffs.

Depends on the initial amount of cash given to the traders, however, we observe different level of market efficiency, measured by the average price deviation and allocative efficiency. With low initial cash endowment to the traders, the market does not show much convergence. Absolute price deviations and allocative efficiencies do not improve much across the periods. When providing a larger amount of cash, however, we do observe some convergence in Experiment 3 and 4. A concrete example will help to illustrate how the market reaches equilibrium. In Experiment 3, for instance, type I and type II 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 II to buy from type I. The uninformed agents, without any private information, will have a reservation price approximately equal to \$1 regardless of their dividend profile. It is approximate because their beliefs, conditioned on the market prices, can affect their estimates of the price. Hence, we can conjecture that the transaction prices will range from \$1 to \$3. Note that type II insiders will bid the highest price (close to \$3) and they will never sell the shares out. The rest will try to buy or sell at roughly \$1 but type II insiders will be responsible for most of the buying. Consequently supply diminishes and price goes gradually to \$2. Not surprisingly, we also observe close to 100% allocative efficiency in the market. However, the large bid-ask spread shows that many traders are still interested in trading at prices far from the REE price, and there is little improvement in this spread across periods.

Information dissemination in a market with diverse dividends (Experiment 4) is unsuccessful by our learning agents. This contrasts with the laboratory markets studied by Plott and Sunder (1982), where after a few repetitions, insiders began to realize that the actual equilibrium price could be different from what their dividend profiles say, and adjust their trading strategy accordingly. Similarly, uninformed traders could also infer the equilibrium price from market conditions. The key distinctions between these experimental markets and our simulation are human traders' knowledge of the existence of other dividend payoffs, and their ability to learn about the association of the equilibrium price and price given the state.[8] The lack of these two make convergence impossible in our simulation.

In the market of diverse information and diverse dividends (Experiment 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 similar experiments with human subjects, Plott and Sunder (1988) show that information aggregation was unsuccessful in a market with diverse dividends, attributing 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 1), or replacing the single three-state security with three state-contingent claims. In a separate study, Forsythe and Lundholm confirmed similar results and added that information aggregation can be successful if the diverse dividend structure 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.

[8] In the context of Experiment 3, uninformed traders from group I pays as high as \$1, for one share of the stock. Given their dividend payoffs, they can make a profit of \$1 and break even in states 1 and 3 . However, they will soon realize that in both states, the prices are high than \$1, which they cannot afford. On the other hand, when the stock does trade below \$1, it would be certainly in state 2 when the stock pays zero dividend. By going through this kind of thinking, human traders can associate the equilibrium price with price given the state.

4.3 Momentum Traders

While the rational expectation model is successful in the homogeneous preference case, we are interested in testing its robustness. Momentum traders are added to the market to introduce extra noise or even false information to the signal perceived by the empirical Bayesian traders. In Experiment 5, we set up three markets with different compositions of partially informed empirical Bayesians and momentum traders. Figure 23 presents the mean absolute price deviations in the three markets of different compositions of traders. The 10 versus 10 market shows the best performance. In fact, its average prices are at approximately the same level as those in Experiment 1. With *fewer* empirical Bayesians, price convergence is more inferior in the other two markets as one would expect. Interestingly, the presence of extra number of momentum traders does not deteriorate the price efficiency significantly. In this kind of market, the empirical Bayesians, which effectively aggregate information, are crucial to the efficiency of the market. The empirical distributions of the moving average price have a high dispersion, but states are still discernible (figure 24).

We would expect the empirical Bayesians to take advantage of the irrational momentum traders and thus end up with a higher level of end-of-period wealth. The farther the price deviates from the theoretical prices, the higher is the gain of the empirical Bayesians. We observe that the wealth differential increases initially (till period 6 or so) and declines afterwards. The initial increase is due to the learning of the empirical Bayesians. After some point, the market becomes more efficient, and in turn the price becomes more informative and less deviated from its theoretical values. Thus the relative advantage of the empirical Bayesians diminishes.

5 Conclusions

We study information dissemination and aggregation by simple learning agents in simulated markets with different setup. The results from the simulations are consistent with those in the experimental markets, except for the case of information dissemination under diverse preferences. The failure in this particular case exposes the limitations of our simple learning agents. It also reflects the underlying complexity of the asset markets, albeit their unrealistically simple setup.

To study more realistic market scenarios in the future, we will introduce more sophisticated market and information structures. At the same time, we will place the learning agents in an evolutionary environment, by parameterizing the learning algorithms and letting these parameters mutate and evolve, subject to selection pressure in the market environment. The dynamics of the evolved agents will yield important insights into the persistence and stability of different types of trading behaviors. It is possible that we find a continually changing population as in Lindgren (1992) or a parasite strategy that appears to occasionally take advantage of others as in Rust et al. (1992), or converge to a kind of Nash equilibrium.

References

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.

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.

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.

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

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

LeBaron, B. (forthcoming 1998), 'Ragent 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.

Lindgren, K. (1992), Evolutionary phenomena for social dynamics, *in* C. G. Langton, C. Taylor, J. D. Farmer & S. Rasmussen, eds, 'Artificial Life II', Addison-Wesley, Reading, MA, pp. 295–312.

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

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.

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.

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

Youssefmir, M. & Huberman, B. A. (1995), Clustered volatility in multiagent dynamics, Technical report, Santa Fe Institute, Santa Fe, NM.

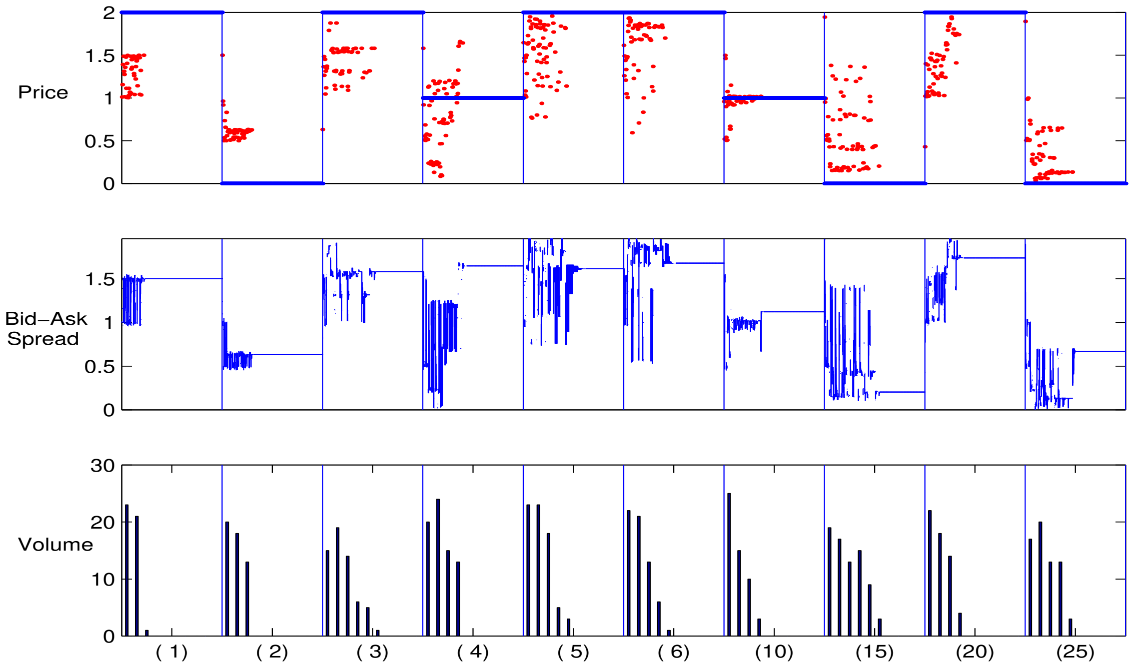

Figure 2: Information aggregation with identical preferences. The prices, bid-ask spreads and volume in some earlier periods of a typical experiment.

