Market Efficiencies and Drift: A Computational Model

John Dickhaut
Emeritus, University of Minnesota
Professor of Accounting and Economics, Chapman University
Economic Science Institute
jdickhau@chapman.edu

Baohua Xin
Assistant Professor
Rotman School of Management
University of Toronto
bxin@rotman.utoronto.ca

Current Draft: October 2008

Keywords: market efficiency, drift, computational model

Data availability: Data are available upon request

* We have benefitted from discussions with Sudipta Basu, Joyce Berg, Peter Bossaerts, Amy Choy, Sheri Dickhaut, Murray Frank, Frank Gigler, Kareem Jamal, Bruce Johnson, Chandra Kanodia, Jack Kareken, Tom Rietz, Jack Stecher, Greg Waymire, Rick Young, and workshop participants at Emory University, Purdue University, the University of Iowa, the University of Minnesota, the University of Pittsburgh, the University of Wisconsin at Madison, Norwegian School of Economics, the Economic Science Association, the Southern Economic Association, the American Accounting Association 2006 Annual Meeting and EIASM Managerial Accounting Conference.

Abstract

Accounting and finance researchers show semi-strong efficiency or lack thereof by using sequences of prices from CRSP and COMPUSTAT data for which there is not a model of how these prices came about through individual decisions. One needs a setting in which the prices (including bids and asks) as well information about individuals making the choices are both available. To begin to bridge the gap between theory and data we extend work done by experimental economists on the double auction and model price formation that is/is not semi-strong efficient.

I. Introduction

"What individual has chosen prices? In the formal theory, at least, no one. They are determined on (not by) social institutions known as markets, which equate supply and demand. The failure to give an individualistic explanation of price formation has proved to be surprisingly hard to cure."

Arrow, 1994

For 40 years accountants using archival data have sought an understanding of the relationship between stock price behavior and information (and in particular accounting information). In this research the issue described by Arrow is always apparent, namely there is no "individualistic" explanation of how prices in archival data are formed. Accountants show semi-strong efficiency or lack thereof by using sequences of prices from CRSP and COMPUSTAT data for which there is not a model of how these prices came about through individual decisions. The fact that such a puzzle has been unsettled for so long suggests that it will not be resolved completely overnight. One needs a setting in which the prices (as well as bids and asks) and information about individuals making the choices are available. To begin to bridge the gap between theory and data we extend work done by experimental economists on the double auction to model price formation that is/is not semi-strong efficient.

We use a general computational model of price formation with risky assets to give insight into what kind of structure would support many of the (at times apparently conflicting) results reported. It is a model which in principle applies to any setting with a finite set of actors who behave according to any well known risky preference structure (expected utility, prospect theory, rank dependent utility theory) and is void of strong (rational expectations) informational knowledge on the part of actors in the economies. The model is general yet it is computational in that it is not closed form and yields testable propositions when parameters are explicitly specified. In this sense it offers more than just experimental and/or simulation results but rather an entire framework to examine theoretical predictions across a wide set of parametric specifications and contexts.

A brief description of the dynamic process

We describe the economy at an arbitrary point in time, t, in the life of the risky asset. At that point in time there are a finite number of traders, n. Each trader has an endowment of the risky and a riskless asset. The traders' risk attitudes are all different. Traders may also differ in their information; some traders may be better informed than others. The traders are not assumed to have any direct knowledge of the other types of traders in the environment at any particular point of time.

We model how a sequence of trades comes about in a double auction. In a double auction agents post bids and asks. Such bids and asks can start at any amount (here we assume that agents will confine their amounts between the high and low payoffs.) The critical feature of the auction is that the bids and asks are subject to an improvement rule. That is, if there is an outstanding bid and ask at point in time t, agents may only submit a $t+1$ bid that is higher or a $t+1$ ask that is lower than the outstanding bid or ask. The auction processes only one possible improvement at a time, and a legitimate bid/ask submitted by the first agent will become the new outstanding bid or ask at $t+1$. A trade occurs when a bid equals the outstanding ask or vice versa. If a trade does not occur, all agents must decide what to do at $t+1$ in an attempt to be the outstanding bid/ask at $t+2$.

At any point in time, besides having some information on the payoff of the risky asset, an agent also sees the history of previous bids and asks and how many bids and asks were or were not taken for each ask. There are two types of information processing that a market participant engages in. First he analyses the implications of his own information on the value of a traded asset. Secondly he analyses the implications of publicly available prices, bids and asks to determine his action choice in the market. That action can be either making a bid or ask or taking an outstanding bid or ask based on his assessment of value. Because agents can have different risk attitudes as well as information, agents can have different valuations of the risky asset. When there are different valuations there is potential for an advantageous trade between two agents if the price agreed upon is between the valuations. An agent need not bid/ask his valuation; however if there is an opportunity to trade at an advantageous price (i.e. buy at a price less than valuation or sell at a price greater than valuation), the agent will try to take advantage of the opportunity.

When each agent has decided precisely what to do there will be an expected profit from a potential trade. Traders are more motivated to consummate trades the higher their expected trader profit. Thus they will rush to take their action faster the higher their expected trading profits. However another trader may get to the market first. The chances any particular trader takes his action first are a monotonic function of the ratio between his expected profits and the sum of all traders expected profits. Once an action is taken by a trader the market incorporates that action by either executing a trade or posting a new ask or bid. Then the same process is repeated.

Background Literature

Experimental economics has spawned an extensive literature the aim of which is to understand how individuals behave in double auctions. These auctions have been objects of study because they bear resemblance to real world trading institutions such as the NYSE and the Chicago Mercantile Exchange. There are some limits on this interpretation and we shall discuss these limits in our conclusions. Fundamental work on this institution was conducted in a series of experiments summarized by Plott (1982) and Smith (1982) which demonstrated that a competitive equilibrium (prices, quantities, and consumer surplus) can be achieved by prices formed in the double auction.

Results on the double auction experiments did not come without an accompanying puzzle: "Why did the prices arrive where they did?" After all, subjects had only limited information; they did not know the private valuations of other traders, nor the equilibrium price, but only the rules of the game and their own valuations. The work of Easley and Ledyard (1993), Cason and Friedman (1996) and Wilson (1987) provided theoretical insight into how this convergence might occur. Gode and Sunder (1993) provided a dramatic example of how using a very simple computer representation of each individual agent (“a zero-intelligence agent”), it was possible for such agents to arrive at the overall competitive equilibrium consumer surplus. Such agents were not concerned with any historical information about previous behavior in the market so that prices did not necessarily inform them what the next observable price would be. Gjerstad and Dickhaut (1998) explored how the historical information in the markets (bids, asks and trades) might be used by their computerized agents to form beliefs in a way that could produce price paths like those which are observed in classic double auction experiments. This latter approach of asking how modeling individual agents could lead toward understanding observables such as price paths in markets is the one adopted here. The approach can be used to formulate theory and from this theory inform researchers how potential new experiments could be designed. This confluence of what are often referred to as agent based models and experimentation is nicely elaborated in Duffy (2004).

In related literature Grossman (1976, 1978) and Grossman and Stiglitz (1980) show that equilibrium exists only if there is noise in the price system which prevents traders from inferring perfectly the information from prices, for example, shocks in the supply of the risky asset. They focus on the equilibrium (or disequilibrium) analysis – all traders move simultaneously and price incorporates all the available information instantaneously. There are several significant differences in our approach. First we model the dynamics of price/bid/ask formation instead of focusing on the equilibrium results. Secondly we use the double auction instead of Walrasian auctioneer, the latter being far removed from the auction structures that are studied in both laboratory and archival studies. Lastly, the traders in our model are heterogeneous (in terms of preferences) and there is no shock in the supply of the risky asset. In the Grossman and Grossman et. al. models, there is no trade if everyone knows the revealed information when it is revealed, while in ours, risk sharing will always allow trade to take place.

Marcet and Sargent (1989a, 1989b) study a class of linear stochastic models in which the actual law of motion depends on the agents' perceived law of motion. They show the convergence of recursive least squares learning schemes to a rational expectations equilibrium. Marcet and Sargent use a description of how beliefs about the relationship between choices and outcomes are arrived at by repeated processing. Through repeated processing of information and allocations based on that processing the predicted law of motion becomes the true law of motion.

In our model beliefs are learned but in a direct form. This learning involves examining prices, bids and asks and building a conditional set of likelihoods a particular bid or ask will be taken based on a particular mechanism.

Gjerstad and Dickhaut (1998)'s agents traded in riskless assets (i.e. assets with a constant payout) yet many archival market studies based on prices from the New York Stock Exchange are concerned with how risky assets are priced, particularly in response to the arrival of information. Fama (1970) distinguishes between semi-strong and strong form efficiency. Under semi-strong (strong) form efficiency profitable trades cannot be achieved by individual agents after the release of public (private) information. Such efficiency is often understood to be a natural consequence of the behavior of agents whose expectations are rational. When observed in archival data the returns after information is released are flat and in numerous instances it is impossible to build advantageous trading algorithms on such prices. Figure 1 taken from Fama, Fisher, Jensen and Roll (1969) (FFJR) shows the positive change in cumulative residuals (excess returns) up to the point of announcement of a stock split (date 0) for a portfolio of firms and a nearly flat curve thereafter. A stock split is public information and an interpretation of the residuals before date 0 is that their change reflects private information coming into the market while the behavior of the residuals after time 0 indicates there are not advantageous trades to be made, hence semi-strong form efficiency.

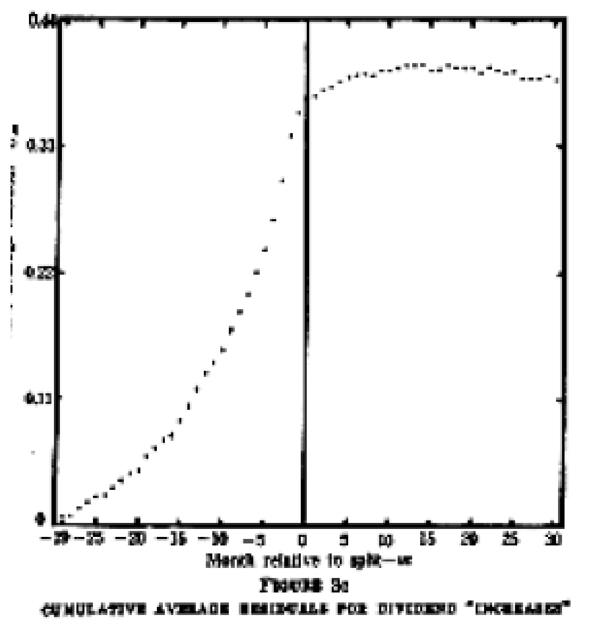

Figure 1 FFJR Semi-Strong Market Efficiency

We show, by adapting the Gjerstad and Dickhaut trading algorithm to trade in a risky asset, how it is possible to produce prices which are semi-strong efficient but not strong-form efficient.

Figure 2 is taken from Ball and Brown (1968) and shows the positive (negative) abnormal performance index up to the point of earnings announcements (date 0) for various portfolios of firms and relatively flat curves thereafter. An earnings announcement is public information and an interpretation of the residuals before date 0 is that their change reflects private information coming into the market while the behavior of the residuals after time 0 is consistent with a notion of a mild drift.

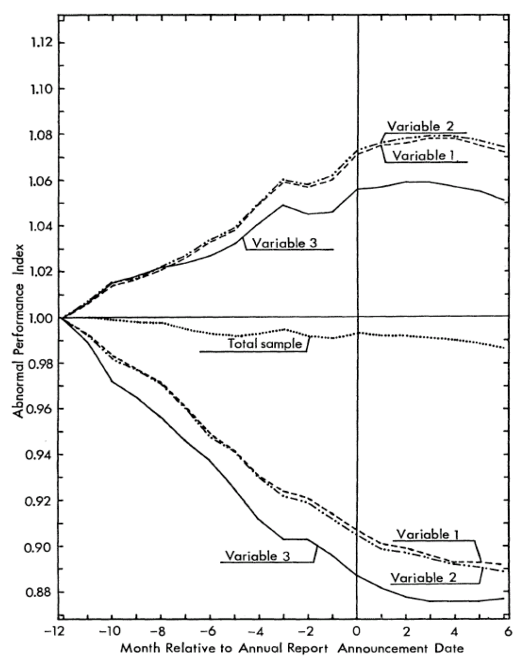

FIG. 1 Abnormal Performance Indexes for Various Portfolios

Figure 2 Ball and Brown Drift

Point 0 is when public information in the form of the earnings announcements arrives for a collection of firms. Each line represents different cumulative average residuals for a collection of forecast/earnings combinations. For example, bad forecast followed by bad earnings is the last plotted line in the diagram. As can be seen the drift after announcement is not independent of the type of information generated. This particular study does not have an explicit theory associated with it explaining a drift. We are able to show that this type of drift can be well approximated by the interaction of agents in our model of price formation.

II. A Description of the Model

Figure 3 describes the four basic steps of the model. In step 1 each agent determines his value. In step 2 the agent uses market information to assess the likelihood of being able to trade. In step 3 the optimal action of the agent is discussed and in step 4 the stochastic arrival of agents to the market is examined. We examine each step in detail.

