SELF-GENERATING TRADE AND RATIONAL FADS: THE RESPONSE OF PRICE TO NEW INFORMATION

by

James Dow
Gary Gorton

40-89

RODNEY L. WHITE CENTER FOR FINANCIAL RESEARCH
The Wharton School
University of Pennsylvania
Philadelphia, PA 19104-6367

SELF-GENERATING TRADE AND RATIONAL FADS:

THE RESPONSE OF PRICE TO NEW INFORMATION*

James Dow
London Business School

Gary Gorton
The Wharton School
University of Pennsylvania

December 1989

Abstract

The dynamic behavior of security prices is studied in a setting where two agents trade strategically and learn from market prices. Each trader receives a private signal about fundamentals, the significance of which depends on the signal received by the other trader. In trading, each agent wants to deceive the other trader into revealing his signal, while not revealing his own signal. We show that trade is self-generating because agents learn the value of the asset only through observation of the market price. Uninformed agents, technical analysts, can also trade by charting past prices. These chartists ensure market efficiency. Equilibrium price paths of the model may display reversals in which all traders rationally revise their beliefs, first in one direction and then in the opposite direction even though no new information has entered the system. A piece of information which is initially thought to be bad news may be revealed, through trading, to be good news. This fad-like behavior results from rational strategic interaction and Bayesian inference. In this model security prices do not follow a martingale.

---

*We are grateful to Andy Abel, Franklin Allen, Fischer Black, Colin Camerer, Julian Franks, Douglas Gale, Larry Harris, Prem Jain, Joel Hasbrouck, Rich Kihlstrom, Murgie Krishnan, Anthony Neuberger, Krishna Ramaswamy, Ailsa Roell, Richard Rogalski, Chester Spatt, Rob Vishny and participants at seminars at the NBER Summer Institute, Philadelphia Federal Reserve Bank, New York University, The Wharton School, London Business School, University of Florida, Bristol University, the University of Wisconsin-Madison, Carnegie-Mellon University, Northwestern University, the European Finance Meetings, and the Hilliard-Lyons Symposium at Indiana University for many helpful suggestions. The second author thanks the NSF (#SES-8618130) and the Geewax-Terker Research Program in Financial Instruments for financial support.

I. Introduction

Keynes (1936) suggested that asset prices could be determined more by fads than by fundamentals. He compared the stock market to a beauty contest in which the goal is to predict the choices of the other judges, rather than basing the choice on 'fundamentals.' Recent empirical studies on financial markets have raised the possibility that asset pricing deviates from fundamentals because of the interaction among the traders in the market. These empirical studies uphold a long-standing tradition of conflict between the two points of view in the way that economists have viewed asset price behavior: on the one hand, prices should reflect fundamentals; on the other hand, prices may deviate from fundamentals because agents are primarily interested in the beliefs of other agents, and pay insufficient attention to the underlying economic realities.

If asset prices and fundamentals diverge, but eventually converge, then asset prices must display mean-reverting patterns over some horizon. Evidence of negative serial correlation in weekly stock returns has been found for individual stocks, although weekly returns on portfolios of stocks show positive autocorrelation (e.g., French and Roll (1986), Lo and MacKinlay (1988)). Lehmann (1988) forms portfolios of weekly "winners" and "losers" and shows that a strategy of shorting winners and buying losers will be profitable. Such evidence is widely viewed as being consistent with "fads" in asset prices. In an efficient market, new information is viewed as entering asset prices very quickly. Formally, this is described by the martingale hypothesis of asset prices. Thus, price movements which do not seem to correspond to the arrival of news appear anomalous. Mean reversion in asset returns, for example, is inconsistent with market efficiency.

Even if the diffusion of information into asset prices is protracted, mean reversion appears irrational. The reason is that information is typically modelled using a paradigm drawn from the statistical theory of point estimation, in which agents are imagined as receiving a signal equal to the true value of a parameter plus an error term (or, more generally, a signal which is positively associated with the true value). This "truth plus noise" paradigm implies that the best estimate of the true value of the parameter should be formed by averaging the individual signals. In applications to financial markets, the parameter of interest is the fundamental value of the asset. Consequently, agents who exchange information about their valuation will tend to converge to a common posterior valuation. It is not possible, in this framework, for exchange of information to lead to all valuations changing in the same direction. Thus fads are viewed as being inconsistent with rational behavior.