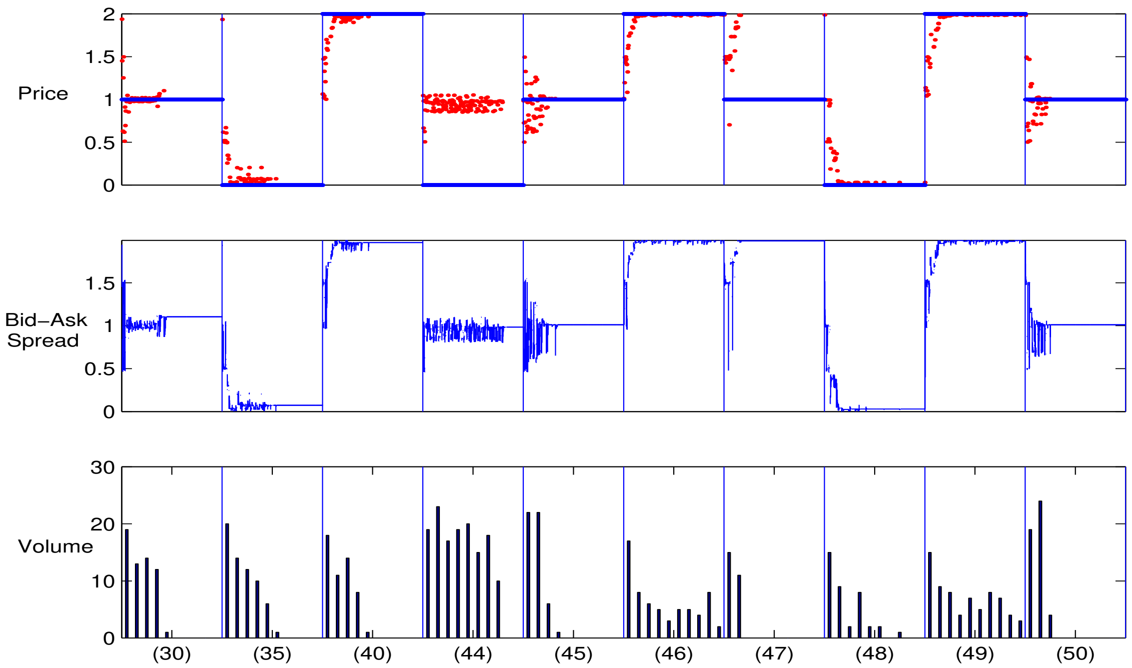

Figure 3: Information aggregation with identical preferences. The prices, bid-ask spreads and volume in some later periods of a typical experiment.

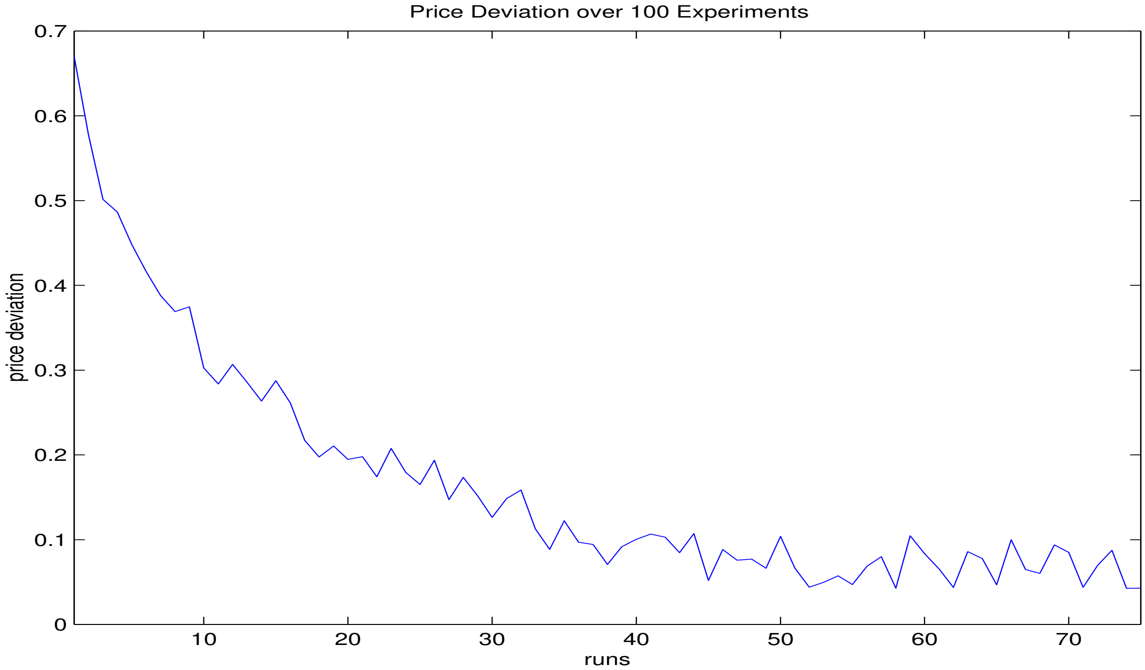

Figure 4: Information aggregation with identical preferences. Absolute price deviations from the theoretical price averaged over 100 experiments.

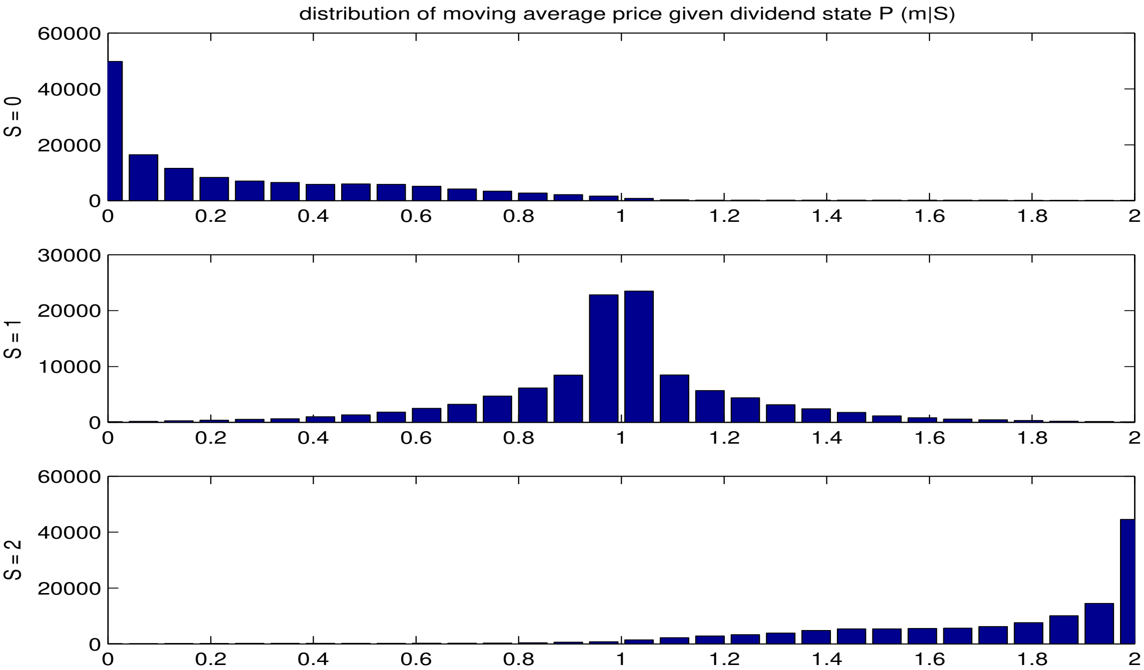

Figure 5: Information aggregation with identical preferences. Distributions of prices given the state of nature.

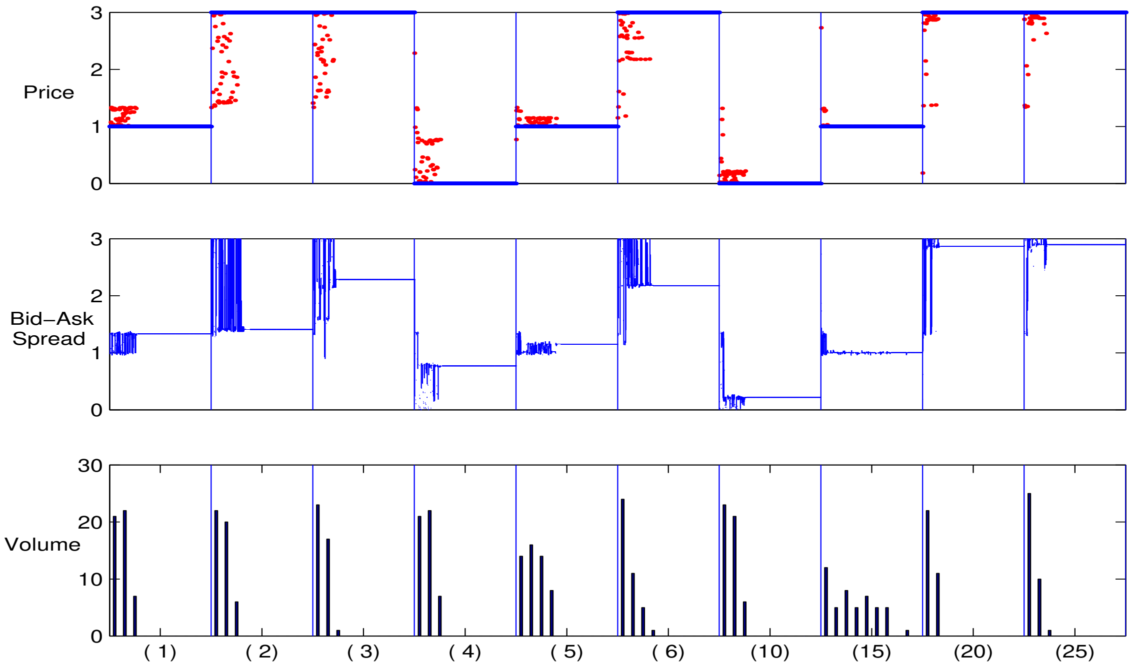

Figure 6: Information dissemination with identical preference. The prices, bid-ask spread and volume in some earlier periods of a typical experiment.