Step 1: Valuation

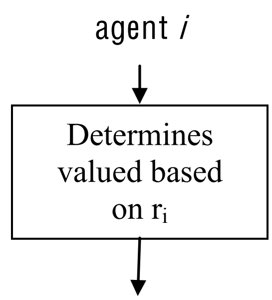

Every trader comes to the market with a specific risk attitude. Not all traders have the same risk attitude. To make this assumption specific we will describe a risk attitude via a risk parameter later on. But there is no restriction conceptually about how the description of risk is arrived at. A trader will use the information he has to arrive at his assessment of the reservation value of the traded asset (i.e. the lowest market price at which he would sell and the highest market price at which he would buy). The method for making this assessment could be using any type of decision making model such as expected utility theory, prospect theory or rank dependent utility theory[1]. The actor may even choose to include previous market price in the

[1] To implement the structure we will assume this assessment is made by computing expected utility theory given available information and then deriving the related certainty equivalent. We derive here the implications of having ten agents with different risk coefficients and who also have the same underlying CARA utility function. The approach can be extended to a general two state model where we assume two payoffs and a reservation value. Note that any structure that has a reservation value for the gamble will also work. For simplicity we employ a model which makes the calculation of the reservation value a direct function of payoffs and risk characterization although there is no reason to make the model this specific. It is always assumed that the reservation value will be in between the high and low payoffs. It may be that the agent has two reservation values, one related to being a seller, and the other related to being a buyer. As long as the agent is still interested in the difference between the reservation values and the price received in the market the structure can work out in the same manner.

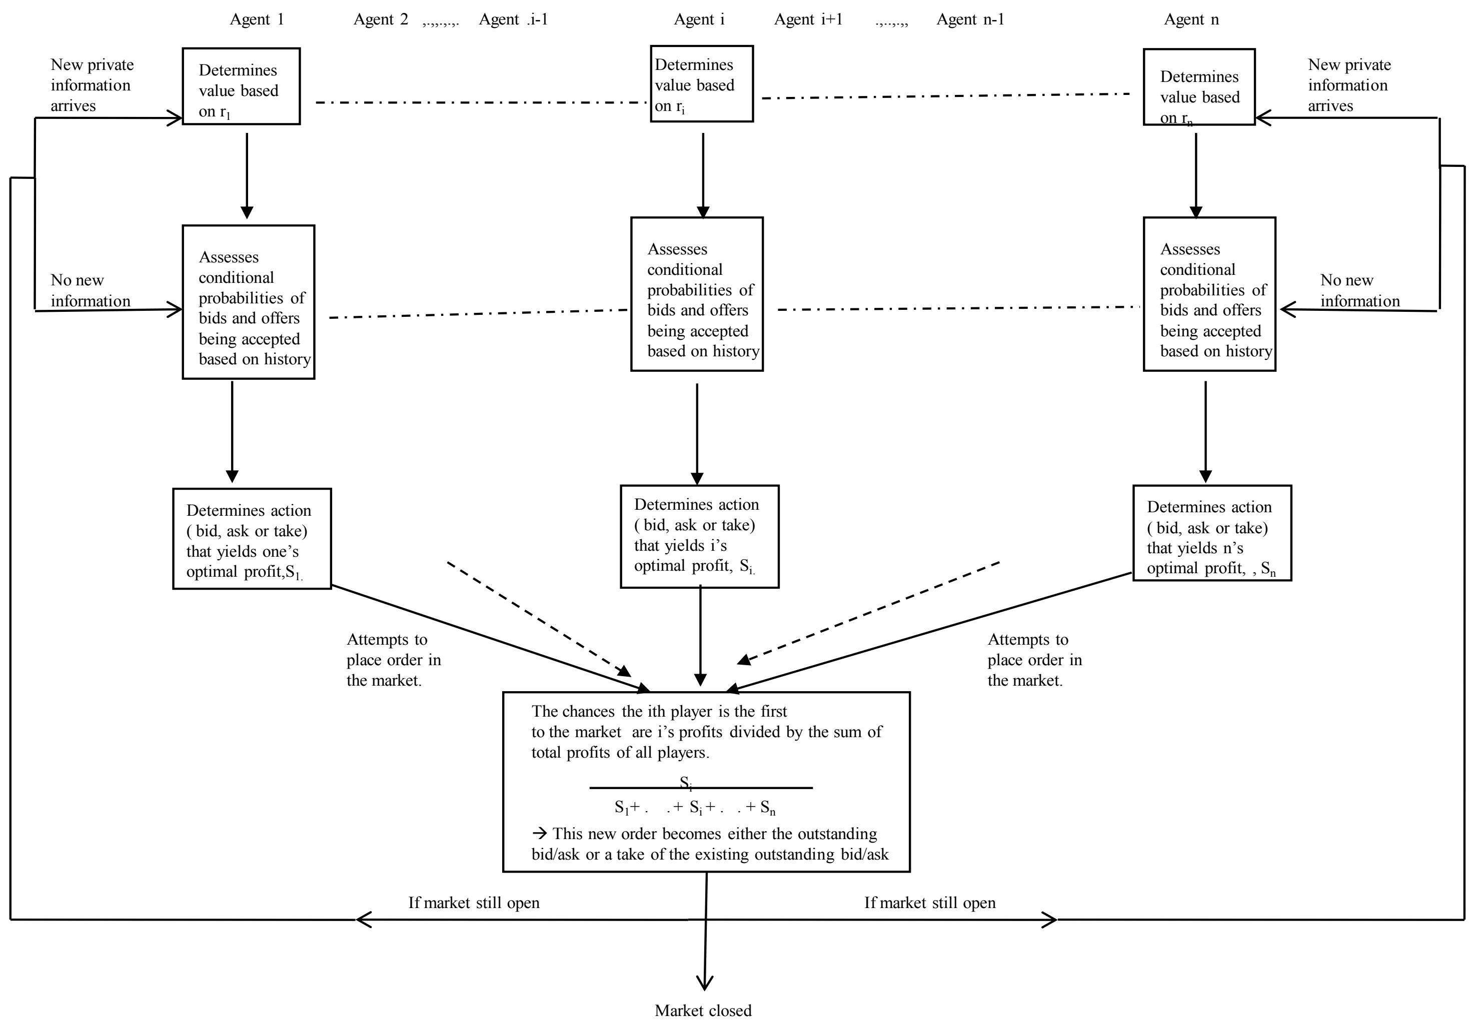

Figure 3 Layout of the Model

formulation. Note again as long as certainty equivalents can be stated for each probability payoff pair (and wealth combination) the theory works the same way although the implementation becomes more challenging. Because traders have different risk attitudes trade is likely to occur.

A trader in the market may or may not have personal private information (i.e., either be informed or uninformed). A trader's private information could in principle derive from word of mouth or personal observation. Information is not purchased. It lands on the trader able to be at the right place at the right time. Informed traders do not know how many other traders are informed. All informed traders share the same beliefs about the dollar outcome from the asset when computing their estimate of value. This is implemented by assuming they use Bayesian revision relative to the arrival of information.[2] Uninformed traders will have processes for generating beliefs that will include some portion of all available public information. Uninformed market participants do not know when informed market participants receive information. Because prices, bids and asks arise both from new information to informed market participants as well as risk sharing considerations, uninformed traders will be unable to disentangle these effects and base their assessment of value only on released public information. We implement this assumption by assuming that uninformed market participants use Bayesian revision relative to only the arrival of public information.

In this first step agents are just determining their value but they still must decide in steps 2 and 3 what is the best action they can take in the market (i.e. appropriate bid, appropriate ask, or a take).

Step 2: Agent computations of likelihoods

An agent assesses the conditional probabilities of bids and offers being accepted based on history.

[2] Note that pure Bayesian is not absolutely necessary for the general structure of the model. An agent could differentially underadjust or overadjust. This characteristic of course would need to be built specifically into the adjustment process of the model. See Robustness Section for other ways of updating.

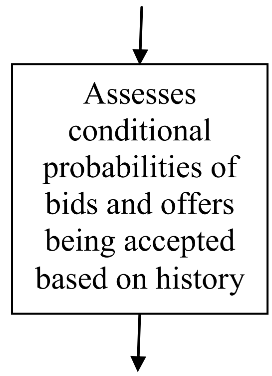

In this stage the agent is gathering additional inputs for the decision regarding what action to take in the market (bid, ask or take of an available bid or ask outstanding). The agent will use the certainty equivalent of his computed value from the previous stage, as well as the available market information (prices, bids and asks). At any point in time, t, there is a set of messages. Each of these messages was sent in the form of a bid, ask or taken bid or taken ask by a specific agent. The message at time t will reflect an action by one of the agents in the economy, and that action will maximize that agents expected surplus. [3]

The additional information the agent computes at this stage will be the likelihood a particular bid or ask will be taken by someone else in the market. In absence of differential information traders still do not know what types of people are in the market and so engage in trades based only on available bid, ask and trade information and attempt a pseudo arbitrage based on that information and their private knowledge of their own risk. This is stimulated by the ability to retrade. While this pseudo arbitraging is going on differential information can arrive to traders. There will be haves and have nots. People with new information revalue, people without do not revalue but keep their old value. These types do not know anyone else received new information and what that information might be. They infer from price the advantageous trade relative to that price. We assume agents have limited memory.

At this stage while a trader knows his value of the asset he knows nothing about the risk distribution of other trader types nor whether the most recent market observations were a

[3] The timing of messages is independent of the type of risk preferences the agent has for representing the basic gamble as well as the belief revision policy.

consequence of differential information or risk sharing. Thus every trader will base his assessment on the likelihood of a particular bid (ask) being accepted on recent observed information in the market about whether that bid (ask) would be successful or unsuccessful. We assume because he simply has such limited information and no basis for assigning priors the only information that he can base his action regarding a bid or ask on is in the recent prices, bids, and asks. He then resorts to a simple counting system (which can be unconsciously done by the normal brain). For example for a specific dollar bid, B, the likelihood it will be taken is based on:

(1) the number of times this specific dollar bid has been successfully taken, $\#BT$;
(2) the number of times an ask below this bid has been made, $\#A<B$; and
(3) the number of times a bid above this bid has not been taken, $\#BNT \geq B$.

Then the likelihood is computed using the following formulation

$$\frac{\#BT + \#A < B}{\#BT + \#A < B + \#BNT \geq B}$$

A similar counting approach applies to the calculation of the likelihood of being taken for a specific dollar ask, A: $\frac{\#AT + \#B > A}{\#AT + \#B > A + \#ANT \leq A}$ where

(1) $\#AT$ is the number of times this specific dollar ask has been successfully taken;
(2) $\#B>A$ is the number of times a bid above this ask has been made; and
(3) $\#ANT \leq A$ is the number of times an ask below this ask has not been taken.

Why does a trader use such a strategy? Given the paucity of information in the environment we are suggesting the trader resorts to using primitive mechanisms that have evolved through millions of years. In referring to primitive mechanisms we are referring not only to human evidence but also evidence that has been more precisely gathered with animals other than humans. There is now ample evidence that the brain has the ability to represent numbers as if taken from a line segment and that this ability has existed for 1,000's of years. For example rats are capable of distinguishing the numbers 1, 2, 3, etc. based on whether a reward given when a rat presses a lever exactly 1, 2, or 3 times. Monkeys can also do simple auxiliary mathematical operations such as addition etc. This behavior occurs conjointly with activation in the posterior parietal sulcus and dorsaleral prefrontal cortex of the monkey brain. Humans have a very extensive system for representing numbers inexactly in horizontal intraparietal sulcus, and of course have mechanisms for adding, subtracting, multiplying and dividing. These operations can be carried on by the brain regardless of whether the subject is conscious of performing these operations. Recently Fiorello et. al. (2003) has produced evidence that specific cells of the monkey can be tuned to different conditional probabilities of reward that exist in the environment. Thus we are not assuming that the mechanisms goes on consciously in the mind of the trader but rather they can lead to a trader having a feel for what is the best bid or ask even though the trader may not be aware of where such feelings came from. Gigerenzer (1995) trains subjects to behave in a Bayesian fashion based on variations on the types of counting systems used here although he assumes more explicit information than we have.

Step 3: Determination of optimal action

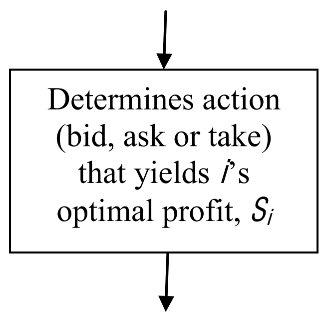

In this stage the agent determines the expected monetary profit for each action he might take[4]. The expected profits for particular agent i for bid B will be $P(B)\,(V_i - B)$, where for an observed B, $P(B) = \dfrac{\#BT + \#A < B}{\#BT + \#A < B + \#BNT \geq B}$; while for previously unobserved B's the calculation for that $P(B)$ is determined by interpolation. It can be proved that the resulting $P(B)$ is non-decreasing in B.

Similarly the expected profits for particular agent i for bid A will be $P(A)\,(A - V_i)$, where for an observed A, $P(A) = \dfrac{\#AT + \#B > A}{\#AT + \#B > A + \#ANT \leq A}$; while for previously unobserved A's

[4] We omit here the construction that would be necessary if the agent had a different certainty equivalent for buying and selling.

the calculation for that $P(A)$ is determined by interpolation. It can be proved that the resulting $P(A)$ is non-increasing in A.

The agent behaves as if doing four calculations

$$\underset{B}{\text{Max}} \quad P(B)\,(V_i - B)\,;$$

$$\underset{A}{\text{Max}} \quad P(A)\,(A - V_i)\,;$$

$V_i - OB$; and

$OA - V_i$

where OB (OA) is the existing outstanding bid (ask).

Determining the maximum profit over these calculations leads to the selection of an action which will be a particular B, a particular A or OB or OA, where the maximum profit is greater than 0. Each subject can make a bid, an ask or take an outstanding bid or offer. Thus he .may be a buyer or a seller. He chooses the optimum of these applying the assessed conditional probability a particular bid or ask being accepted based on history that he derived earlier.