Additional evidence that asset pricing may depend, at least in part, on what agents believe other agents believe has been suggested by findings that trade begets trade. French and Roll (1986), Barclay, Litzenberger, and Warner (1988), among others, have found that stock return variances are higher during trading hours than during nontrading hours. Barclay, Litzenberger, and Warner also find an increase in price variability and volume associated with increased trading hours. Also, Meese (1986) found that "the ratio of return variance during trading hours to nontrading hours is about three for a sample of daily dollar-yen exchange rates." This evidence has been interpreted as suggesting that trade is "self-generating," that is, the opportunity to trade seems to lead to further trade. This seems to contradict the basic logic of rational expectations that at any point in time the price simultaneously aggregates all the existing information and communicates it to all agents.

Rational expectations is an equilibrium concept that allows agents to condition their actions on the equilibrium price which their actions simultaneously help to form. While rational expectations provides a convincing reduced form for the learning process it leaves no scope for the researcher to investigate empirical manifestations of learning as the learning happens in real time. From a formal point of view it would seem that rational expectations equilibrium could be implemented as the equilibrium of a game in which players submit demand functions. Dubey, Geanakoplos, and Shubik (1987) have shown that this approach is problematic because the strategy space is too complex. Not only is it implausible to suppose that agents could choose entire demand functions continuously over time, but hypothesizing that an auctioneer is a processor of these extremely detailed messages begs the question of how complex information is aggregated. Furthermore, actual trading institutions do not seem to correspond to this mechanism. These considerations suggest investigation of models where agents choose actions from smaller spaces, in particular, where their quantities are not price contingent.[1]

How is information communicated? We study how information diffuses into asset prices in a setting in which agents' actions and the formation of prices occur sequentially in real time. In other words, there is a time lag between an agent's choice of an action and the observation of the price. In the setting we study traders each receive private signals. Then these traders have the opportunity to trade repeatedly in a market with many small traders who have price sensitive liquidity preferences. Because they are price sensitive, the market is not infinitely deep. Therefore, even a single trader receiving a private signal will spread trade over all periods much like a monopoly owner of a durable resource who must choose an intertemporal pricing profile. The structure is similar to Kyle (1985) except that Kyle has only a single informed trader.[2]

With multiple informed traders there are two additional effects. First, each trader, knowing that the other informed traders are acting on the basis of their information, will make inferences from past market prices about the true value of the asset and combine this inference with his own signal. We say that trade is "self-generating" when the outcome of a round of trading, as summarized by the price, generates a further round of trading and determines the quantities traded in the next, even when no new information enters the model. Secondly, this learning process creates strategic possibilities. In particular, since other traders are learning from prices, each trader will want to conceal his own information, and further, will actually want to deceive his rivals into drawing the wrong inferences. For example, if a trader receives a signal that the asset is valuable, he would want to sell the asset so as to mislead other traders into thinking the asset is not worth much, and advantageously depress the future price. Of course, this effect conflicts with the trader's desire to profit immediately by buying the asset and, in equilibrium, the trader optimally balances these two effects.

Clearly, if informed strategic traders can learn from past prices, then so can outside observers. Such technical analysts would be prepared to enter the market and trade on the basis of their inferences until expected profits fell to zero. Of course, while this will dilute the value of the informed traders' information, they will always retain an informational advantage over the technical analysts. The idea that the price is equal to the expected value of the asset, conditional on all available public information is the standard notion of market efficiency. We express this rationality condition in our model by requiring that a technical analyst be indifferent between entering the market and trading or not.