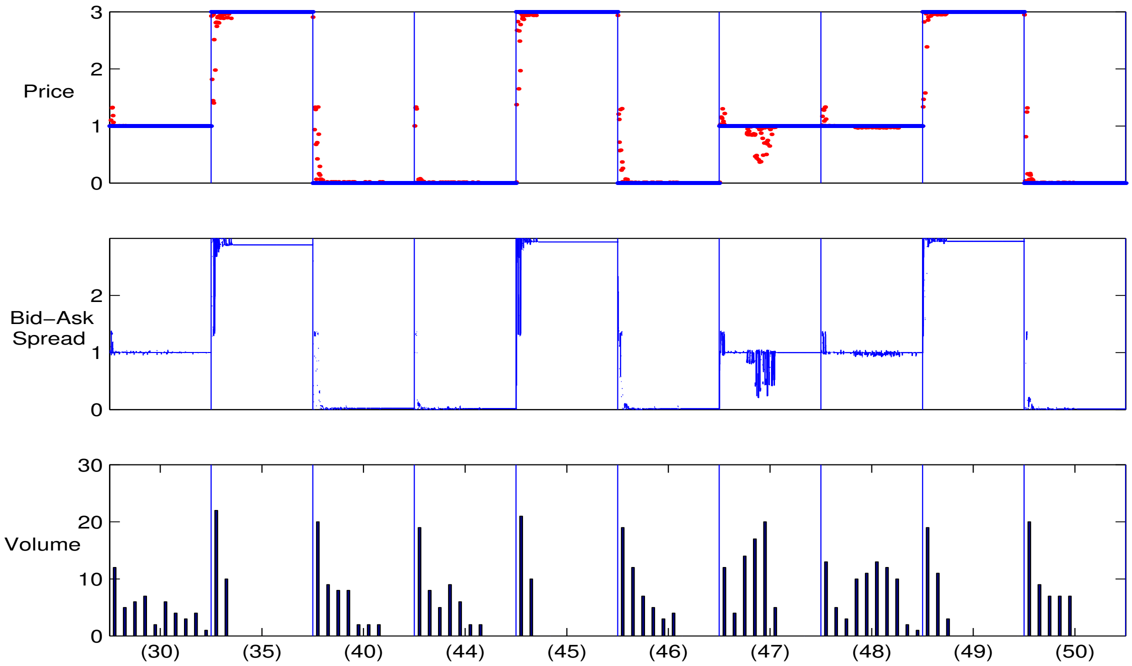

Figure 7: Information dissemination with identical preference. The prices, bid-ask spread and volume in some later periods of a typical experiment.

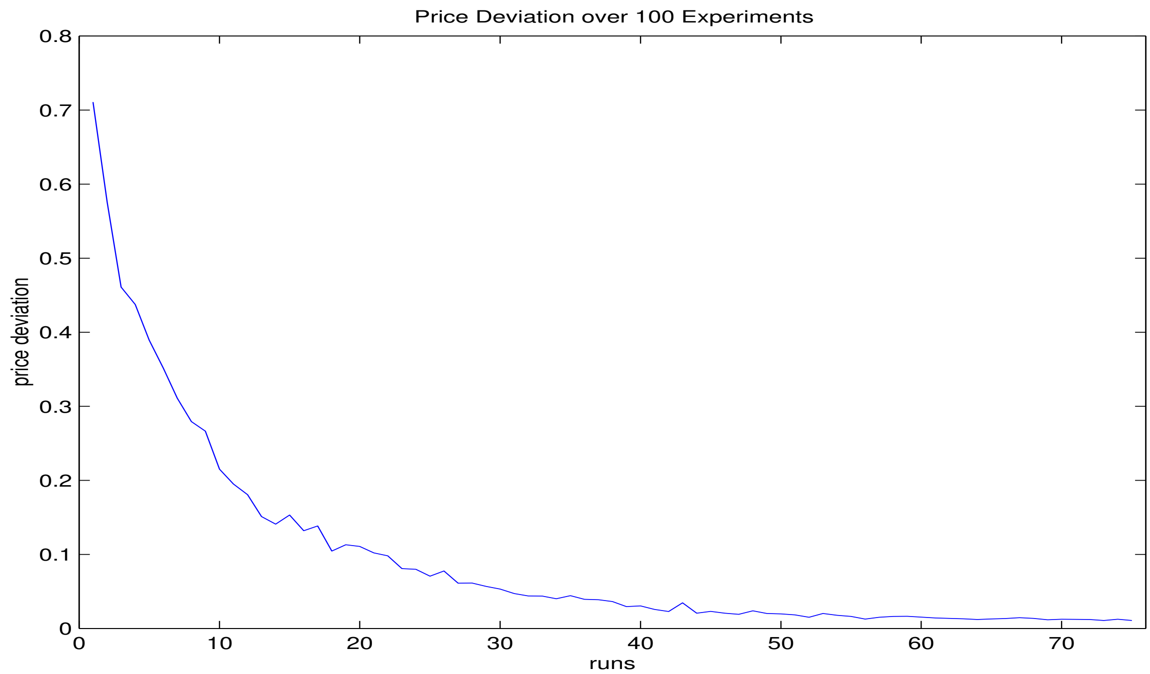

Figure 8: Information dissemination with identical preference. Absolute price deviations from the theoretical price averaged over 100 experiments.

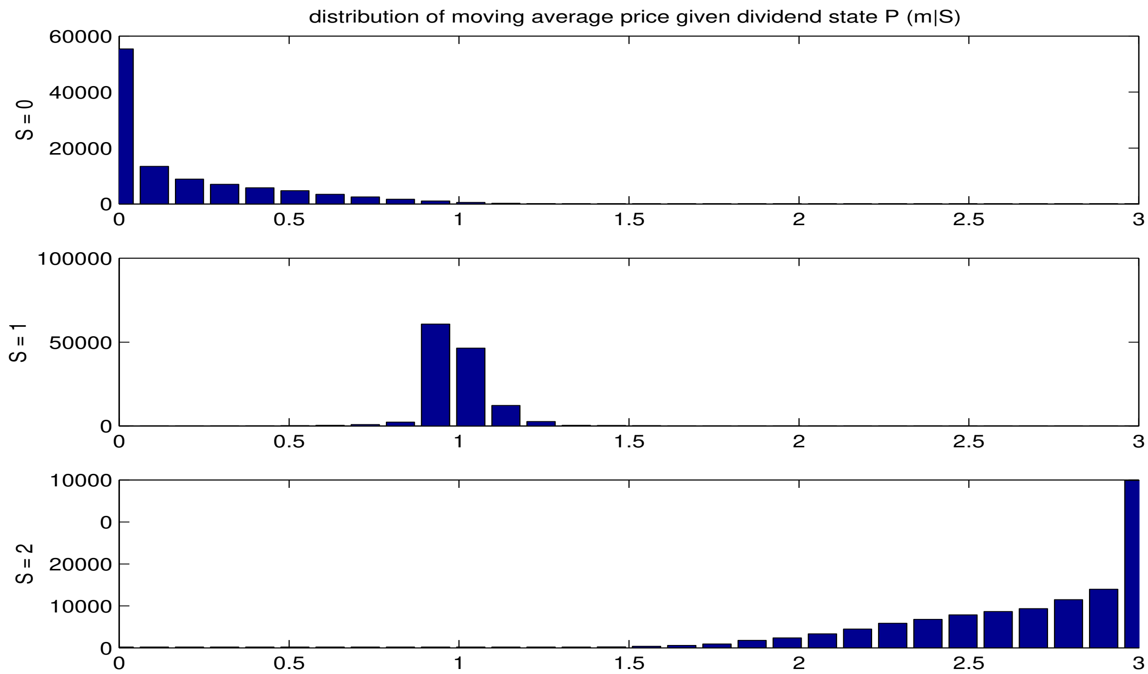

Figure 9: Information dissemination with identical preference. Distributions of prices given the state of nature.