Step 4: Whose bid/ask is posted first – a decentralized process

The chances the ith player is the first to the market are i's profits divided by the sum of total profits of all players

$$\frac{s_i}{s_1 + \cdots + s_i + \cdots + s_n}$$

$\rightarrow$ The new order becomes either the outstanding bid/ask or a take of the existing outstanding bid/ask

To summarize, through the first three steps, each player has arrived at what he would do if he could act first in the auction. That is he knows whether to 1) make a bid, 2) make an ask, 3) accept an outstanding bid or 4) accept an outstanding ask. A player will only take one of these four actions. The expected profit for each action can be computed.

The auction mechanism will acknowledge only one action of only one player. Once this action is incorporated into the message structure of the auction, then all players are notified and the bidding process begins again. What is important is the determinate of who moves in the auction and that we are not assuming the mechanism makes the decision of who moves first.

It is crucial that the auction be decentralized if we are to argue prices are a consequence of individual choice. This means there needs to be a behavioral theory regarding how such decentralization takes place. No single person (or mechanism) knows all of the profits or valuations of each actor. We assume that on average the more profit a single player makes the more likely he will jump into the auction, i.e. higher profits to a player make that player want to act faster. Interpreting this proposition strictly would mean the person with the highest expected profit would always move first. However we assume some noise in the response which could be attributed to the circumstances of individual players. Some players may be closer to a phone to transmit their bids. Some may be paying closer attention. In experiments some may be distracted by something in their environment. Thus our behavioral approach suggests that subjects probabilistically move first the higher their profits. This assumption is captured in the following way. If S_i is the expected profit of the optimal action for the ith player then the likelihood that player i will move first is

$$\frac{S_i}{S_1 + \cdots + S_i + \cdots + S_n}$$

Any monotonic function defined on the S_i will result in a similar construction that will generate noise in trades. The important point here is if traders move to get in the auction based on their profits then the entire process is decentralized.

A general theory of reaction times is that an alternative with higher profits results in arriving at a decision faster when compared with a 0 profit alternative (Doing Nothing). Such an approach can be incorporated into an established theory of reaction time such as Ratcliff's

(See Dickhaut et al. 2008), there it is shown the further away a choice is from *0* profits compared to an alternative the faster the reaction time. Furthermore the notion the higher reward is the faster the reaction time has strong support with respect to primitive brain development. Tremblay and Shultz (1999) have found that specific neurons in orbital frontal cortex not only respond when a reward is received but furthermore respond faster the higher the reward. The latter evidence is suggestive that the speed of response can be determined without conscious awareness and in fact be the determinant of "*as if*" behavior of the monkey and/or the human.

The mathematical formulation of our theory that developed in this section is relegated to the appendix.

III. Results from Computational Economies

Characteristics of all economies

In the computational economies discussed in this section we follow the general hypothesis arising from experimental economics that properties of large scale economies can be approximated with a small number of agents in the double auction (Smith 1962). For our graphical results, we always use either 10 agents or 20 agents. Each agent is endowed with the same number of risky and riskless assets. In the specific economies agents have different risk coefficients but across economies, the collection of risk coefficients is held constant.

Possible states and signals are always the same and priors are held the same on all states in the first three economies. There are two possible public information signals in economies 1 and 2 and the realization of public information is the same across economies. Furthermore the conditional probabilities of public information given the signal are always the same. These features of the economies are characterized in Tables 1 and 2. Economies 3, 4 and 5 are described later.

Table 1 State Characteristics

Possible States	40	80
Prior Probabilities of States	0.2	0.8

Table 2 Conditional Probabilities of Public Signal Given States

	State = 40	State = 80
Public Signal 40	0.8	0.2
Public Signal 80	0.2	0.8

(I) Computational Economy 1 - Strong vs. Semi-Strong Market Efficiency (Approximate Fama, French, Jensen and Roll (1969) -- FFJR)

The first study shows that with Public Information Semi-Strong Form Efficiency arises without Strong-From Efficiency. There are 1000 points of history that are potential trading dates. At any of these points of history an outstanding bid and/or ask as well as the previous history exists. The trader may have private or public information at any point. (The same piece of) Private Information arrives on different traders at the 100th, 200th and 300th point of history – 2 agents receive private information at 100; 3 more at 200; 3 more at 300 and the private information is made public to everyone at the 400th point. A person who receives private information is simply at the right place at the right time.

Each agent can determine a certainty equivalent relative to his information. Because risk preferences are not identical different agents have different certainty equivalents. When informed agents receive a signal they entertain a set of beliefs regarding the likelihood of a particular private signal given a state and update their probabilities of the state accordingly. Such an assumption is meant to capture the idea that a market participant will have access to informed sources concerning how to think about the likelihoods of payoffs and can arrive at a relatively precise belief about the value of a holding a share. For the Economies studied here the appropriate conditional probabilities are represented in Table 3.

Table 3 Conditional Probabilities of Private Signal Given States

	State = 40	State = 80
Private Signal = 40	0.8	0.2
Private Signal = 80	0.2	0.8

The information structure, the timeline and the belief updating are summarized in the Scenario 1 in Section 1.1.2 of the Basic Model in the Appendix.

Having a certainty equivalent each trader must make a decision concerning what to do in the market. He would like to sell above his certainty equivalent or buy below his certainty equivalent. At a particular point in time a trader uses the outstanding bid and ask as well as previous bids, asks and market prices to assess the likelihood each of his possible actions (bid or ask) will be successful in the market in the sense that a bid or ask will be taken. He can also take an existing bid or ask. For example if a trader were at point in time 200 he would see all the information portrayed in figures 4 and 5 prior to point 200. A trader also knows which bids and asks were successful in leading to an exchange so he can see for a particular bid just how frequently it was chosen. Note we assume the trader knows nothing about other traders in the market such as the distribution of traders or whether at any point in time someone other than himself has received private information. Using only the limited trader information about whether particular bids or asks were taken gives the trader some insight (although the lack of specificity of the problem makes the person non-Bayesian) and an ability to calculate relative frequencies of what will happen[5]. This is "as if" behavior in the sense that the trader may not be fully aware that his brain is actually making these computations.

Each trader is guided by maximum profit in the trading market relative to the certainty equivalent. In periods 0 to 100 all traders, though not consciously aware, have the same identical information for determining their value. In this case priors were .2 and .8 on the states 40 and 80 respectively, which are final payoffs. Note that prices need not be constant because trades at a new price for one trader *can change the expected gains to trade* for an action (bid, ask or take) of the other traders. Also a single trader cannot buy or sell indefinitely because of endowment restrictions. The higher the potential profit the more likely a person will be the next in the market. Intuitively this is because those with the highest profit have the most to lose if they do not move and another trader moves first.

[5] The exact description of traders in this economy is given in the appendix and while we used a specific form of the utility function here the theory is not at all constrained to this function or to expected utility maximization for determining the certainty equivalent of a share.

Computational Economy 1 - Bid/Ask Results

We first focus on the bid and ask process to give a clear indication of how the double auction process works in conjunction with the choice process of participants in the auction. In Figure 4 the gray points are asks and the black points bids. There is a tendency toward some scalloping (i.e. a downward string of asks or upward string of bids) in the data, particularly in the early stages of the data. This scalloping most often represents a single agent who tries to make an advantageous trade and then is disciplined by the market in the sense that the market will not "jump at" the over (under) priced ask (bid). A participant represented by a scallop has a much higher expected profit from his action than other traders have from their actions so that participant is most likely to jump in first and try to take advantage of the profit opportunities. The improvement rule forces a potential seller (or buyer) to reduce asks (or increase bids) until a trade is made. A particular seller keeps lowering his ask until his profits are much more in line with those of other traders in the market. As more traders enter then we see competition in bidding behavior. After a trade the outstanding bid and ask are reset and we can see scalloping again. The bids and asks become more and more concentrated because the agents most frequently tend to use asks and bids that have had some degree of success in generating exchanges in the past. In this setting sellers appear to undercut each other but in a formal sense they are not strategic: their actions do not anticipate other traders' responses. The sellers simply make their optimal ask based on the information that is available in the market, i.e. bids, asks, and takes. This competition is not arising from considering what other agents might be thinking but rather it is a form of arbitrage relative to the market bids and asks. After there is new information (i.e. either Private or Public) there is a gradual reduction of the variance of bids and asks up to the point of the next information announcement.

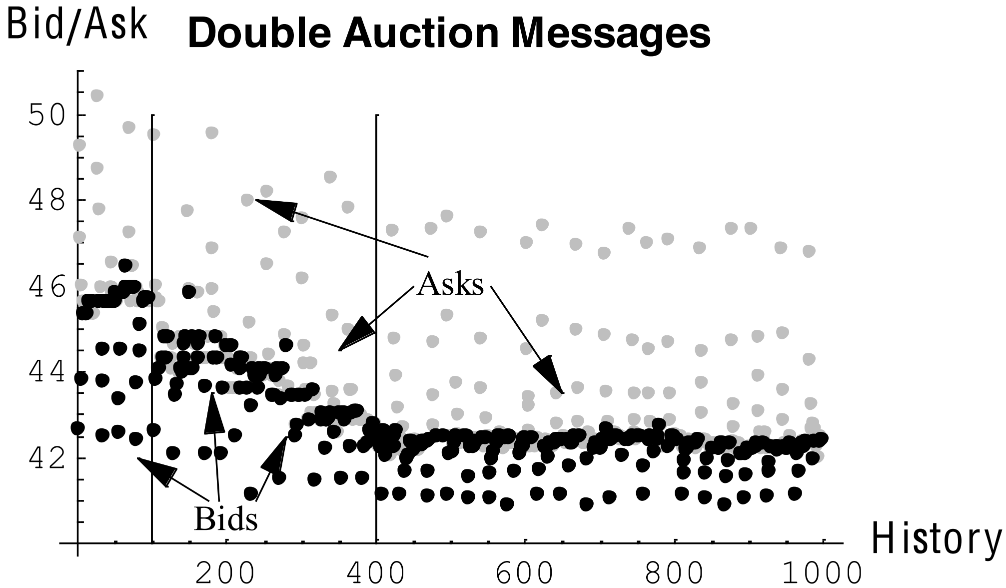

Figure 4 Double Auction Messages from Economy 1

We argue that in this scenario the market is via its operations always discovering information in a Hayekian (1948) sense, i.e. an order is emerging in the market in a way that none of the participants anticipate (i.e. the distribution of bids and asks is an unintended consequence of the actors in the economy).

Computational Economy 1 - Pricing Results

The time series for price in this economy is displayed in Figure 5. Trades are represented by points in the graph. Trades in the interval 0 to 100 reflect gains to risk sharing on the part of traders. At point in time 100 the private signal begins to be absorbed in market price and we see price gradually shifting downward as more and more traders learn the private message. At point in time 400, public information is released and prices settle down to a steady state but not a constant price because any trade can affect the profit opportunities of others. Up until time 400 it is possible to see the private information leaking into prices. During this time there are always two sources of uncertainty when trading, the potential existence of private information and uncertainty about whether a particular action will lead to a consummated transaction.

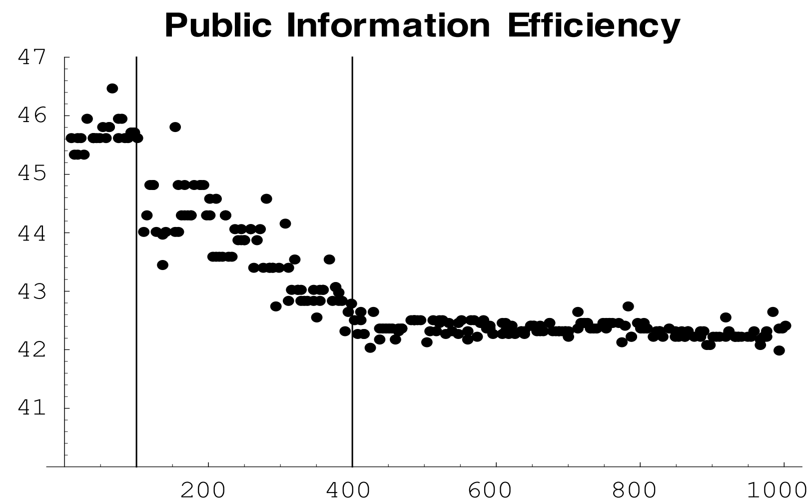

Figure 5 Pricing Messages from Economy 1

We interpret our result as being consistent with semi-strong form efficiency but not strong form efficiency. The graph has a qualitative similarity with the earlier graph for Fama, Fisher, Jenson and Roll (1969). To get a better feel for this result we bootstrapped this procedure 100 times yielding a distribution of regression coefficients of prices on time. For comparison, we put the diagrams together with the Ball and Brown (1968) computational economy described in the next section. Note the distribution is much more symmetrically located around zero after public information release than before public information release.

(II) Computational Economy 2 (Approximate Ball and Brown [1968] -- BB)

To approximate the Ball and Brown results we keep the scenario the same as the FFJR case but assume that the private information is not equivalent but only stochastically related to the public information. Roughly speaking the public signal is correlated but not perfectly correlated with the private signal. There are again 1000 points of history that are potential trading dates. At any of these points of history an outstanding bid, ask or exchange may occur. Private Information arrives at the 100th point of history; two agents receive private information at 100; three more at 200; and three more at 300. At point in time 400 Public Information (40 here) is made available to all the traders in the market and traders adjust their values accordingly. Note each trader still makes bids and asks relative to the trades and bids and asks that have occurred previously but now the less informed agents use the same public signal to arrive at their valuation while the fully informed use both their private signal and the public signal. Agents in the setting in which public information is only partially correlated with private information have the same distribution of risk preferences when the private signal is fully revealed in the private signal. The bid/ask data for the correlated setting has the same general properties as the complete revelation setting although it appears that offers are disregarded a bit more as time increases after the public information announcement.