What is a fad? Individual traders in a market may receive different private information, and therefore, may rationally update their beliefs as a result of learning about each other's beliefs. If they learn each other's beliefs, their posterior beliefs will tend to converge. It seems impossible for all agents to revise their beliefs in the same direction in response to their private information, and then to update their beliefs in the opposite direction upon learning each other's private information. For example, if one trader gets information that an asset is valuable, and another trader gets information that the same asset is not worth much, and they then learn each other's signal, they should revise their beliefs in opposite directions. The trader who received a high signal, and then learns that the other trader received a low signal, will revise his valuation downwards, while the trader who received a low signal will revise his beliefs upward. On the other hand, if traders who receive information decide that this information implies that an asset is worth very little, and then, through communicating with each other, change their minds and value the asset highly, this is commonly viewed as inconsistent with rational behavior. Therefore, we define a fad to be a reversal in price resulting from exchange of information.

In this essay we suggest that information can usefully be viewed as a more complex commodity than simply 'the truth plus noise.' Often, in financial market settings, information is difficult to interpret, and the significance of a piece of information may not be understood without other pieces of information. For example, suppose an analyst learns that a company president, who has been a good manager in the past, is changing jobs, moving to another company. Is this good news or bad news? Perhaps by combining this information with some other information, he can decide whether the manager is being forced out because of bad performance, or whether his excellent ability has attracted the attention of his new employers.

As a simple example of our information structure, consider two firms making production decisions. One firm makes computer hardware, and the other firm makes computer software. Each firm has a choice of two standards for its products. If the two firms choose the same standard, then their products will be more useful, and hence, the firms will be more highly valued. Now imagine two analysts, each of whom privately learns the choice made by one of the firms. The signal that each analyst receives is of little value by itself, but if it could be combined with the other analyst's signal it would be highly valuable.

This paper develops two distinct themes. The first concerns the formation of prices by a repeated sequential process of action, price formation, learning. The second theme is that information structures which differ from the "truth plus noise" paradigm may be both plausible descriptions of reality, as well being capable of yielding answers to the anomalies mentioned above. We stress that while we have combined these themes, each could be developed separately. When they are combined in our model fads can occur as a result of self-generating trade.

With our information structure, a price reversal can be the equilibrium manifestation of efficient communication of information. Suppose two traders each receive a signal in a setting similar to our example described above. Each trader's prior belief about the signal received by the other trader determines whether these private signals alone are marginally good news or marginally bad news. Recall that the coordination-like information structure implies that two pieces of apparently bad news actually signal a high value of the asset. Say that both traders receive bad news, and so, they sell the asset. This will depress the price and lead the two traders to infer each other's signal, and conclude that the asset is valuable. Then they will buy the asset and the price will rise. In other words, even though no new information has entered the system the price path reverses direction.

The paper proceeds as follows. In Section II we review and detail the literature on common knowledge from which we draw the information structure to be analyzed. Section III presents a variety of examples of our information structure. Section IV introduces the model and obtains the equilibrium. Section V discusses belief and price reversals. Section VI discusses extensions and concludes.

II. Common Knowledge, Trading and Information

In the model we study communication occurs via the process of trade between profit maximizing agents. Information transfer requires that agents actually trade. The relationship between trade and information has been studied in the literature on common knowledge. (See Aumann (1976), Geanakoplos and Polemarchakis (1982), and Geanakoplos (1988).) This literature is relevant to our paper for two reasons. First, although general communication processes have been studied in this literature, Milgrom and Stokey (1982) have shown that no trade can occur when it is common knowledge that the information exchange mechanism is a purely speculative process. Paradoxically, although these information exchange mechanisms seem realistic, there is no reason for them to ever be initiated since they cannot lead to trade. Our goal is to investigate self-generating trade in a model where information exchange occurs via actual trade. Secondly, the common knowledge literature is relevant because our information structure is derived from that literature.