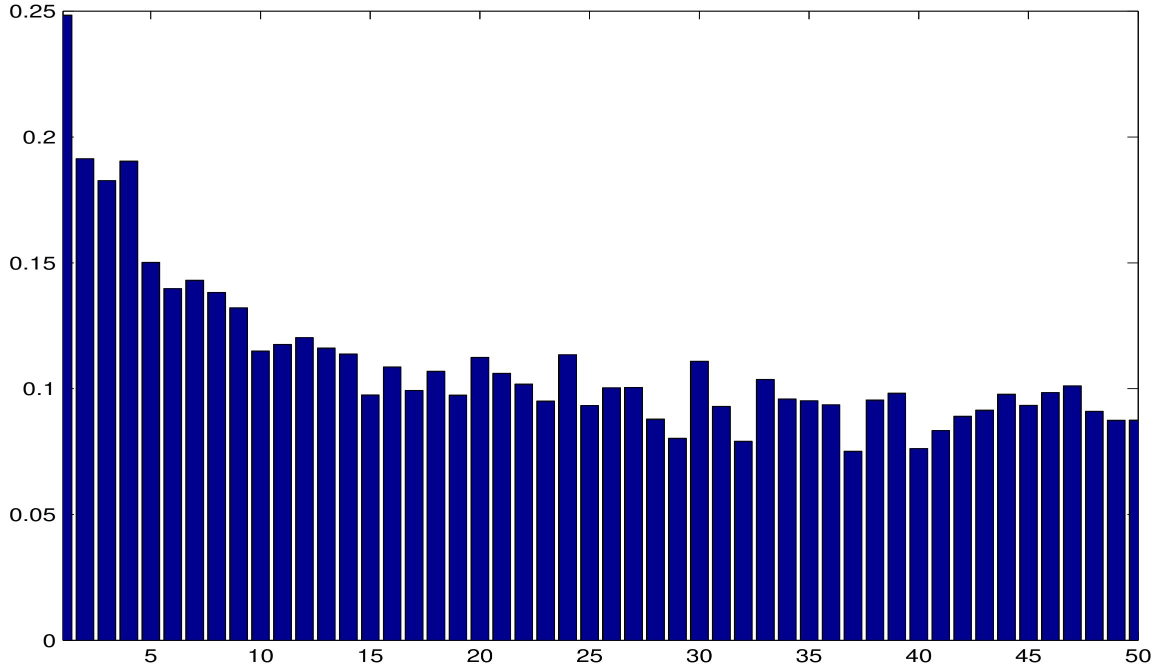

Figure 10: Information dissemination with identical preference. The figure shows the differences in wealth between insiders and uninformed traders averaged over 100 experiments.

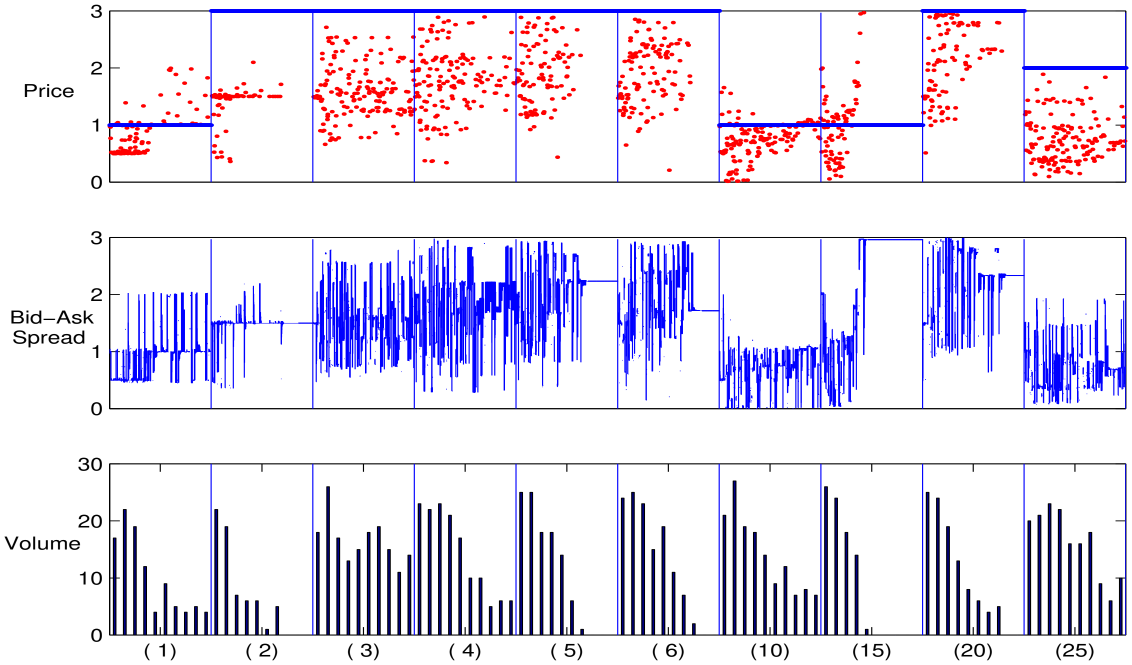

Figure 11: Information aggregation with diverse preferences. The prices, bid-ask spread and volume in some earlier periods of a typical experiment.

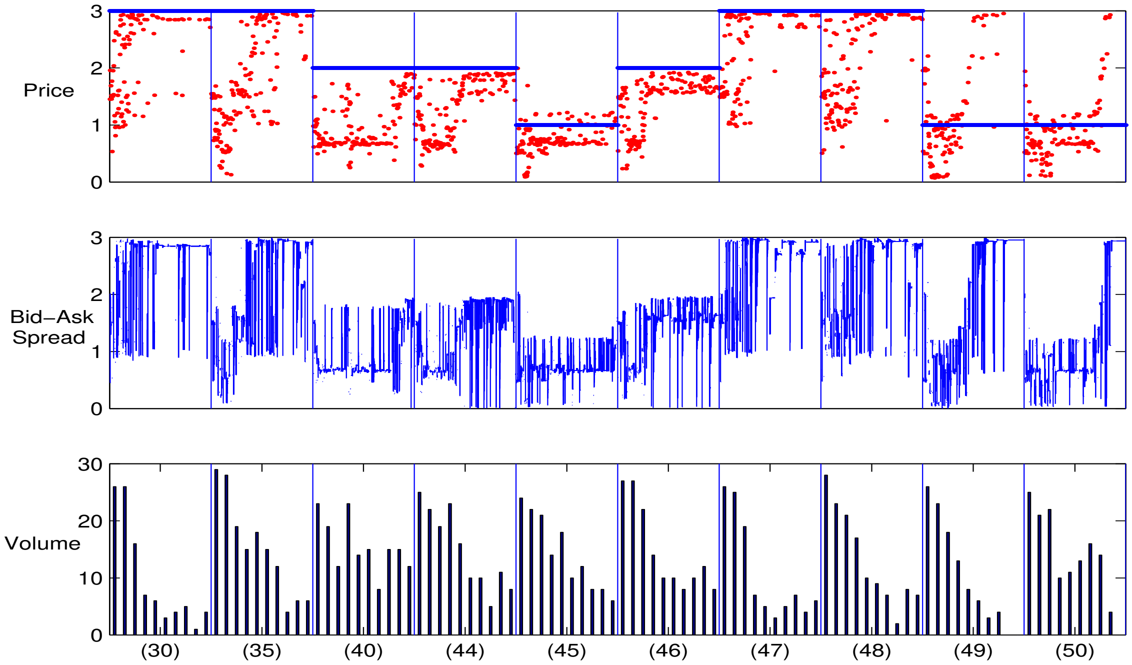

Figure 12: Information aggregation with diverse preferences. The prices, bid-ask spread and volume in some later periods of a typical experiment.