The information structure, the timeline and the belief updating are summarized in the Scenario 2 in Section 1.1.2 of the Basic Model in the Appendix. For this Economy (and Economy 3 later) the appropriate additional conditional probabilities are represented in Table 4.

Table 4 Conditional Probabilities of Public Signal Given State and Private Signal

(1) State = 40

	Private Signal = 40	Private Signal = 80
Public Signal = 40	0.9	0.7
Public Signal = 80	0.1	0.3

(2) State = 80

	Private Signal = 40	Private Signal = 80
Public Signal = 40	0.3	0.1
Public Signal = 80	0.7	0.9

The double auction price data follows:

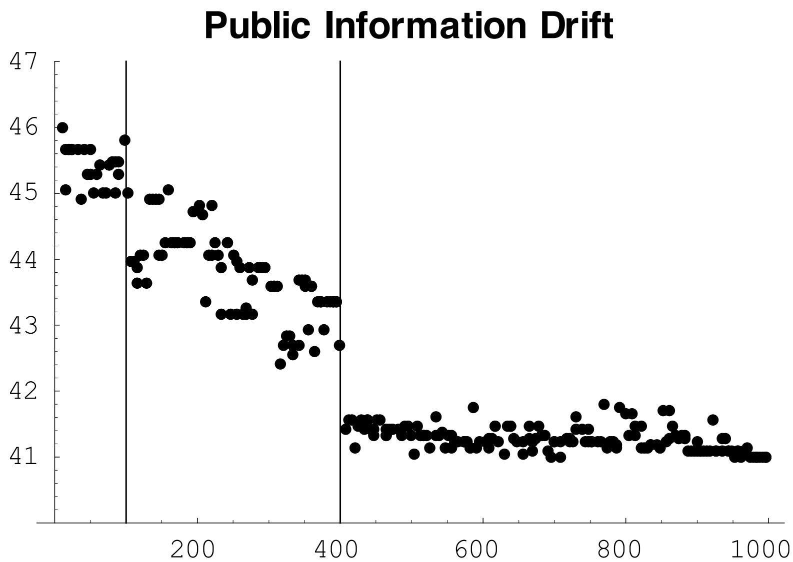

Figure 6 Price Messages from Economy 2

There is an apparent drift in these data. As in computational economy 1, we also bootstrapped this procedure 100 times yielding a distribution of regression coefficients of prices on time. The distributions of the regression coefficients of prices against time for the case when the private information is fully revealed in the public signal and the case when the private information is only correlated with the public signal are displayed in Figure 7.

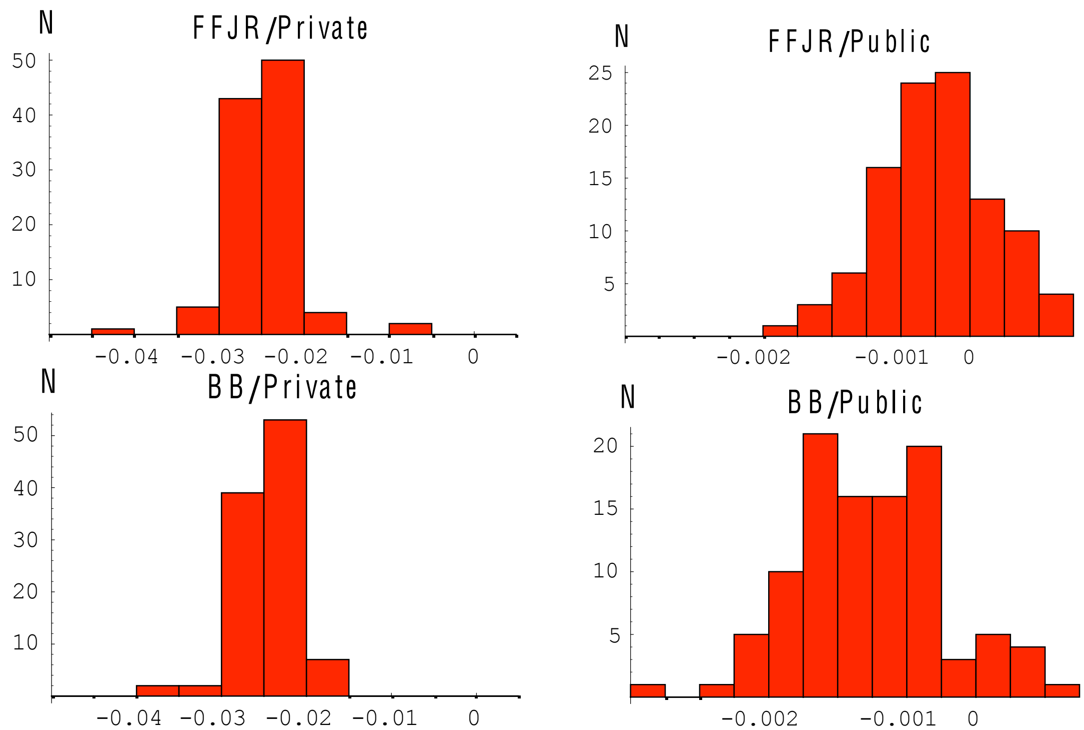

Figure 7 Distributions of the Regression Coefficients of Prices against Time

The first row is the coefficients corresponding to the Private Information times (101 to 400) and Public Information times (401) to 1000 for FFJR (Economy 1). The distribution of coefficients has shifted significantly to the right in going from private to public information. The second row is for the BB Private and Public Information (Economy 2). Since both private information graphs are drawn on the same scale and both public information graphs are drawn on the same scale it is clear that prior to public information the distribution of coefficients is virtually identical. On the other hand after release of Public information the distribution of coefficients under BB is shifted to the left and is significantly different from that under FFJR at the .001 level using a Kolmogorov Smirnov test.

We also compared an estimate of the autocorrelation process using GARCH techniques. By expanding the number of agents in the economy to 50, we were able to show that when asymmetric information was present in a setting that the prices followed an GARCH (2,1) process and when all information was public an GARCH process could not be detected. This is consistent with what have been documented in experimental data (eg. Bruguier et al. 2008).

(III) Computational Economy 3 - Forecasts (Approximate Foster, Olsen and Shevlin [1984])

We now examine another information structure that generates drift and is characteristic of drift as often seen in accounting data. This approach can be developed within the discussion of Public vs. Private Information and assuming a relationship between Private and Public Information described in Economy 2. Here we build on Foster et al. (1984) which directly examines responses to forecasts and later earnings announcements. Several important features of the Foster et al. study are important. First if forecasts reveal bad information, prices are driven downward and then if a later announcement reveals bad information prices are driven down further. We also wish to capture the notion that after the public information release prices may not have fully adjusted to all the available information in the economy and there is the potential for some ex-post announcement drift. In our story the drift arises because there is also private information that arrives prior to the release of public information and that private information is not fully redundant with public information. Thus we are employing much of the structure outlined in Economy 2 and the timeline describing that scenario. We add one more piece of information in which all asymmetry between parties is resolved.

We approach this task in the following way. First we examine prices based on all agents having the same information and heterogeneous risk preferences. Then a public forecast is released and all agents realign their portfolios based on risk sharing issues. Following this event a private information release occurs. Later on the public information is released. The private information is only partially correlated but not fully redundant with the public information so it is possible market price might not fully respond to all the outstanding public information. At some later point in time there is yet another public release such as a target price (or possibly another earnings announcement) which implicitly reveals the previously unrevealed portion of the private information. The information structure, the timeline and the belief updating are summarized in the Scenario 3 in Section 1.1.2 of the Basic Model in the Appendix. For this particular computational study the appropriate conditional probabilities are represented in Table 5.

Table 5 Conditional Probabilities of Forecast Given Public Signals

	Public Signal = 40	Public Signal = 80
Forecast = 40	0.8	0.2
Forecast = 80	0.2	0.8

Computational Economy 3 - Pricing Results

Now we examine what happens to observable prices that arise from employing our price formation model (Figure 8). When agents receive the same forecast which is a piece of public information, prices quickly jump to a new price and remain stable (period [76,150]). Then private information release occurs and agents trade on their own account and prices move downward but in a manner in which prices drifted downward in Figure 5. Thus we have evidence again of lack of so-called strong form efficiency. At point in time 225 public information release occurs but now, unlike the release in described in Figure 5 the private and public information are not fully redundant. In such a setting prices drift after the arrival of public information. Prices only adjust further once the information that is known privately but not publicly is revealed.

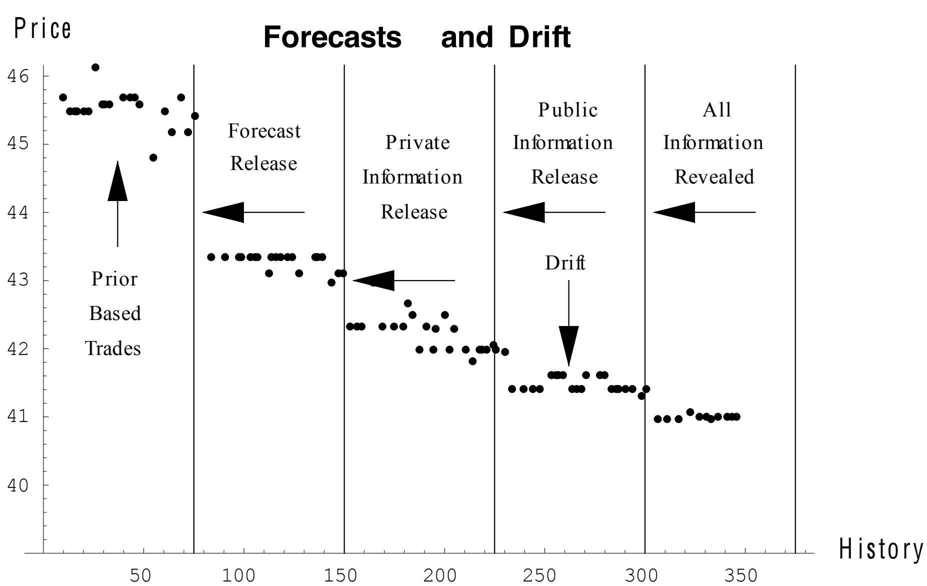

Figure 8 Price Messages from Economy 3

IV. Robustness

Having discussed basic results regarding the existence of semi-strong efficiency, there are a number of ways to examine the robustness of our price formation dynamic. One is to simply replicate the same economy for a large number of trials. This approach was taken for describing the initial studies of semi-strong efficiency and failure of semi-strong efficiency and the behavior of the economies were within a well described tolerance level (See Figure 7 for example). In addition the parametric specifications of the utility functions, the endowments of agents and initial probabilities of traders have been randomly varied and consistently replicate. Here we report the tests of robustness related to the ability to find different systematic results in the sequence of prices including results relating to volatility and results capturing non-rational models of value.

Volatililty – Economy 4

There are numerous results in the literature regarding volatility. We focus on a volatility result for which there is theoretical basis. Veronesi (1999) proposes that volatility in response to information regarding a particular event is conditional on the priors of that event. In particular given the same information event the market with the higher prior probability will have more volatility in the market price when the same information is revealed. In the Veronesi model prices move from one equilibrium level to another. There is no mechanism by which individual agents formally enter into the price setting process through making bids or asks. In our model it is possible to produce Veronesi's volatility result when traders are actively engaged in trading in an auction; we can then show how differences in measured volatility arise as a consequence of differences in prior probabilities in that auction.

Figure 9 portrays the outcome of applying our price formation mechanism under Veronesi’s assumptions. In one economy agents begin with high priors (.8) of the high state outcome (80) while in the other economy agents begin with low priors (.2) of the high state occurring. The remaining aspects of these two economies are identical and reflected in Tables 1 and 2. We assume that in each setting a public low signal (40) is received at trading point 50. The interior graph plots price on the vertical axis and history of trades on the horizontal when all agents had a prior of 0.2 and received a public signal of 40. When everyone receives an identical piece of information the adjustment in price is very quick even though individual agents are trading only on their private account and in no way attempt to assess a distribution of future prices. Note that there is some variability after the adjustment occurs. In the interior graph, the ticks on the vertical axis increase at a rate of 0.2 which suggests the volatility is low. If we measure volatility as the standard deviation of price after the release of the information it is 0.036. On the other hand the description of a market reaction when there are high priors is portrayed in the outer graph. In it agents trade on their own account and prices adjust quickly. Once again there is no anticipation by agents of what other agents will do and in this sense they do not behave strategically. The range of the outer graph is 40 to 50 which seems to suggest that the volatility for prices after receiving information is greater than when the priors are 0.2. Volatility as measured by standard deviation of prices after the information release is now 0.155. So the Veronesi claim is reflected in this data.