An event is common knowledge if everyone knows it has occurred, everyone knows that everyone knows it has occurred, and so on. An example can illustrate the idea of common knowledge (Barwise (1981) and Moses, Dolev and Halpern (1986)). Suppose three children are in a room. Two children have mud on their foreheads; the third does not. Each child cannot observe whether he has mud on his own head. An adult enters the room and says: "At least one of you has mud on your head. All children with mud on their heads should leave the room." The two children with mud on their heads can each observe two other children, one with mud and one without. They cannot, therefore, infer anything about their own foreheads. The third child observes two other muddy heads and so he cannot infer anything either. Then a second round of inferences occurs. A child who previously observed two other children, one with mud and one without, can put himself in the place of the child he saw with mud on his head. He thinks: "That child did not leave the room and so he must have seen at least one other child with mud on his head. But he saw me and the child with no mud and did not leave the room. Therefore I must have mud on my head." After this second round, therefore, the two muddy children will leave.

The example illustrates the difference between knowledge and common knowledge. Before the adult's announcement, each child knew that there was at least one muddy child, but this fact was not common knowledge. For example, neither of the muddy children could rule out the possibility of himself being clean. If he were clean himself, then the single muddy child would not know that there was at least one muddy child. Therefore, each muddy child does not know that the other muddy child knows that there is at least one muddy child.

In the computer science literature, as in the above example, communication is direct. In economics communication occurs indirectly by inference from agents' actions: an offer to trade can itself convey information. The winner's curse in auctions is an example since winning, itself, conveys information about the valuations of the other bidders and hence about the value of the object. In Akerlof's 'lemons market' the fact that a seller is willing to sell the used car, itself, indicates that the car is a lemon. (See Akerlof (1970).) Similarly, in Rothschild and Stiglitz (1976), if an agent is willing to buy insurance, then it is likely that that individual is a high risk person. Finally, suppose you offer a car dealer a price which is immediately accepted, without hesitation. Then you would likely conclude that you were not fully informed about the car's value, and that your offer was too high.

In the above examples, the gains from trade may counteract possible 'lemons' problems. In a purely speculative financial market there is no countervailing gain. Tirole (1982) and Milgrom and Stokey (1982) have shown that purely speculative trade cannot occur if it is common knowledge that ex ante allocations are Pareto-efficient. In other words, an agent approached and offered a trade knows, by virtue of the fact that it is beneficial to the person making the offer, that the trade cannot be beneficial to him. These 'no trade' results apply to models of asset pricing in which new information is assumed to become common knowledge. Thus, these models are valuing assets in a framework in which no trade would occur. If no trade occurs, then there are no predictions about the volume of trade, and no explanation of how the information gets into the equilibrium prices.

The 'no trade' results have been linked to a process of information transfer in which one party learns from the other party's offer to trade. Geanakoplos and Polemarchakis (1982) display a process under which agents by simple iterative communication of their beliefs converge to a common posterior. Similarly, Sebenius and Geanakoplos (1983) consider an iterative mechanism for information transfer which is a series of tentative offers to gamble. If one trader offers a gamble to the other, this rules out some of the signals that the trader could potentially have received. Similarly, a trader's acceptance of an offer to gamble, rules out some possible signals. By successively stating whether they would take the bet or not two agents eventually learn the true state of the world, and hence do not gamble or trade.

Since it is common knowledge that no trade will occur in the above models, the information transfer processes are redundant. In other words, there is no point for agents to participate in the communication process since they already know that the outcome will not affect them. Nevertheless these communication mechanisms are interesting and suggestive. However, they have two special features. First, the offers or statements are not binding commitments. Second, the players are honest when, in general, it will benefit them to give deceptive responses to the offers. One way of making responses credible is to require that trade actually occur, that is, that each trader must 'put his money where his mouth is.' Of course, in a market setting, the size of the bet is also a variable.

Our objective is to embed a communication mechanism into a realistic trading process. In order for trade to be possible, we must have a nonzero sum game. This, of course, is the exact issue confronted by Grossman and Stiglitz (1976, 1980) and later by Kyle (1985). We adopt Kyle's solution, namely, we suppose the presence of 'liquidity' traders who are willing to pay a premium for immediacy. We then consider a market setting in which two players, similar to those discussed above, each receive a signal and engage in strategic trade. In the next section, the examples of the information structure are given.

III. Examples

In this section we present two examples which illustrate the type of information structure we study in this paper. These examples are stylized, but hopefully they are close enough to reality to carry some conviction. The underlying structure required is that the two pieces of information received should be complementary. The examples are intended to show that the required complementarity can plausibly happen.