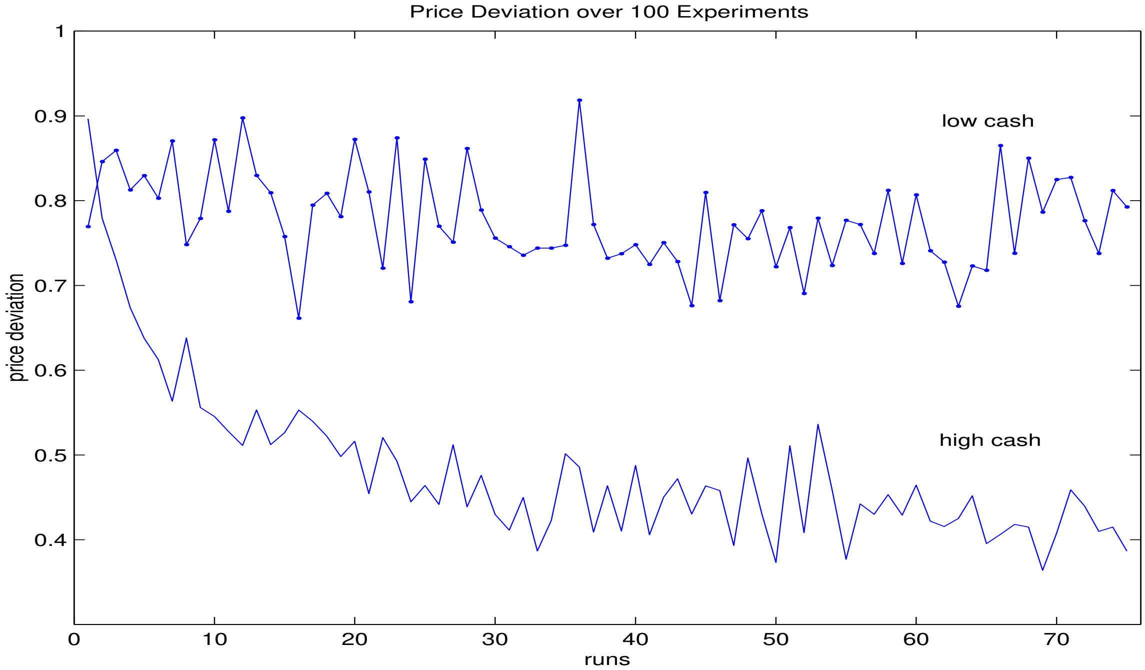

Figure 13: Information aggregation with diverse preferences. Absolute price deviations from the theoretical price averaged over 100 experiments.

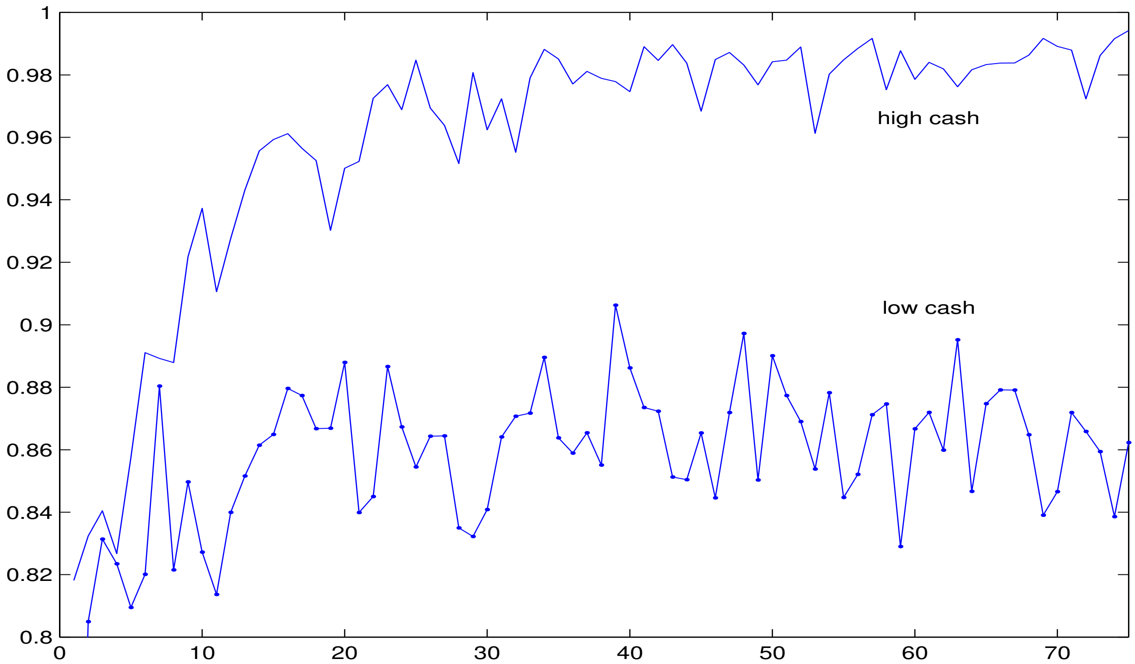

Figure 14: Information aggregation with diverse preferences. Allocative Efficiency averaged over 100 experiments.

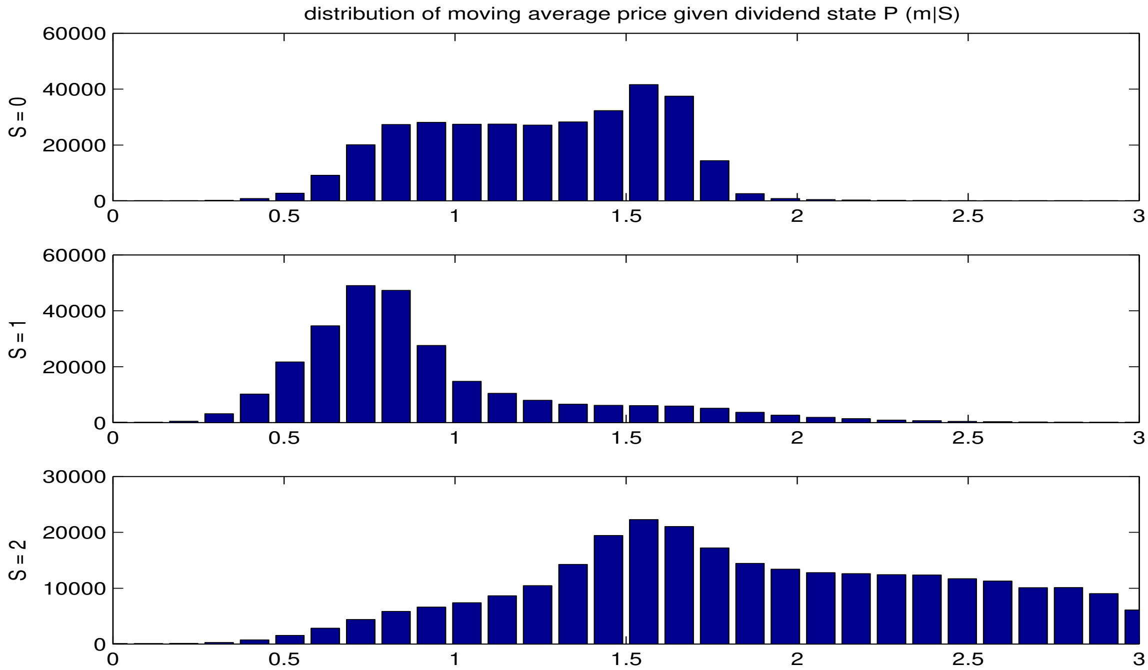

Figure 15: Information aggregation with diverse preferences. Distributions of prices given the state of nature.