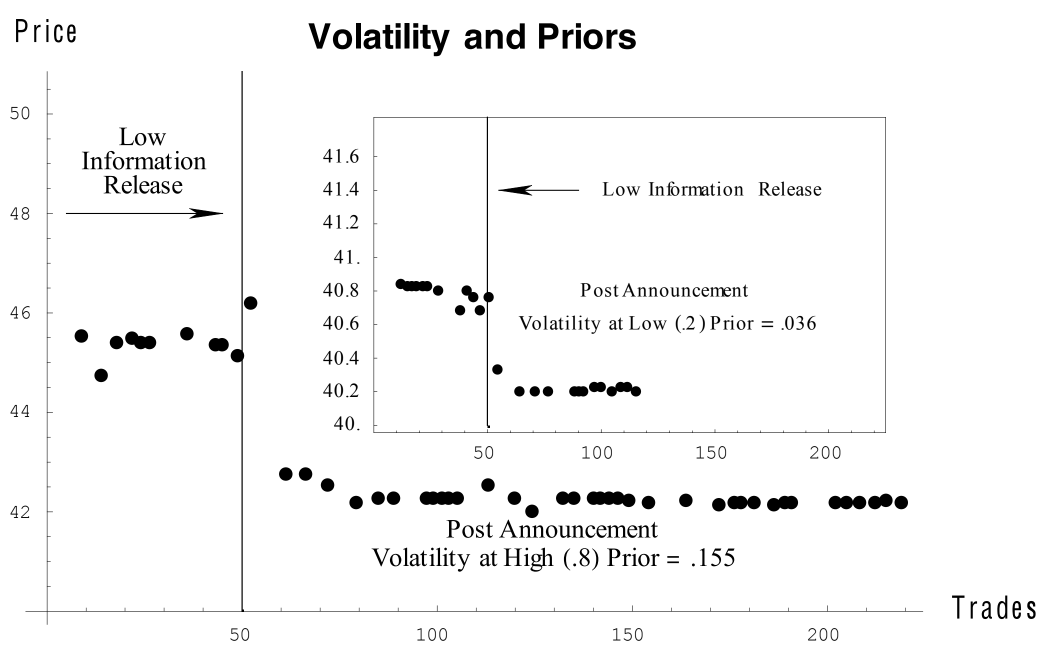

Figure 9 Price Messages Showing Volatility

Non rationality in revising probabilities in valuation – Economy 5

We next turn to what happens if market agents incorporate in a non-Bayesian manner in their valuations. One popular explanation of market anomalies, including in particular post-announcement drift centers on the theory that informed agents are overconfident about the quality of their private information. Such overconfidence leads to a distortion in market prices because informed agents overreact to new information and such overconfidence is further augmented by what is referred to as self-attribution. One such model belongs to Daniel et al. (1998). In this model, informed traders receive superior information and this overconfidence leads such informed traders to become overconfident in their valuation which produces large jumps in prices. Seeing such a change in prices leads the overconfident traders to become even more confident and self-attribute this overconfidence to their own knowledge which creates a further change in prices. Post announcement drift arises when prices correct for the overconfidence of those with superior information and is finally resolved when the information reaches all the market participants.

Like the Veronesi model, how prices changes is an equilibrium model and there is no attempt to specify how such equilibria are reached. Using our price formation model and using valuation procedures described by Daniel et al., we are able to trace out price paths that mimic the Daniel et al. story and hence show how it is possible for individual agents to engage in exchange in such a way as to achieve the Daniel et al. conclusions. In this economy, there are 2 types of investors, informed and uninformed. The informed in this structure are the “irrational”. When they receive private information they think it is better than it really is. In particular in our setting although the true nature of the private signal has the conditional probabilities described in Table 6, the informed individuals think the information is characterized by Table 7.

Table 6 Conditional Probabilities of Private Signal Given States

	State = 40	State = 80
Private Signal = 40	0.7	0.3
Private Signal = 80	0.3	0.7

Table 7 Conditional Probabilities of Private Signal Given States for Informed Overconfident Agents

	State = 40	State = 80
Private Signal = 40	0.8	0.2
Private Signal = 80	0.2	0.8

In Daniel et al., the informed agents are all risk neutral, the uninformed are risk averse and the informed agents behavior determines what the price will be when new information is received. We are able to show how the story works in this setting with a differential risk aversion across agents.

In the Daniel et al story the informed agents' overconfidence produces changes in prices and the agents self-attribute the price change to themselves. Such self-attribution makes the informed even more extreme in their belief about the quality of their own information. This is captured in Table 8.

Table 8 Conditional Probabilities of Private Signal Given States for Informed Overconfident Self-attributing agents

	State = 40	State = 80
Private Signal = 40	0.9	0.1
Private Signal = 80	0.1	0.9

Informed agents continue to set price and such self-attribution drives the prices even further downward, then public information which is fully redundant with the private information arrives. At this point informed agents knowing the new information is public begin to lose some confidence but not so much as to arrive at the conditional probabilities described in Table 6. Rather they make a partial adjustment which in our case will be represented by Table 9. The resulting state in which beliefs are not fully adjusted to those of Table 6 means that prices only partially adjust to the release of Public Information. At a later point in time the informed investors fully discount their overconfidence.

Table 9 Partially Adjusted Conditional Probabilities of Private Signal Given States for Informed Agents

	State = 40	State = 80
Private Signal = 40	0.8	0.2
Private Signal = 80	0.2	0.8

Prices (Figure 10) under the price formation mechanism follow the Daniel et al. story. Private information hits informed investors at 50th point in time and there is a rather severe drop in prices in the first set of overconfidence driven trades. The self-attribution phenomenon produces a further drop until public information arrives at 150. At that point in time prices begin to drift upward but do not fully respond to the public information (thus we have an example of drift). When the informed finally realize that their overconfidence is unwarranted prices eventually reach a stable state in the periods 230 to 250.

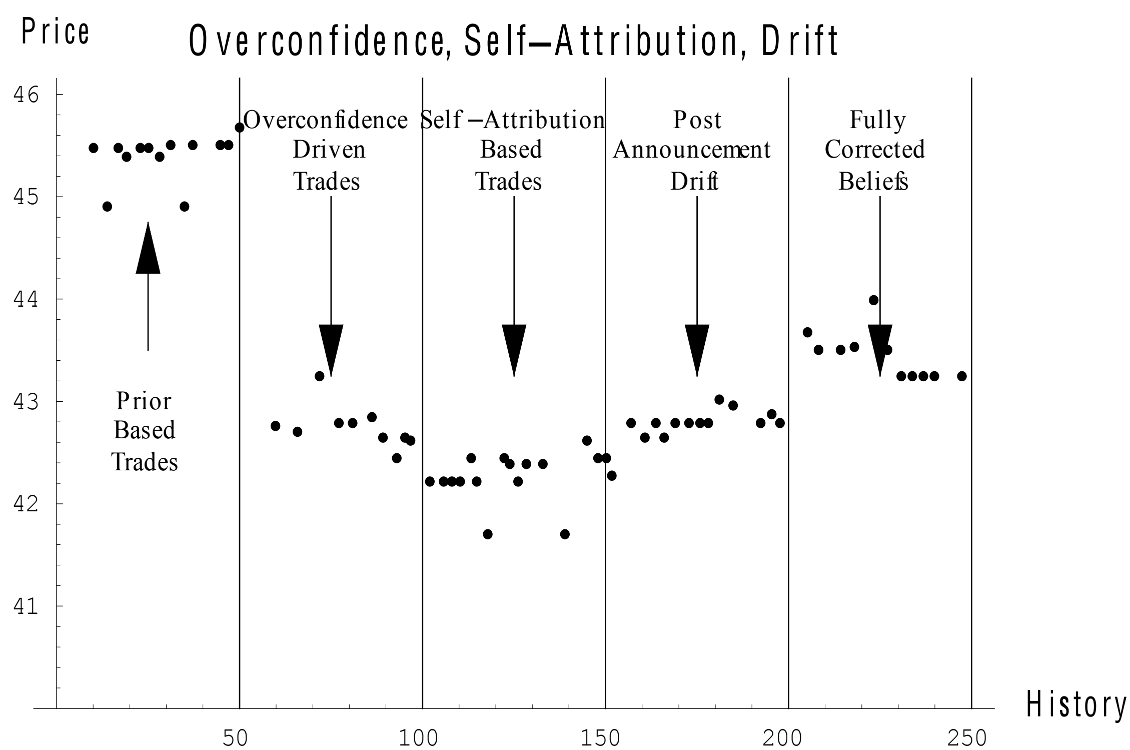

Figure 10 Price Messages from Non Bayesian Investors

V. Conclusions and Caveats

This paper builds a model of price formation and shows that the model has the ability to generate predictions that lead to a variety of phenomena including semi-strong market efficiency, semi-strong market inefficiency, volatility and drift generated by overconfidence. Rather than specify just prices, the model specifies the generation of all messages as they reach the market. As a step in understanding price formation with risky assets it provides something very different from the results of experiments. It in fact yields the ability to test parametric predictions of the theory in a variety of experimental settings where tastes can be estimated or induced. O'Brien (xxx)'s as well as Dickhaut et al. (1993)'s results of inducing preferences auctions suggest the plausibility of tracking the predictions of such a model. O'Brien does not explicitly attempt to specify the evolution of all of the messages in the market. Recently Gjerstad (2006) in a non risky setting has been able to explicitly test the predictions of this price formation process in a case where individuals' values were certain and there was no asymmetric state information.

But we have developed more than a model to predict prices in the laboratory. We explored fundamental findings from accounting and finance. Notions of market weak and strong form efficiency have been replicated extensively in archival settings. These results are presented in a 1969 review by Fama. Our contribution is to provide a theory derived from laboratory experimentation that can explain such results and provide conditions under which they might occur. Unlike standard accounting and finance theory we assume that the choices of what to do in the market involve how to submit bids, asks and take outstanding bids and asks by individual traders. Thus we attempt to build from behavior by an individual in a specific setting, the double auction (for related discussion, see consolidated limit order books below.)

We outline in the appendix how the computational model presented here can provide a basis for structuring an extended examination of multiple securities in which each security is sold separately as in a typical exchange like the NYSE. Such an approach assumes it is possible to assess (under any theory of individual choice such as EU, prospect theory or rank dependent utility) the marginal value of adding any stock to the existing portfolio. Beyond this level of additional complexity it is also possible to use this framework to begin to address additional findings discussed in the second Fama (1990) review article. The latter findings include size effects, book to market effects. The framework outlined here allows for such a representation.

In particular if the number of firms issuing shares in the economy is expanded it is reasonable to allow additional factors such as size to enter individually into the valuations described in step 1 of our model. In terms of inside information it is possible to allow specific subsets of traders to always get the information. As new findings emerge then it is possible to include them in the model and arrive at potential implications.

Of course to the extent that phenomena are unobservable to the archivalist such a model would not be tested directly from classical data such as that which occurs on the CRSP and Computstat tapes. However the exact specification of conditions under which the phenomena would occur would create a large number of settings that would then be subject to laboratory tests. Thus we envision a natural movement from the archival data to computational theory to laboratory test. And of course there would be a movement back again to more specific archival tests.

To date we have worked only with the experimental paradigm of the double auction which has been championed as a mechanism that has many similarities to large scale exchanges. Frequently however large exchanges involve a system of dealers. We note that in reality many of the differences between a dealer system and typical large scale double auction exchange have been eroded by the existence of CLOBs and Hybrid CLOBs such as the NYSE and NASDAQ in which the collective limit orders are known (Jain, 2003). Thus an interesting and useful extension of this study will be an attempt to assess how the price formation differences between CLOBs and the double auction diminish as the existence of the open book increases.

We believe an important aspect of the model is that as long as there are heterogeneous preferences the model does not reach an equilibrium per se in this classical sense. The reason is a purchase by one trader can affect the adjusted beliefs of another player so that they will be willing to offer a new bid which might be taken thus altering the existing belief distribution which can lead to a new trade etc. So prices can easily vary around a particular asymptote. This is a feature of the model which replaces the classical idea of noise traders in theoretical models with a rational expectations flavor. A look at the close of traditional double auction experiments suggests that there can still be some noise after many periods of trade for the same asset. Also the asymptotes themselves can vary as well as the way these asymptotes are reached which permits the comparison of different models of individual choice (including expected utility theory, rank dependent utility theory) in the determination of prices.

Given our general emphasis on information, we believe this work is an important step in attempting to integrate theoretical, experimental and archival work in accounting. In the appendix we outline how many additional variables can be implemented in this setting such as a large number of traders, short-selling and multiple securities. We hope that this construction can serve as a template for these more extended considerations.

References

Arrow, K.J. 1994. Methodological individualism and social knowledge. *American Economic Review* 84 (May):1-9.

Ball, R., and P. Brown. 1968. An empirical evaluation of accounting income numbers. *Journal of Accounting Research* 6 (Autumn): 159-178.

Bernard, V.L., and J.K. Thomas. 1989. Post-earnings-announcement drift: Delayed price response or risk premium. *Journal of Accounting Research* 27 (supplement): 1–35.

Bernard, V.L., and J.K. Thomas. 1990. Evidence that stock prices do not fully reflect the implications of current earnings for future earnings. *Journal of Accounting and Economics* 13 (December): 305–345.

Cason,T.N., and D. Friedman. 1993. An empirical analysis of price formation in double auction markets, in *The Double Auction Market: Institutions, Theories, and Evidence.* D. Friedman and J.Rust, eds., 253-283. Reading, MA. Addison-Wesley.

Cason, T.N., and D. Friedman. 1996. Price formation in double auction markets. *Journal of Economic Dynamic Control* 20 (August): 1307-1337.

Daniel, K, D. Hirshleifer, and A. Subrahmanyam. 1998. Investor psychology and security market under- and overreactions. *Journal of Finance* 53 (December): 1839-1885.

De Bondt, W., and R. Thaler. 1985. Does the stock market overreact? *Journal of Finance* 40 (July): 793-808.