Example 1: Cray Research

Cray Research is a leading manufacturer of supercomputers, which up until now has remained closely identified with its founder, Mr. Seymour Cray. Mr. Cray has recently been engaged in a large research project to develop a new computer technology, based partly on the use of gallium arsenide chips. This is an ambitious project which, if successful, will probably lead to the next generation of supercomputers and give the company a strong competitive position in the market. But it is also risky and, if unsuccessful, will not only cost the company a large amount of money but may divert research resources away from development of products based on the existing technologies. Because the technology being investigated is radically different, if sucessful it will supersede the existing technology, but if not, the research will provide very few improvements to the existing technology. This raises the possibility that the company may decide not to pursue the project, and that Mr. Cray will leave Cray Research to work on it elsewhere.

We model the value of the shares as follows. If Mr. Cray stays and the technology is successful, the shares will have a high value since the company will have the state-of-the-art technology. If he stays and the project fails the shares will have a low value since the company will have spent a large amount of money on the research, and will have neglected the development of its existing technology. If he leaves and the technology proves to be viable, the company will also have a low value since it will be left with an obsolete product. Finally, if he leaves but the project fails, the company will have a high value since it will retain its position as a leading supercomputer manufacturer.

Example 2: Foreign Exchange Rates

An important application of our model concerns public information. The usual view of publicly announced information is that it is immediately incorporated into price and, therefore, it cannot result in trade. Researchers, however, have found a positive relationship between trading volume and the public release of information. For example, Bamber (1986) shows that the announcement of unexpected firm earnings leads to increased trading activity. This could be explained if the combination of a public announcement with private knowledge made the announcement equivalent to a private signal. If different traders have different private knowledge or expertise, then they will interpret public signals in different ways. By learning the interpretations of other traders, they can make use of their specialized knowledge. When all the private interpretations of all the traders is generally known, the asset can be priced efficiently making use of the combined expertise of the individual analysts.[3]

In this example, we consider the value of the US dollar in an environment where there is an informal agreement among major industrial countries to keep exchange rates within certain ranges. Suppose that the dollar has been drifting downwards, and that there is currently uncertainty about whether the agreement will be maintained by concerted action among the central banks of these countries, or whether the authorities will allow the dollar to fall below the target range. Further suppose that for political reasons the maintenance of the agreed rate requires a degree of cooperation among the governments of the countries concerned. Thus, relatively minor international disputes concerning other issues such as trade policy, or defense policies within a military alliance, will hinder exchange rate cooperation and cause the dollar to fall in value.

The public information that triggers the exchange rate adjustment process is some event which has a bearing on such an issue. For instance, suppose that the US Congress is about to consider new trade legislation and that a draft of the proposed legislation is published. This legislation appears to have implications for imports of Japanese automobiles, but it is not easy to predict whether it will in fact have the effect of restricting imports, or whether the protectionist clauses are likely to be weakened before the bill is passed. Thus, some expert understanding of the political situation is needed to interpret the news and judge whether it signals a period of protectionist policies, or whether the appearance of protection is an artefact of everyday internal political wrangling. An American analyst can judge this more accurately than a Japanese analyst. However, it is not easy for American analysts to judge the reaction in Japan to the proposed piece of legislation.

If the current political mood in Japan is protectionist, a US law which restricts car imports may not be unwelcome to the Japanese government since it may weaken the pressure to open up certain Japanese markets to foreign trade. But if the Japanese government has been preparing to liberalize trade restrictions, a protectionist move by the US may cause a disagreement between the two governments. A Japanese analyst can judge fairly accurately which of these would be the case, but is less able to predict whether the US legislation really is protectionist or not. If the legislation does not actually reflect a shift to a more protectionist policy and the US government is actually in favor of increased trade liberalization, the situation is reversed. Protectionist attitudes by the Japanese government would lead to disagreement between the countries, while liberal trade policies would permit cooperation.

The combination of a public signal with private expertise has an effect equivalent to the two types of analysts receiving private signals.