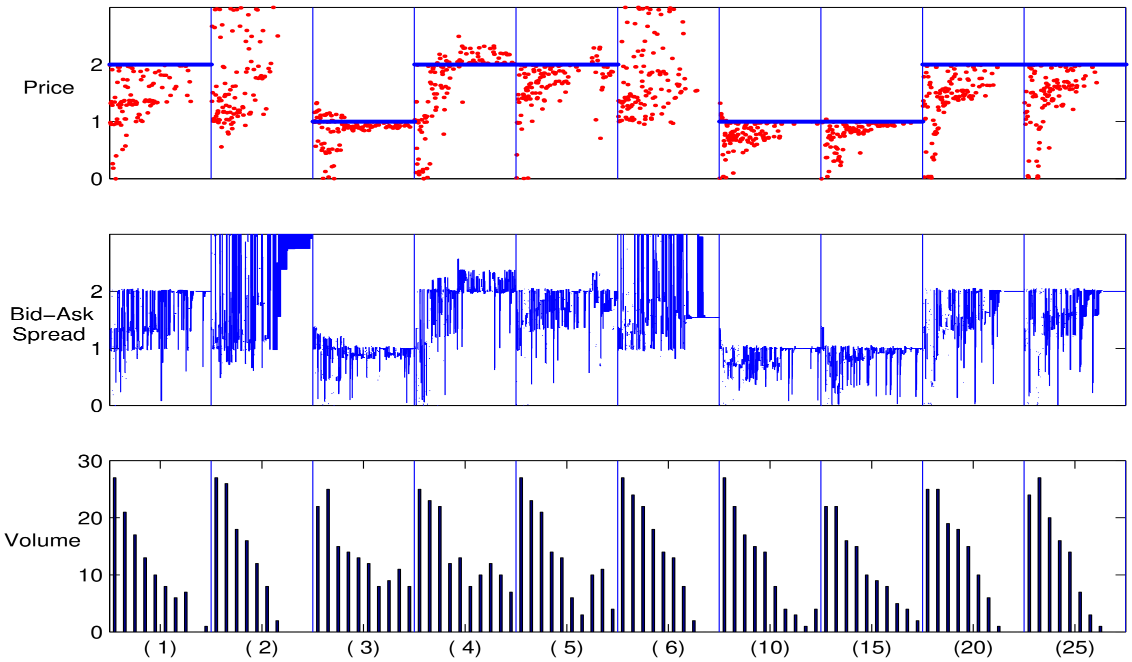

Figure 16: Information dissemination with diverse preferences. The prices, bid-ask spread and volume in some earlier periods of a typical experiment.

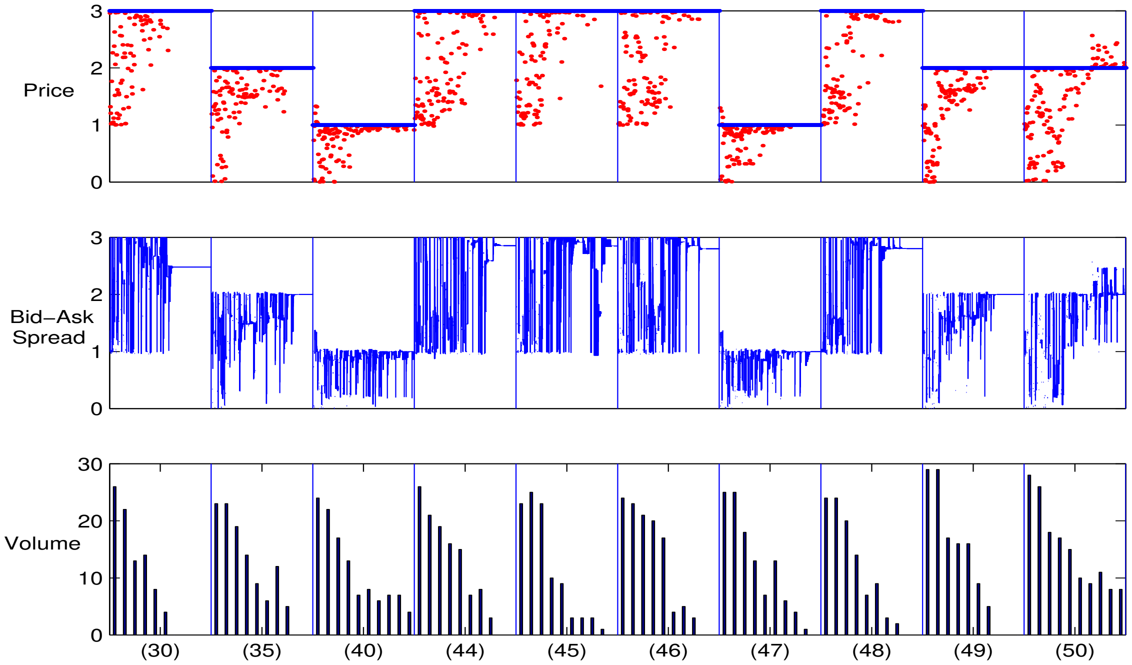

Figure 17: Information dissemination with diverse preferences. The prices, bid-ask spread and volume in some later periods of a typical experiment.

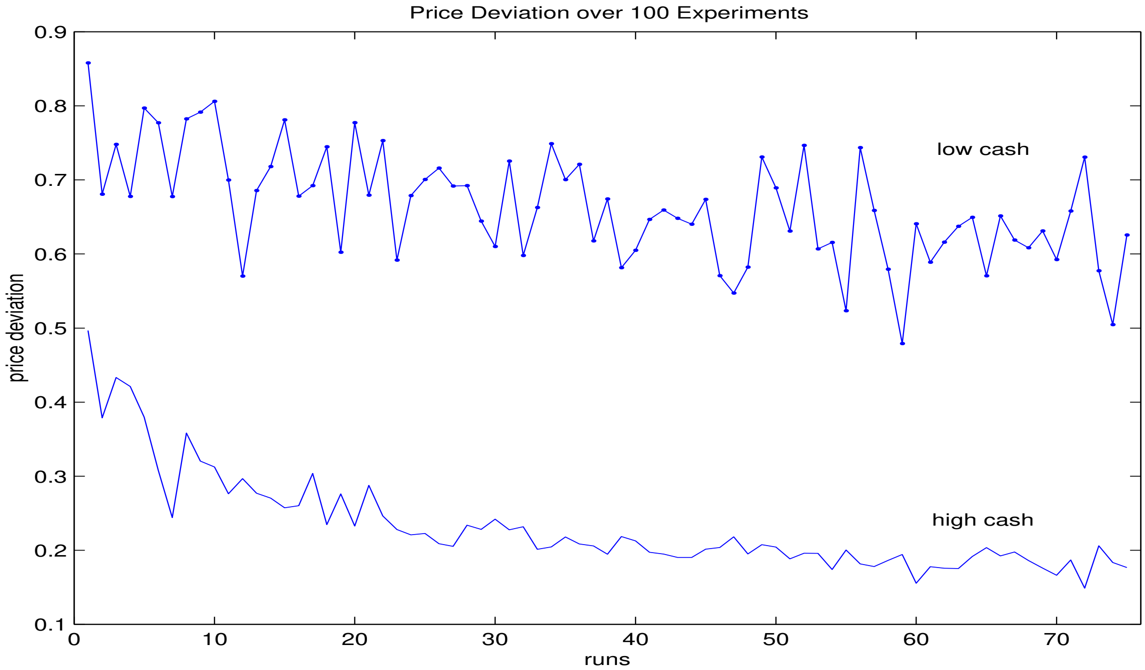

Figure 18: Information dissemination with diverse preferences. Absolute price deviations from the theoretical price averaged over 100 experiments.

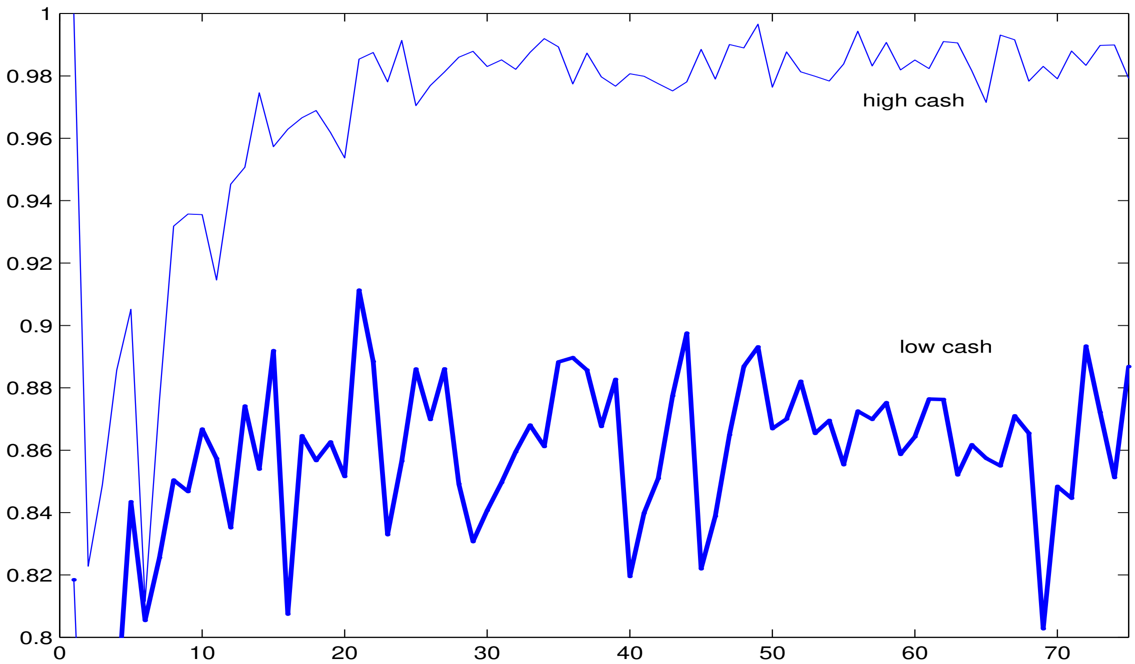

Figure 19: Information dissemination with diverse preferences. Allocative Efficiency averaged over 100 experiments.