Dickhaut, J., O. Lungu, V. Smith, B. Xin and A. Rustichini. 2008. A neuronal theory of human economic choice. Working paper, University of Minnesota.

Dickhaut, J., A. Martin, D. Senkow and G. Sevcik. 1993. The price/quantity sealed bid/offer auction with pro-rata rationing: experimental evidence. *The Economic Journal* 103 (May): 547-569.

Duffy, J. 2004. Agent based models and human subject experiments. Working paper.

Easley, D., and J. Ledyard. 1993.Theories of price formation and exchange in double oral auctions, in *The Double Auction Market: Institutions, Theories, and Evidence*. D. Friedman and J.Rust, eds. 63-97. Reading, MA. Addison-Wesley.

Fama, E.F. 1965. The behavior of stock market prices. *The Journal of Business* 38 (January): 34-105.

Fama, E. 1998. Market efficiency, long-term returns, and behavioral finance. *Journal of Financial Economics* 49 (September): 283-306.

Fama, E., L. Fisher, M. Jensen, and R. Roll. 1969. The adjustment of stock prices to new information. *International Economic Review* 10 (February): 1-21.

Fiorillo, CD., N. Philippe, P. Tobler, and W. Schultz, 2003. Discrete coding of reward probability and uncertainty by dopamine neurons. *Science* 299: 1898-1902.

Foster, G., C. Olsen, and T. Shevlin. 1984. Earnings releases, anomalies, and the behavior of security returns. *The Accounting Review* 59 (October): 574-603

Friedman, D. 1991. A simple testable model double auction markets. *Journal of Economic Behavior and Organization* 15 (January): 47-70.

Gigerenzer, G., and H. Ulrich. 1995. How to improve Bayesian reasoning without instruction: frequency formats. *Psychological Review* 102: 684-704.

Gjerstad, S. 2006. The competitive market paradox. *Journal of Economic Dynamics and Control* 31 (May): 1753-1780.

Gjerstad, S. and J. Dickhaut. 1998. Price formation in double auctions. *Games and Economic Behavior* 22 (January): 1-29.

Gode, D., and S. Sunder. 1993. Allocative efficiency of markets with zero intelligence traders: Market as a partial substitute for individual rationality. *Journal of Political Economy* 101(February): 119-137.

Grossman, S. 1976. On the efficiency of competitive stock markets where trades have diverse information. *Journal of Finance* 31 (May): 573-585.

Grossman, S. 1978. Further results on the informational efficiency of competitive stock markets. *Journal of Economic Theory* 18: 81-101.

Grossman, S., and J. Stiglitz. 1980. On the impossibility of informationally efficient markets. *American Economic Review* 70 (June): 393-408.

Hahn, F. 1988. Auctioneer, in *The New Palgrave: A Dictionary of Economics*, 136-138. The Macmillan Press Limited.

Hayek , F.A. 1948. *Use of knowledge in society.* reprinted in *The Essence of Hayek.* C. Nishiyam and K.R. Leube, eds. Hoover Institution Press. Stanford, CA.

Jain, P. 2003. Institutional design and liquidity at stock exchanges around the world. Working paper, Indiana University.

Jegadeesh, N., and S. Titman. 1993. Returns to buying winners and selling losers: implications for stock market efficiency. *Journal of Finance* 48 (March): 65-91.

Kohn, M. 2004. Value and exchange. *The Cato Journal* 24 (Fall): 303-339.

Ketcham, J., L. Smith, and A.W. Williams. 1984. A comparison of posted-offer and double auction pricing institutions. *Review of Economic Studies* 51 (October): 595-614.

Lakonishok, J., A. Shleifer and R. Vishny. 1994. Contrarian investment, extrapolation, and risk. *Journal of Finance* 49 (December): 1541-1578.

Marcet, A., and T. Sargent. 1989. Convergence of least squares learning mechanisms in self-referential liner stochastic models. *Journal of Economic Theory* 48: 337-368.

Marcet, A., and T. Sargent. 1989. Convergence of least-squares learning in environments with hidden state variables and private information. *Journal of Political Economy* 97 (December): 1306-1322.

Michaely, R., R. Thaler, and K. Womack. 1995. Price reactions to dividend initiations and omissions. *Journal of Finance* 50 (June): 573-608.

O'Brien, J., and S. Srivastava. 1991. Dynamic stock markets with multiple assetd: an experimental analysis. *Journal of Finance* 46 (December): 1811-1838.

Plott, C. 1982. Industrial organization theory and experimental economics. *Journal of Economic Literature* 20 (December): 1485-1527.

Shanthikumar, D. 2003. Small and large trades around earnings announcements: Does trading behavior explain post-earnings-announcement drift? Working paper, Stanford University.

Shin, H. S. 2005. Endogenous Disclosures and the Post-Earnings Announcement Drift. Working paper, London School of Economics.

Smith,V.L. 1962. An experimental study of competitive market behavior, *Journal of Political Economy* 70 (April): 111-137.

Smith, V.L. 1982. Microeconomic systems as an experimental science. *American Economic Review* 72 (December): 923-955.

Tremblay, L., and W. Schultz. 1999. Relative reward preference in primate orbitofrontal cortex. *Nature* 398 (April): 704-708.

Veronesi, P. 1999. Stock market overreaction to bad news in good times: A rational expectations equilibrium model. *The Review of Financial Studies* 12 (Winter): 975-1007.

Wilson, R.B. 1987. On equilibria of bid-ask markets in Arrow and the Ascent of Modern Economic Theory. G.W. Feiwel, ed. 375-414. New York, New York University Press.

Wolfram, S. 2002. *A new kind of science*. Champaign, IL. Wolfram Media.

Appendix

1. The Basic Model

1.1 Environment at time t

1.1.1 Preferences and Valuations

The model can be applied to an arbitrarily large finite set of states and signals. For concreteness we focus on a world where there are two states, and each time information is released there are only two possible signals. We assume an economy with n agents, all of whom will consume at a future time, T. There are two types of assets, riskless and risky.[6] At T the risky asset pays off one of two amounts of the riskless asset, $\{x_L, x_H\}$, where $x_H > x_L$. At any time, t, agent i holds a belief $P_{H,i,t}$ $(0 \le P_{H,i,t} \le 1)$ regarding the probability of the high outcome. To clarify the ideas here we assume each agent i has a constant absolute risk averse utility function defined relative to the holdings of the riskless asset at time T[7]:

$$U_i(w_i) = -e^{-\lambda_i w_i}$$

where w_i denotes the i's wealth at T and $0 \le \lambda_i < \infty$ is agent i's risk parameter. Given λ_i, $P_{H,i,t}$, and $\{x_L, x_H\}$, the expected utility of the risky asset to i at t will be

$$EU_i(\lambda_i, P_{H,i,t}, \{x_L, x_H\}) = (1 - P_{H,i,t}) \bullet (-e^{-\lambda_i x_L}) + P_{H,i,t} \bullet (-e^{-\lambda_i x_H})$$

The valuation of the gamble in units of the riskless asset is that value of $c_{i,t}$ which solves

$$-e^{-\lambda_i c_{i,t}} = (1 - P_{H,i,t}) \bullet (-e^{-\lambda_i x_L}) + P_{H,i,t} \bullet (-e^{-\lambda_i x_H})$$

where $c_{i,t}$ represents that amount at which agent i is just willing to buy or sell at time t.

1.1.1.1 Remarks

In the above setting the agents may or may not have homogeneous beliefs. Later, we consider settings where the market might not fully incorporate information. Furthermore t is meant to be any arbitrary t less than T, thus t might be the initial time ($t=0$), or it may represent the time at which private or public information is received or the time of a forecast.

1.1.2 The different information structures

[6] The model can go beyond two types of assets. Please refer to Section 2 of the appendix where we allow for multiple assets.
[7] Any structure that can define certainty equivalents for each probability payoff pair (and endowment) will also work here. So for example prospect theory, etc.

We focus on several prototypical information scenarios. In all scenarios all agents start out with the same prior beliefs. In scenarios 1 and 2 privileged agents receive private information and later public information is released. Scenarios 1 and 2 are captured by the following timeline.

Period Begins	*Private Information Received*	*Receipt of Public Information*
$t = 0$	$t = t_{private}$	$t = t_{public}$

In each setting there are initial prior beliefs for the low and high state represented by the vector

$$S = \begin{pmatrix} P_L \\ P_H \end{pmatrix}$$

At time $t_{private}$ selected agents are informed of a private signal which can take on the values L or H. The likelihoods of the private signal given low and high states are given by the matrix

$$SM1 = \begin{pmatrix} P_{L|L} & P_{H|L} \\ P_{L|H} & P_{H|H} \end{pmatrix}$$

where $P_{i|j}$ is the conditional probability of message i given state j.

The scenarios differ at time, t_{public}. In scenario 1 the public information is identical to the private and so the likelihood of the public signal given the private signal is simply equivalent to the private signal given the state. In Scenario 2 the public signal is a separate random drawing and thus the public information is only partially correlated with the private information. An appropriate probabilistic representation of these facts is

$$SM2M1_L = \begin{pmatrix} P_{L|L,L} & P_{L|H,L} \\ P_{H|L,L} & P_{H|H,L} \end{pmatrix} \text{ and } SM2M1_H = \begin{pmatrix} P_{L|L,H} & P_{L|H,H} \\ P_{H|L,H} & P_{H|H,H} \end{pmatrix}.$$

where $P_{i|j,k}$ represents the probability the public signal is i given the private signal is j and the state is k.

In this construction we assume that agents after seeing information perform Bayesian revisions to arrive at revised probabilities. Thus in Scenario 1 an agent would calculate the probability of the $High$ state given Low private and public messages as

$$\frac{SM1_{2,1} \bullet S_2}{SM1_{2,1} \bullet S_2 + SM1_{1,1} \bullet S1}$$

On the other hand in Scenario 2 the agent would calculate the probability of the *High* state given a *Low* private and public message based on the dependency between the messages as

$$\frac{(SM2M1_H)_{1,1} \bullet SM1_{2,1} \bullet S_2}{(SM2M1_H)_{1,1} \bullet SM1_{2,1} \bullet S_2 + (SM2M1_L)_{1,1} \bullet SM1_{1,1} \bullet S_1}$$

This structure is reflected in Table 10.

Table 10 Information Structures without forecast

Scenario 1	Initial Beliefs $S = \begin{pmatrix} P_L \\ P_H \end{pmatrix}$	Private Information $SM1 = \begin{pmatrix} P_{L\|L} & P_{H\|L} \\ P_{L\|H} & P_{H\|H} \end{pmatrix}$	Public Information Fully Redundant with Private Information $SM2M1 = SM1 = \begin{pmatrix} P_{L\|L} & P_{H\|L} \\ P_{L\|H} & P_{H\|H} \end{pmatrix}$
Scenario 2	Initial Beliefs $S = \begin{pmatrix} P_L \\ P_H \end{pmatrix}$	Private Information $SM1 = \begin{pmatrix} P_{L\|L} & P_{H\|L} \\ P_{L\|H} & P_{H\|H} \end{pmatrix}$	Public Information Partially Correlated with Private Information $SM2M1_L = \begin{pmatrix} P_{L\|L,L} & P_{L\|H,L} \\ P_{H\|L,L} & P_{H\|H,L} \end{pmatrix}$ $SM2M1_H = \begin{pmatrix} P_{L\|L,H} & P_{L\|H,H} \\ P_{H\|L,H} & P_{H\|H,H} \end{pmatrix}$

Scenario 3 adds an additional event, the receipt of a forecast of the public signal, and is captured by the following timeline.

Period Begins	*Public Forecast*	*Private Information Received*	*Receipt of Public Information*
$t = 0$	$t = t_{forecast}$	$t = t_{private}$	$t = t_{public}$

Table 11 captures the likelihood structure associated with these events for this scenario.

Table 11 Likelihood Structures with forecast

Scenario 3	Initial Beliefs $S = \begin{pmatrix} P_L \\ P_H \end{pmatrix}$	Forecast Information $FM = \begin{pmatrix} P^f{}_{L\|L} & P^f{}_{L\|H} \\ P^f{}_{H\|L} & P^f{}_{H\|H} \end{pmatrix}$	Private Information $SM1 = \begin{pmatrix} P_{L\|L} & P_{H\|L} \\ P_{L\|H} & P_{H\|H} \end{pmatrix}$	Public Information Partially Correlated with Private Information $SM2M1_L = \begin{pmatrix} P_{L\|L,L} & P_{L\|H,L} \\ P_{H\|L,L} & P_{H\|H,L} \end{pmatrix}$ $SM2M1_H = \begin{pmatrix} P_{L\|L,H} & P_{L\|H,H} \\ P_{H\|L,H} & P_{H\|H,H} \end{pmatrix}$

where $P^f_{i|j}$ represents the probability of the forecast i given the public signal is j.

In Scenario 3 there are 3 points in time ($t_{forecast}$, $t_{private}$, t_{public}) at which Bayesian updating occurs. Table 12 summarizes the probability of the State being i given various combinations of signals.