IV. The Model

We consider a single asset market with three groups of traders. There are two risk neutral strategic traders, a large number of liquidity traders and a large number of risk neutral technical analysts. Throughout, the strategic traders will be referred to as 'traders.' There are three trading periods, and then a final period in which the value of the asset is realized and consumption takes place.

Within each trading period all agents submit orders and then these orders are executed at the market clearing price. The sequence of events is as follows. In period 0 the market opens but no information has been received. The period 0 price will serve as a benchmark for comparing subsequent prices. Following this round of trading the two strategic traders each receive a private signal about the value of the asset. In period 1 the market reopens. The outcome of this round of trading will reflect the private information, and therefore both the chartists and the strategic traders will learn from the period 1 price. In period 2 the market opens for a final round of trading, giving traders an opportunity to trade based on their inferences from the period 1 price.

The information received by each strategic trader after the period 0 trading round consists of a signal A or B. If both get the same signal, the value of the asset will be high. If they receive different signals the value of the asset will be low. We simplify and suppose that the value of the asset may be one or zero. The chance of signal A, α, is independent of the signal received by the other trader and is higher than $\frac{1}{2}$. Thus A represents "good news" on its own, but in order to evaluate its full implications for the value of the asset it is necessary to know the signal received by the other trader. There is assumed to be a possibility that the true value of the asset becomes publicly known between periods 1 and 2. If this happens, the strategic traders have no advantage in period 2 trading. Define δ to be the probability that this does not occur.

We will use the following notation. Q_t^j will denote the quantity sold by trader j in period t. Thus, a positive value of Q_t^j represents a sale and a negative value of Q_t^j a purchase. Because the game is symmetric we look at the game from Trader 1's point of view and describe Trader 1's best response to Trader 2's strategy. In the discussion below we will focus on the symmetric Nash equilibrium of the game. Trader 1's strategy is a rule for choosing Q_t^1. Thus, a strategy is a quintuple $\{g, g^A, g^B, h^A, h^B\}$, where the first three elements are real numbers, and the last two elements are real-valued functions on the real line. The first element, g, is the quantity sold in period 0. In period 1 the quantity sold depends on the signal received: g^A is the quantity sold if the signal A is received, and g^B is the quantity sold if signal B is received.

In period 2, the quantity sold depends on the signal received, on the beliefs concerning the other trader's signal, and on the prediction of the other trader's period 2 quantity. The period 1 price conveys information about the other trader's quantity, and hence, his signal. It also conveys information about the quantity the other trader will sell in period 2. Thus, the quantity chosen in period 2 is a function of the signal received and of the period 1 price.

The payoff to a trader is the sum of the expected payoffs in each period. The payoff in each period is the difference between the value of the assets acquired and the amount of money paid for them in the case of a purchase, and vice versa in the case of a sale. Throughout we will ignore budget constraints and short sale restrictions in order to concentrate on optimal strategies being restricted only by the information revealing properties of the market price.

The liquidity traders' preferences are represented by a downward sloping linear demand curve in each period. The interpretation of this demand curve is as follows. Each period these traders have transitory new real investment opportunities. They compare the return on an investment in these projects with the return offered by the security market. They may be willing to sell a financial asset which they believe to be undervalued, in order to invest in a superior alternative. Thus, they are willing to pay a premium to satisfy their 'liquidity needs,' but they are not willing to pay any price. On the contrary, this liquidity preference results from rational consideration of the alternative choices. This leads to the downward sloping demand curve of the liquidity traders.[4]

We also assume that the downward sloping demand curve is noisy. This assumption is made for technical reasons, namely, to insure the existence of equilibrium in pure strategies in our model. Unlike the model of Kyle (1985) noise is not an intrinsic feature of the liquidity traders. The randomness in aggregate liquidity trader demand may, however, be interpreted as exogenous liquidity needs that affect the whole population.

The expected inverse demand curve is given by:

$$\pi(Q) = a - bQ \tag{1}$$

where Q is the total quantity sold by the strategic traders. The realized price, P, will vary from the expected price due to fluctuations in liquidity trader demand: $P = \pi(Q) + \varepsilon$, where ε is normally distributed with zero mean and variance σ^2.[5] Subscripts will denote the different periods: thus, P_0 is the period 0 price, etc.