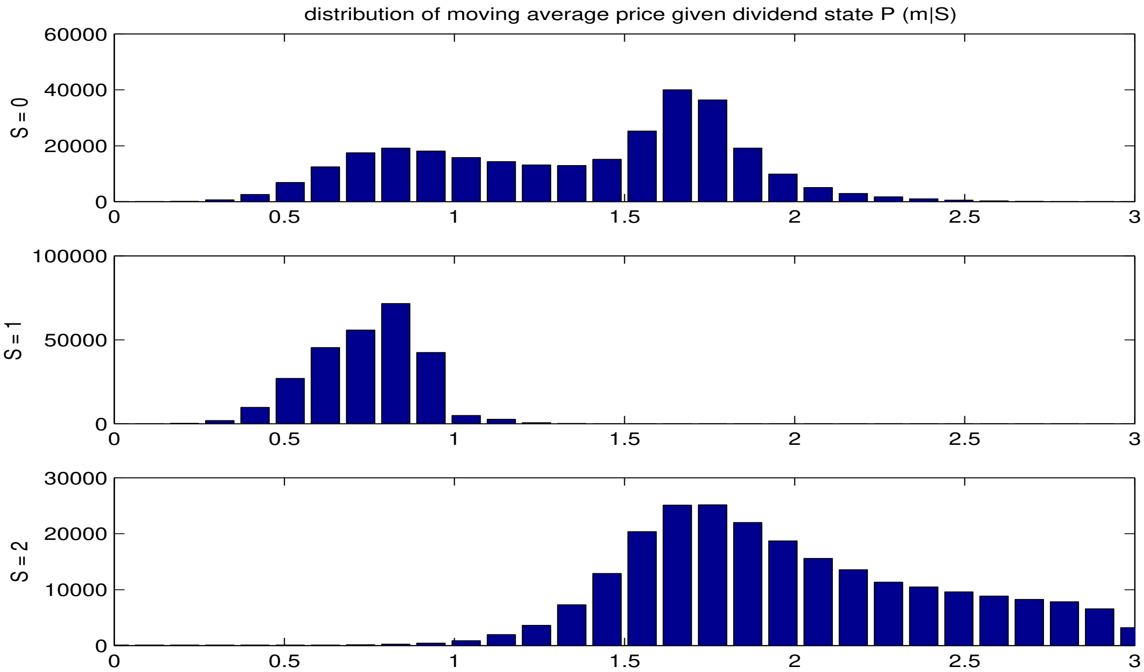

Figure 20: Information dissemination with diverse preferences. Distributions of prices given the state of nature.

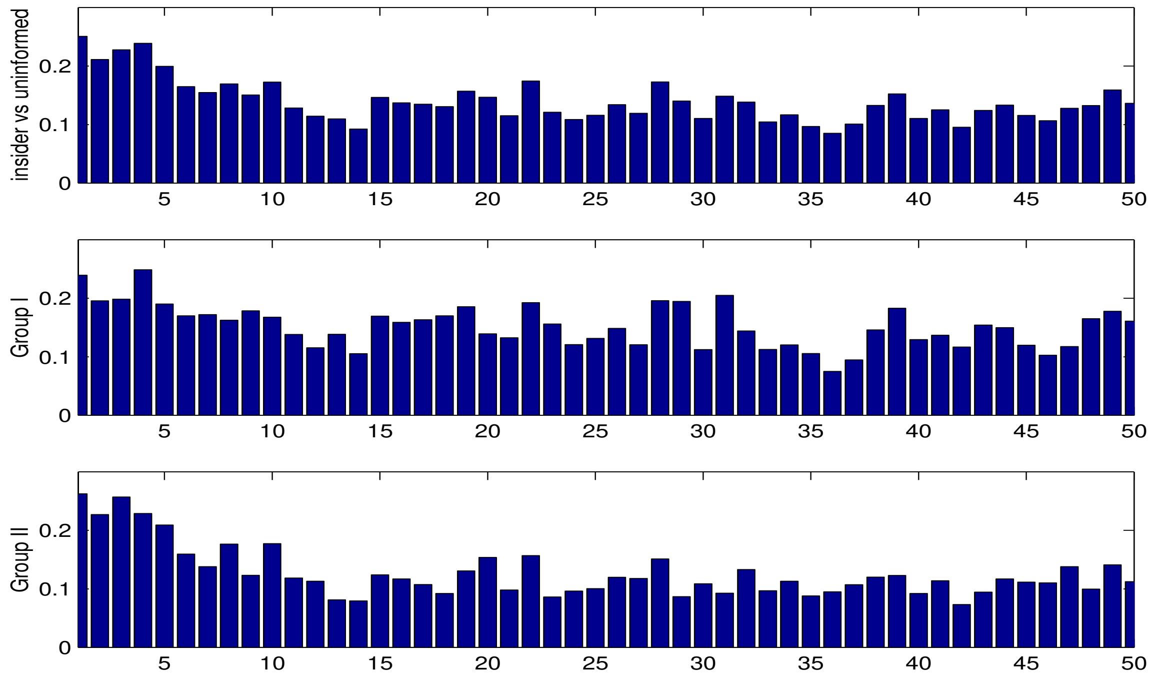

Figure 21: Information dissemination with diverse preferences. Top: The figure shows the average difference in wealth between the insiders and uninformed traders after 100 experiments. The first panel compares the wealth of insiders and uninformed traders in both group I and II; the second panel shows wealth difference in group I; the third panel shows wealth difference in group II.

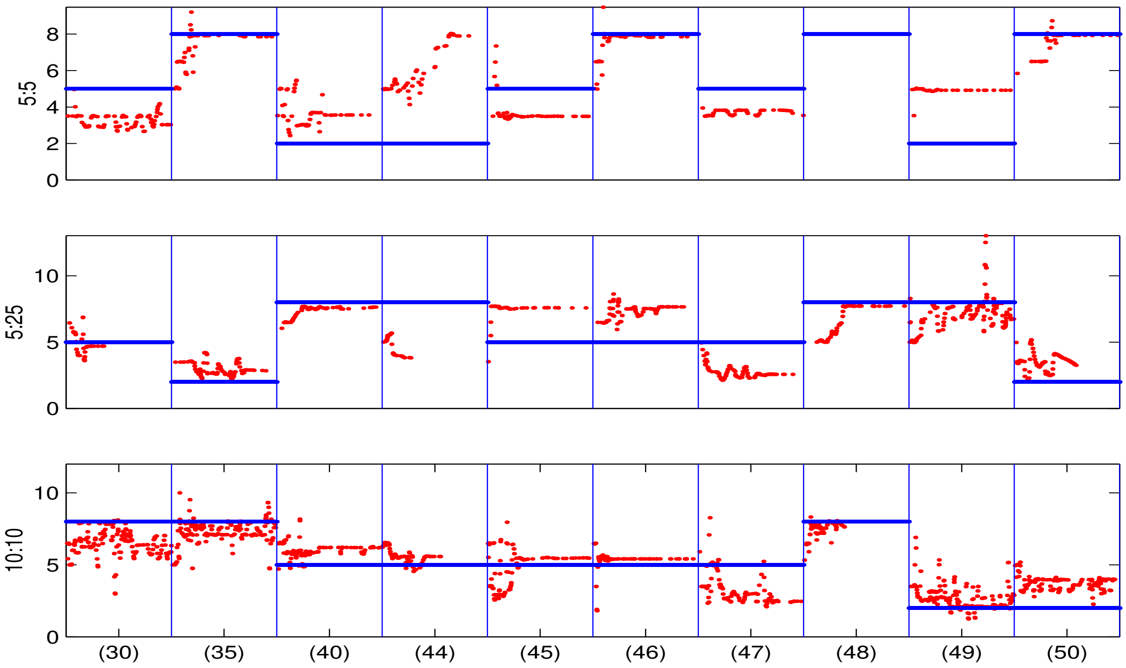

Figure 22: Information aggregation by empirical Bayesian traders in the presence of momentum traders. The prices, bid-ask spread and volume in some later periods of a typical experiment.