Table 12 Belief Updating When There is a Forecast

$t_{forecast}$	$\dfrac{\sum_{k'=1}^{2}\sum_{r'=1}^{2} FM_{j,k'} \bullet (SM2M1_i)_{k',r'} \bullet SM1_{i,r'} \bullet S_i}{\sum_{i'=1}^{2}\sum_{k'=1}^{2}\sum_{r'=1}^{2} FM_{j,k'} \bullet (SM2M1_{i'})_{k',r'} \bullet SM1_{i',r'} \bullet S_{i'}}$
$t_{private}$	$\dfrac{\sum_{k'=1}^{2} FM_{j,k'} \bullet (SM2M1_i)_{k',r} \bullet SM1_{i,r} \bullet S_i}{\sum_{i'=1}^{2}\sum_{k'=1}^{2} FM_{j,k'} \bullet (SM2M1_{i'})_{k',r} \bullet SM1_{i',r} \bullet S_{i'}}$
t_{public}	$\dfrac{(SM2M1_i)_{k.r} \bullet SM1_{i,r} \bullet S_i}{\sum_{i'=1}^{2} (SM2M1_{i'})_{k.r} \bullet SM1_{i',r} \bullet S_{i'}}$

1.2 Institution – The double auction (DA)

In the double auction, potential sellers post asks and potential buyers post bids. The message space defines a set of allowable messages for each agent. In this paper we consider the double auction with a bid-ask spread reduction rule (defined below). In effect, this produces restrictions on any agent's messages as a function of previous messages from all the agents. The DA imposes no restrictions on the sequencing of messages: any agent can send a message at any time during the trading period. Allocation of units is by mutual consent between any buyer and seller. If a seller's ask is acceptable to a buyer, then a transaction is completed when the buyer takes (accepts) the seller's ask. Similarly a buyer's bid may be accepted by a seller.

Definition 1: Asks. An ask a is an amount that a potential seller i is willing to accept from any buyer as payment for a unit of the commodity being traded. To submit an ask of a, seller i sends the message (i, 0, a).

Definition 2: Bids. A bid b by a potential buyer j is an amount that j is willing to pay to any seller for a unit. Buyer j submits this bid by sending the message (0, j, b).

Definition 3: Spread Reduction Rule. The lowest ask in the market at any time is called the outstanding ask and is denoted oa. At any time agents must place asks a satisfying $a < oa$. The highest bid is called the outstanding bid ob. If agents place a bid it must be above the outstanding bid. The outstanding ask oa and outstanding bid ob define the bid-ask spread [ob,oa]. In markets with a spread reduction rule, all bids and asks fall within the bid-ask spread.

Definition 4: Acceptance. If a potential seller i sends the message $(i, 0, a)$ and holds the outstanding ask $oa = a$, then a take of oa by a buyer j is an agreement by j to purchase a unit from seller i at the transaction price $p = oa$. Buyer j accepts the outstanding ask oa by sending the message $(0, j, b)$ where $b = oa$. Similarly, if the outstanding bid ob is held by buyer j', then a take of ob by seller i' is an agreement by i' to sell a unit to buyer j at the transaction price $p = ob$.

Definition 5: Trades. A trade is represented by (i, j, oa) if j takes i's ask or (i', j', ob) if i' takes the bid of j'.

1.2.1 Observed Histories

This section demonstrates a method of describing data to arrive at what needs to get into the calculation of probabilities. A record of bids, asks and takes constitutes a history of transactions. For example consider the following history

$H_2 = \{h_1, h_2\} = \{(3, 0, 45.00), (3, 1, 45.00)\}$

This history captures the fact that first agent 3 submitted an ask of 45.00 and then buyer 1 accepted that ask.

Definition 6: Histories. After n messages have been sent, there is a history H_n of n ordered triples. For any message $m_{n+1} = (m_{n+1,1}, m_{n+1,2}, m_{n+1,3})$ that is sent one of six cases will hold:

Invalid ask or bid. A message $m_{n+1} = (i, 0, a)$ is not valid if $a \geq oa$. An invalid ask will not be included in the history. In effect the institution ignores messages that violate the spread reduction rule. Similarly a message $m_{n+1} = (0, j, b)$ is not valid if $b \leq ob$.

No ask outstanding. If no ask has been made since the last transaction, then there is no outstanding ask, and any ask a is valid. If in addition $m_{n+1,3} > ob$ then $h_{n+1} = m_{n+1}$.

No bid outstanding. Similarly, if no bid has been made since the last transaction, then there is no outstanding bid, and any bid, b is valid. If $m_{n+1,3} < oa$, then $h_{n+1} = m_{n+1}$.

Accept of ob. If $m_{n+1,1} \neq 0$ and $m_{n+1,3} = ob$, then seller $m_{n+1,1}$ is making an offer at ob, so that m_{n+1} is an accept of ob. The buyer's identity is found by looking back in H_n and finding the last h_k with $h_{k,2} \neq 0$, that is, $k^* = max\{k: h_{k,2} \neq 0\}$. Then $(h_{n+1,1}, h_{n+1,2}, h_{n+1,3}) = (m_{n+1,1}, h_{k^*,2}, ob)$

Accept of oa. If $m_{n+1,2} \neq 0$ and $m_{n+1,3} = oa$ then buyer $m_{n+1,2}$ is making a bid at oa, so the m_{n+1} is an accept of oa. The seller's identity is found by looking back in H_n and finding $k^* = \max\{k: h_{k,1} \neq 0\}$. Then $(h_{n+1,1}, h_{n+1,2}, h_{n+1,3}) = (h_{k^*,1}, m_{n+1,2}, oa)$.

Improving ask or bid. If $m_{n+1,3} \in (ob, oa)$ then m_{n+1} is either an improving ask, or an improving bid, and $h_{n+1} = m_{n+1}$.

1.3 Behavior

1.3.1 Overview

At any point in time agent i has a valuation for the risky asset, $c_{i,t}$, which represents that price at which the agent would either buy or sell.

Definition 7: Trading Profits. Given a trade $(i, j, h_{n+1,3})$ in which i sells j a unit of the risky asset with given valuations $c_{j,t}$ and $c_{i,t}$ then the trading profits to the seller and buyer are $(h_{n+1,3} - c_{j,t})$ and $(c_{i,t} - h_{n+1,3})$ respectively.

After observing the most recent message m_n, each agent decides whether to submit a bid or ask and the amount of the bid or ask. Every agent i assumes i's behavior will not affect the likelihood of a particular bid or ask resulting in a trade. Furthermore the agent assumes that the likelihood of a particular bid or ask resulting in a trade can be derived from information that the agent has observed and remembers from history.

We first provide the intuition of how the optimal bid and/or ask is determined by an agent. Suppose agent i is considering asking a. Then the profit if the ask is accepted is $(a - c_{i,t})$. The question for the agent then is "How likely is it that a will be accepted?" In past trading behaviour there may be instances in which a seller offered to sell at a and it was (not) accepted. And there is additional information in the environment. First an ask above a that was accepted is evidence an ask at a will be accepted. Also a bid above a that is made previously is evidence that ask a will be accepted. A similar intuition describes the history of past bids.

1.3.2 Implementation

Definition 8: Remembered History. For $H_n \in \mathcal{H}_n$, we make the following definitions.

Trade function. For a vector H_n, define a function T: $\mathcal{H} \to \{0,1\}^n$ by setting $T_k(H_n) = I_{\{h_{k,1} \cdot h_{k,2} > 0\}}(H_k)$. Then each component T_k of T indicates whether a trade occurred in the k^{th} element of history.

Number of trades. Let $x = (x_1, x_2, ..., x_n)$. For each n, define S_n: $\{0,1\}^n \to N$ by setting $S_n(x) = \sum_{k=1}^{n} x_k$. Then $S_n(T(H_n))$ is the number of trades resulting from the first n messages.

Remembered History. Let L be the memory length of a given agent. For fixed n and H_n, to simplify notation, let $S = S_n(T(H_n))$. Let n' be the position of trades $S-L$ for $S>L$, and let $n'=0$ if $S \le L$. Define $\widetilde{H}_n^{(L)} = \{h_{n'+1}, h_{n'+2}, ..., h_n\}$.

Deletion of oa and ob from history. Let $n''=\max\{k: T_k(H_n))=1\}$, i.e., n'' is the index of the most recent trade. Let n^* be the index of the lowest (most recent) ask in the vector $\{h_{n''+1}, h_{n''+2}, ..., h_n\}$. Let n_* be the index of the highest bid in the vector $\{h_{n''+1}, ..., h_n\}$. Note that if $T_n(H_n) = 0$,

$h_{n,3}$ is either the outstanding ask or the outstanding bid as a consequence of the spread reduction rule, and if $T_n(H_n)$ =1, then there is no outstanding bid and no outstanding ask. If $T_n(H_n)$ = 0 and $h_{n,1}$ = 0 then $h_{n,3}$ is the outstanding bid. If $h_{k,1} \neq 0$ for some $k \in \{n'+1,..., n-1\}$, then $n^* \neq \{\}$ and we define $H_n^{(L)}$ by $H_n^{(L)} \equiv \{h_{n'+1}, h_{n'+2},.h_{n^*-1}, h_{n^*+1},.., h_{n-1}\}$, that is, $H_n^{(L)}$ is $\widetilde{H}_n^{(L)}$ with h_{n^*} removed. The other case -- where $h_{n,3}$ is the outstanding ask is treated similarly. This is done because it is not known at time n if the outstanding ask or bid will be accepted. Then $H_n^{(L)}$ is the history remembered by agents with memory length L who observe the history H_n.

Sets of asks and bids. Let $D_n^{(L)}$ be the set of all asks and bids that have been made in $H_n^{(L)}$, i.e. $D_n^{(L)} \equiv \bigcup_{k\in\{n'+1,...,n\}\setminus\{n_*,n^*\}} \{h_{k,3}\}$.

Definition 9: Ask Frequencies. For each $d \in D_n^{(L)}$, let A(d) be the total number of asks that have been made at d, and let TA(d) be the total number of these that have been accepted. Let RA(d) = A(d)-TA(d) be the rejected asks at d.

For A(d), the counting procedure is as follows. For each $k \in \{n'+1,...,n\}\setminus\{n^*,n_*\}$, if $h_{k,3}$ = d, $h_{k,1} \neq 0$, and $h_{k,2}$ = 0, then A(d) is incremented by 1. If $h_{k,3}$ = d, and $T_k(h_k)$ =1, then h_k is either a taken ask or a taken bid. To determine which is the case, find $m^* = \min\{m \geq 1 : h_{k-m,3} = h_{k,m}\}$. If $h_{k-m^*,1} \neq 0$, then A(d) and TA(d) are incremented by 1. The rejected asks at d are given by RA(d)=A(d) -TA(d).

Definition 10: Bid Frequencies. For each $d \in D_n^{(L)}$, let B(d) be the total number of bids that have been made at d, and let TB(d) be the total number of these that have been accepted. Let RB(d) = B(d)-TB(d) be the rejected asks at d. The interpretations and counting procedures for B(d), TB(d), and RB(d) are analogous to those described in Definition 9 for asks.

Note. In what follows, the sets of asks and bids are frequently denoted D, with the subscripts and superscripts omitted. When agents have finite memory, that will be noted. After n messages have been sent the relevant set of asks and bids is $D_n^{(L)}$ and the relevant history is $H_n^{(L)}$.

1.3.3 Beliefs

Definition 11: Beliefs an ask a will be taken. For each potential ask $a \in D$, define

$$\ddot{p}(a) = (\sum_{d\geq a} TA(d) + \sum_{d\geq a} B(d)) / (\sum_{d\geq a} TA(d) + \sum_{d\geq a} B(d) + \sum_{d\leq a} RA(d))$$

Then $\ddot{p}(a)$ is an agent's belief that an ask a will be taken by a buyer. We assume that agents always believe that an ask at a =0.00 will be accepted with certainty, and that there is some valuation $M > 0$ such that $\ddot{p}(M) = 0$.

Definition 12: Beliefs a bid b will be taken. For each potential bid $b \in D$, define

$$\ddot{q}(b) = (\sum_{d \setminus \leq b} TB(d) + \sum_{d \leq b} A(d)) / (\sum_{d \setminus \leq b} TB(d) + \sum_{d \leq b} A(d) + \sum_{d \geq b} RB(d))$$

Then $\ddot{q}(b)$ is an agent's belief that a bid b will be taken by a seller. We assume that agents always believe that a bid at b=0.00 will be rejected with certainty, and that there is some valuation $M > 0$ such that $\ddot{q}(M) = 1$.

1.3.4 Spread Reduction Rule and Beliefs

The spread reduction rule has the effect of making the probability of a take for an ask $a \geq oa$ equal to 0 (where oa is the outstanding offer (from Definition 4). We denote this modification of $\ddot{p}(a)$ by $\tilde{p}(a)$, where $\tilde{p}(a) = \ddot{p}(a)$ if $a<oa$ and $\tilde{p}(a) = 0$ if $a \geq oa$. Similarly, $\tilde{q}(b) = 0$ for all $b \leq ob$. These facts are incorporated into agents' beliefs in the following definition.

Definition 13: Let $\tilde{p}(a) = \ddot{p}(a) \cdot I_{[0,0a]}(a)$ for each $a \in D$. That is $\tilde{p}(a) = \ddot{p}(a)$ if $a<oa$ and $\tilde{p}(a) = 0$ if $a \geq oa$. For all $b \in D$, let $\tilde{q}(b) = \ddot{q}(b) \cdot I_{[ob,M]}(b)$.

Observations. The function $\tilde{p}(a)$ is nonincreasing and the function $\tilde{q}(b)$ is non-decreasing.

1.3.5 Expected Surplus Maximization

At time t an agent has a certainty equivalent $c_{i,t}$. Should the agent make the ask a where a is taken from the set of previous bids and asks (absent ob and oa), the agent's expected surplus will be

$$E[\pi_{i,t}(a, c_{i,t})] = (a - c_{i,t}) \cdot \tilde{p}(a) \qquad (1)$$

The maximum expected gain over all such previous asks is given by

$$Max_{a \in (ob,oa)} E[\pi_{i,t}(a, c_{i,t})] \qquad (2)$$

Let a_j be the j^{th} member of bids and asks in the interval that solves (2).