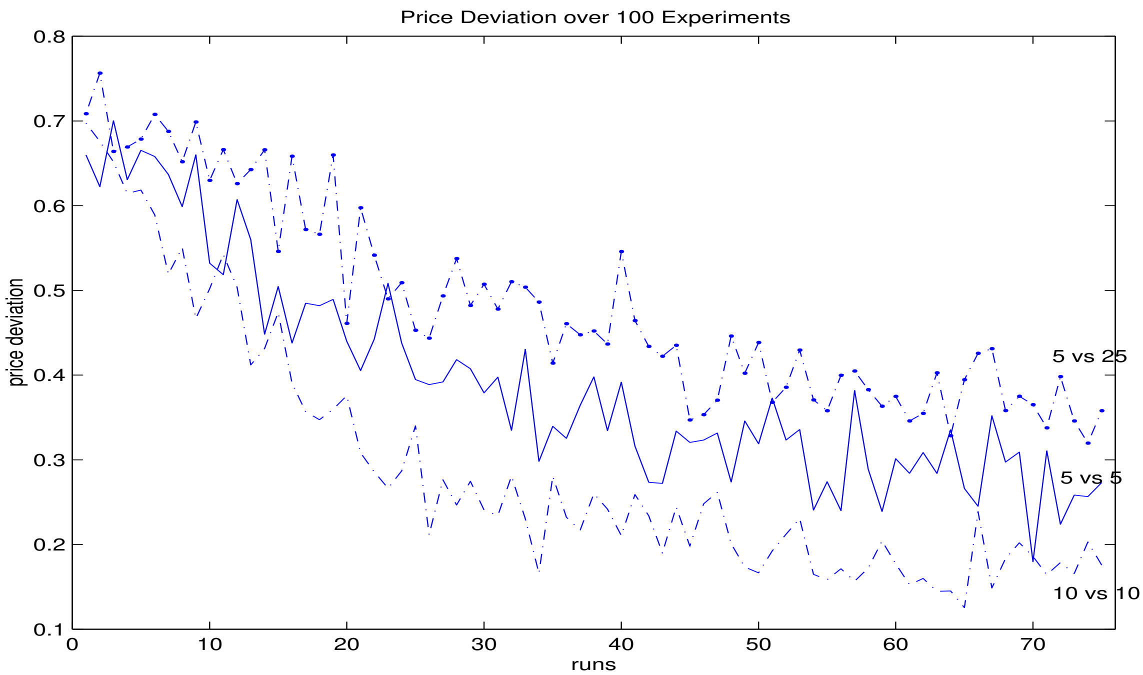

Figure 23: Information aggregation by empirical Bayesian traders in the presence of momentum traders. Absolute price deviations from the theoretical price averaged over 100 experiments.

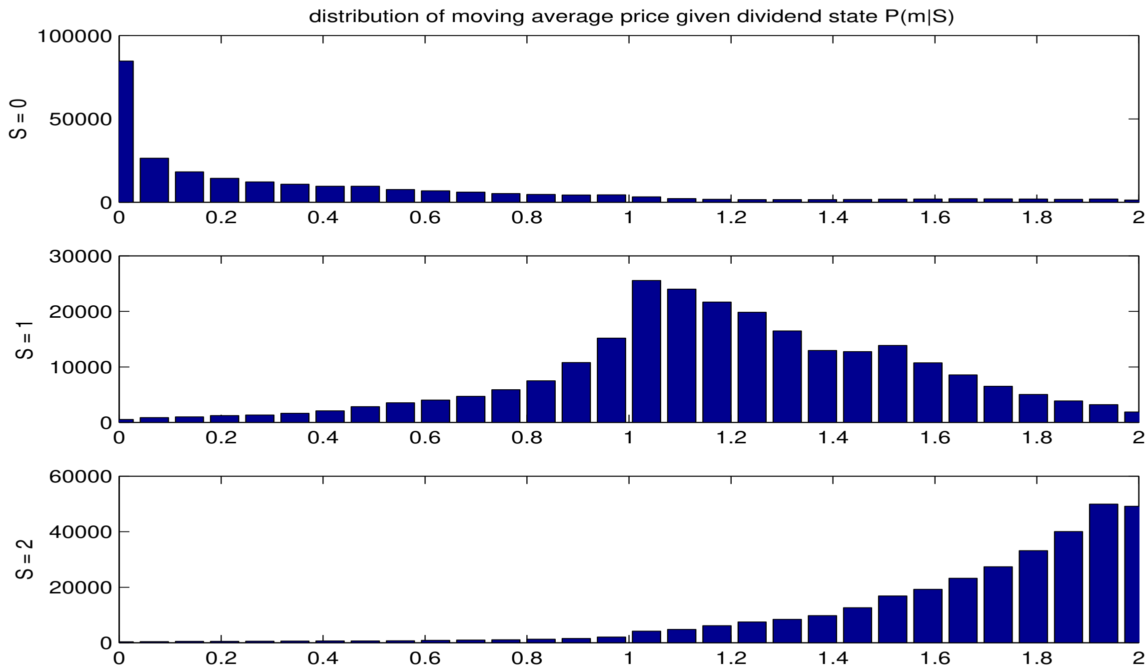

Figure 24: Information aggregation with diverse preferences. Distributions of prices given the state of nature.

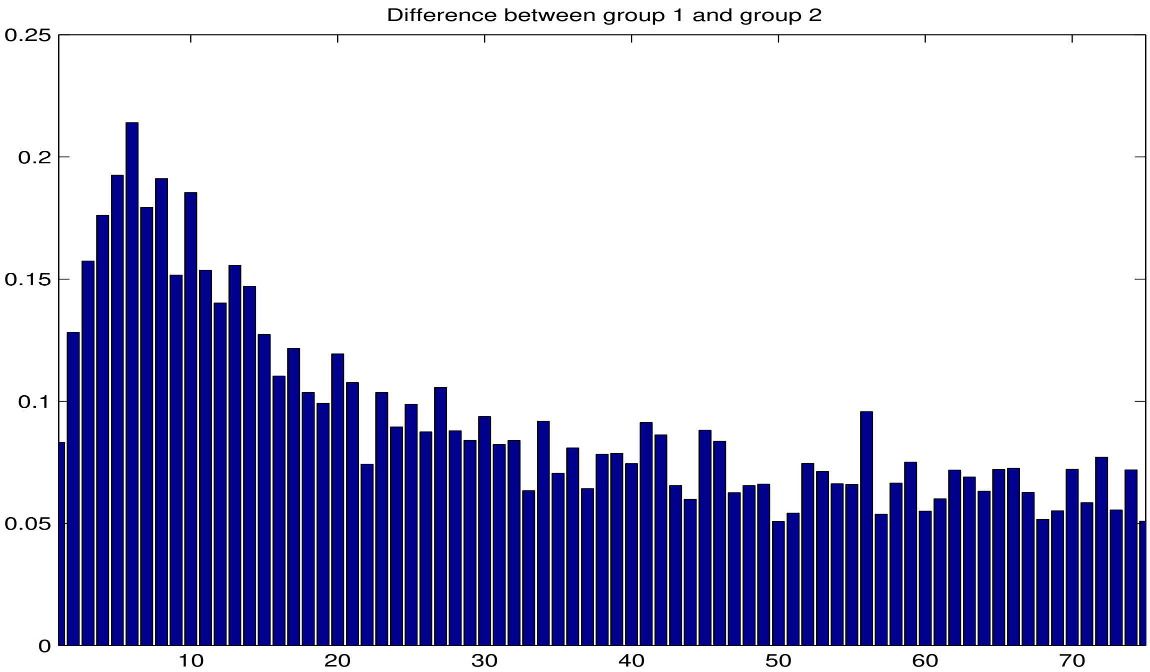

Figure 25: Information aggregation by empirical Bayesian traders in the presence of momentum traders. The figure shows the differences in wealth between empirical Bayesian traders and momentum traders averaged over 100 experiments.