Using linear interpolation the agent assesses probability functions $\tilde{p}$ between a_{j-1} and a_j and a_{j+1} and solves

$$Max_{a \in (a_{j-1}, a_{j+1})} E[\tilde{\pi}_{i,t}(a, c_{i,t})] \qquad (3)$$

where

$$E[\tilde{\pi}_{i,t}(a, c_{i,t})] = (a - c_{i,t}) \cdot \tilde{p}(a) \qquad (4)$$

to find an a^*.

On the other hand the agent may choose to bid. Should the agent make the bid b, the agent's expected surplus will be

$$E[\pi_{i,t}(b, c_{i,t})] = (c_{i,t} - b) \cdot \tilde{q}(b) \qquad (5)$$

The maximum expected gain over all bids is given by

$$Max_{b \in (ob, oa)} E[\pi_{i,t}(b, c_{i,t})] \qquad (6)$$

Let b_j be the j^{th} member of bids and asks in the interval that solves (6).

Using interpolation the agent assesses probability functions $\tilde{q}$ between b_{j-1} and b_j and b_{j+1} and solves

$$Max_{b \in (b_{j-1}, b_{j+1})} E[\tilde{\pi}_{i,t}(b, c_{i,t})] \qquad (7)$$

where

$$E[\tilde{\pi}_{i,t}(b, c_{i,t})] = (c_{i,t} - b) \cdot \tilde{q}(b) \qquad (8)$$

to find an b^*.

Since (3) and (7) can result in negative surplus the agent may choose to do no action at all.

Therefore the overall maximized surplus for agent i at time t is given by

$$S_{i,t} = \max\{ E[\tilde{\pi}_{i,t}(a^*, c_{i,t})], E[\tilde{\pi}_{i,t}(b^*, c_{i,t})], 0\} \qquad (8)$$

1.3.6 Timing of messages

At any point in time, t, there is a set of messages $\{h_1, h_2, ..., h_{t-1}\}$. Each of these messages was sent in the form of a bid, ask or taken bid or taken ask by a specific agent. The message at time t will reflect an action by one of the agents in the economy, and that action will maximize that agents expected surplus. How fast an agent moves to take an action is assumed to be a stochastic function of the amount of expected surplus that an agent can achieve. Let $\tau_{i,t}$ represent the time at which agent i sends a message. Then

$$P(\tau_{i,t} < \tau_{j,t}, \forall j \neq i) = \frac{s_{i,t}}{\sum_{j=1}^{n} s_{j,t}}$$

This probability then determines the likelihood an agents bid will be the one that appears in the history of the double auction.

2. Extensions

The results are based on a single risky asset economy with a simple information and return structure. There is no short selling and there is but a single risky asset. There are several natural ways these restrictions can be relaxed.

2.1 Valuation

Currently the valuation of the asset for agent i is given by $V_{i,t}(I_t, U_i, R_{t+})$ where I_t is the information structure at time t; U_i is the investors utility function, and R_{t+} are future returns. This valuation takes place independently of determining the actual bid in the auction. Thus there is room to incorporate many other functions, including for example functions from behavioural economics in the construction of this value function. The expanded model would read $V_{i,t}(I_t, VA_i, R_{t+}, w_i)$ where VA and w are the valuation and weighting functions discussed in prospect theory. Furthermore the asset could possibly have a much more complicated structure based on n periods of returns, as well as joint distributions on possible return paths as well as a more extended information and return structure; however given valuation construction process in which value of the item is determined prior to submitting a bit or ask, pricing behavior would proceed along similar lines in the double auction.

2.2 Short Selling

Suppose that there are fixed costs to borrowing the shorted sales, margin requirements and a bankruptcy procedure. We assume here that the borrowing could come from another trader. As above the valuation would have to take into account estimates of the chances of default (this could potentially come from price data), the cost of borrowing, as well as the initial endowment. Using this valuation it could then be employed in the same manner as valuations without short selling. Technically there would need to be an enforcement mechanism to the agreement which be an addendum to the double auction itself. It would simply execute any bankruptcy rule that would eventuate as a consequence of not meeting requirements.

2.3 Multiple Assets

Multiple assets with no short selling can be implemented assuming each agent can assess the certainty equivalent of a portfolio. Valuation for each asset becomes the difference in the certainty equivalent of the portfolio with and without the asset in it. Note using this approach a separate sell and buy value can be assessed. Given a valuation then the optimal action for a particular agent in each market can be determined and then the overall optimal action decided.

Then the agent that will move will be determined by the relative profits of the traders. Multiple assets will also involve multiple joint messages for the collections of assets and require a more extensive Bayesian updating (or its behavioral equivalent).

3. Discussion of Equlibria

An equilibrium in the Double Auction with Risky Assets relative to the double auction. We outline why the underlying behavioral strategies, along with the rules of the auction allow us to define an competitive equilibrium relative to the Double Auction with Risky Assets, and furthermore why such behavior is behavior that no agent has a reason to violate given his knowledge of other agents. In this sense we have a strategic equilibrium at point in time t.

3.1 A competitive equilibrium at time t relative to the Double Auction with Risky Assets

The demand in the double auction at any point is a single unit represented by the outstanding bid. The supply in the double auction at any point is a single unit represented by the outstanding ask. Market clearing means that there is no excess demand at the price traded. A competitive equilibrium exists in the double auction with risky assets at point in time t if each agent is maximizing that agent's expected utility with respect to his selection of bid and/or ask and if markets clear.

Agent Maximization: By construction of each agent's beliefs as well as each agent's choice of bid and asks, each agent is maximizing. See Sections 1.3.1 through 1.3.5 in the appendix.

.

Market Clearing: At point in time t there is only one outstanding bid and one outstanding ask. One of two cases holds: if the outstanding ask is greater than outstanding bid then there is no price at which the asset is traded and there is no excess demand; if the outstanding bid equals the outstanding ask then each unit offered is accepted and hence markets clear at that price. Note agents will never enter a bid (ask) greater (less) than the outstanding ask because it would not be surplus maximizing.

3.2 A strategic equilibrium at time t relative to the Double Auction with Risky Assets

Strategic Equilibrium: In a strategic equilibrium each agent is optimizing relative to his knowledge about agents characteristics and actions.

In this setting each agent does not know the distribution of risk coefficients in the economy. Each agent can only observe the bids, asks and takes in the market as well as his/her own private information. Each agent has no knowledge of the information of other agents nor other agents' decision rules. Yet the agent is optimizing relative to his inferred knowledge from prices.

3.3 Comparison of the price formation model with prices that fully reveal available information

In this section we describe equilibria that are calculated assuming that everyone knows all the private information that is held by individual agents. This is a situation in which prices fully reveal all available information. In this setting there will be trade because agents will have heterogeneous risk preferences. Furthermore we construct equilibria given endowment constraints of the traders. This assumption yields benchmark prices to match against our price formation process. In principle all of the equilibria described here could be achieved via a Walrasian auctioneer applied with the assumption that agents truthfully reveal their quantities demanded for every price announced by the auctioneer.

3.3.1 Demonstration

We now show how to determine equilibria in our setting when agents have endowments of riskless and a risky asset. Suppose there are 10 agents with reservation values for the risky asset 45.69, 47.00, 44.85, 44.14, 45.28, 46.85, 45.16, 45.01, 44.80, and 45.23 respectively. Suppose that endowments for each agent are 80 units of the riskless asset and 2 units of the risky asset. Then no individual can sell more than two assets. Furthermore given the minimal selling prices for each of the assets no agent will be able to buy more than 1 asset in this scenario. Thus we can rank the agents' reservation values from highest to lowest to determine a demand curve, and similarly we can determine the minimum at which each of the sellers will sell each of the seller's two units and determine an implicit supply curve. The computations lead to following demand and supply curves which yield an equilibrium price of 45.01 (see Figure 11). Some of the traders will sell and some will buy at this price.

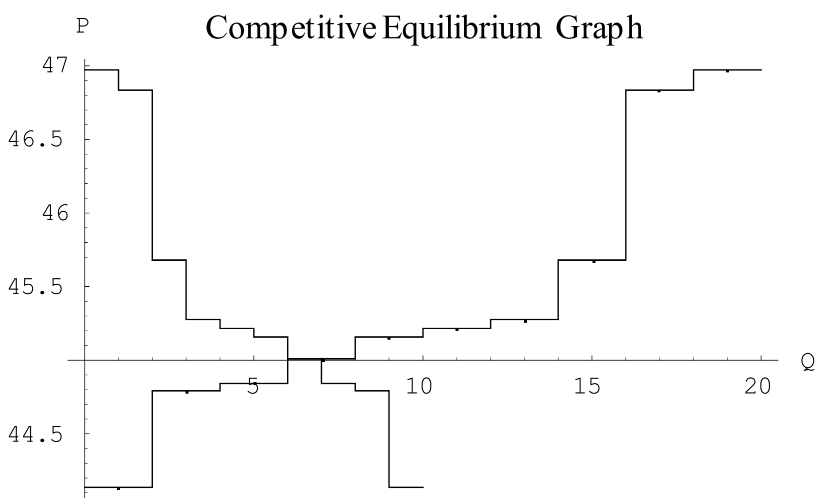

Figure 11 Deriving an Equilibrium Price

It is easy to talk about how price reflects information in this setting. Given we are working with the CARA utility function it is possible to compute and market utility function by looking at the average risk coefficient across all traders. The equilibrium price can be used to infer the underlying state probability by assuming that the equilibrium price is the certainty equivalent of a trader with the average risk coefficient and then solving for the appropriate probability which will be the belief of each trader. In this sense the price reveals the available information.

3.3.2 Comparisons of price formation process with equilibria predictions

Figure 12 shows the outcomes of the three scenarios originally graphed in Figures 5, 6 and 8. We have added equilibrium predictions for the different scenarios. The Public Information Efficiency and drift graphs show on the horizontal lines the equilibrium if every trader behaved as if that trader knew all available information (fully informed). For history varying from 100 to 400 prices range above the equilibrium line but are descending as more traders become privately informed. In periods 0 to 100 all agents have the same information and prices are near the fully informed equilibria. In the Forecast and Drift picture there is differential information from 150 to 225 and from 225 to 300. In both cases prices do not achieve the fully revealing price. Drift is a somewhat different story here however. Traders, trading on their own behalf for risk sharing purposes simply do uncover all the information in price. There are not irrational beliefs only beliefs are based on something less than complete market information.

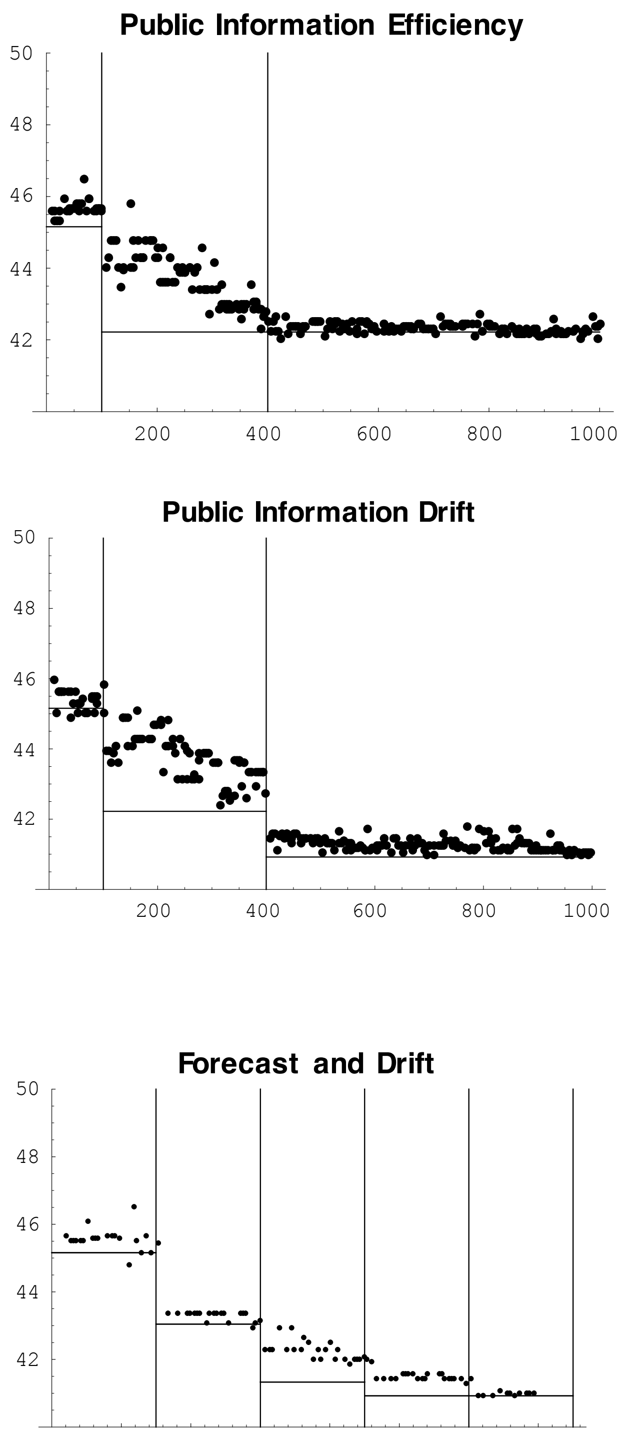

Figure 12 Price Data with Equilibrium Prices