The chartists are agents who trade on the basis of publicly available information, i.e., past prices. There is free entry of such technical analysts into the model. They choose quantities in each trading period to maximize expected trading profits. We do not explicitly model their chosen quantities. Rather, we describe the effect of their trading as a shift in the demand curve faced by the strategic traders in each trading round.[6] The details are explained below.

In each trading period, the orders from the different groups of agents are aggregated and executed at the market clearing price. Nash equilibrium of the model requires that each strategic trader choose a strategy that maximizes his payoff, given the other strategic trader's strategy, and given the strategies of the other agents in the model. We restrict attention to symmetric equilibria.

Nash equilibrium for the technical analysts is equivalent to a market efficiency condition. In other words, since there is free entry of chartists, the expected price in each period equals the expected value of the asset, conditional on past prices (but not on private signals). This market efficiency condition is reflected in the intercept of the demand curve faced by the strategic traders. Thus, a_0 is the intercept in period 0 and a_1 is the intercept in period 1. In period 2 the intercept $a_2(P_1)$ is conditioned on the price history.

We characterize the properties of the equilibrium expected price path for given realizations of the private signals. This then allows us to investigate the time-series properties of prices averaged across the different constellations of private signals. Notice that because the liquidity traders' demand slopes down the market is not infinitely deep. Therefore, a trader who receives private information has an incentive to spread his trades over both remaining periods. Nevertheless, he will trade in the first period. In doing so he will provide information about the signal received to the other trader, and to the chartists. Symmetrically, however, he will be able to make inferences from the period 1 price about the signals received by the other trader. Thus, there is an additional incentive for traders to restrict their period 1 quantities, namely, a desire to deceive other traders.

A. Equilibrium in Period 0

To introduce the basic structure of the model we start by considering equilibrium in period 0. Because no information has yet arrived this period is essentially separate from periods 1 and 2 and will serve as a benchmark for describing the equilibria in subsequent periods.

The equilibrium in period 0 is a standard Cournot-Nash equilibrium. The payoff to Trader 1 is:

$$[\pi(Q_0^1 + Q_0^2) - [\alpha^2 + (1 - \alpha)^2]Q_0^1. \qquad (2)$$

This is just the expected difference between the price and the value of the asset multiplied by the quantity sold. Trader 1 must choose a quantity, Q_0^1, to maximize his payoff. The reaction function is given by the first-order condition that marginal cost equal marginal revenue:

$$\alpha^2 + (1 - \alpha)^2 = \pi(Q_0^1 + Q_0^2) + Q_0^1\pi'(Q_0^1 + Q_0^2)$$

$$= a_0 - b(Q_0^1 + Q_0^2) - Q_0^1 b \ . \qquad (3)$$

The Nash equilibrium condition for $g = Q_0^1 = Q_0^2$ is defined by:

$$g = [a_0 - \alpha^2 - (1 - \alpha)^2]/3b \ . \qquad (4)$$

And the market efficiency condition requires that, in period 0, $a_0 = \alpha^2 + (1 - \alpha)^2$. Then the price will be equal to the expected value, and the equilibrium solution will be given by $g = 0$. This means that the strategic traders will not trade until the private information arrives.[7]

B. Learning from the Price

In equilibrium, agents will learn from the period 1 price. Therefore, in order to describe the equilibrium strategies, we must give the details of this learning process.

Because the traders' period 1 quantities depend on the signals they received, the period 1 price will reveal information about these signals. This information will affect the period 2 equilibrium in two ways. First, the strategic traders will use this information to update their beliefs about the signal received by the other trader, and hence, about the value of the asset. Furthermore, when a strategic trader observes the period 1 price he knows how much he put on the market and will take this into account in estimating the quantity sold by the other trader. Second, chartists can use the price information to update their beliefs about the value of the asset. However, their inference is less precise. They received no signal of their own, so they must estimate the signals received by both of the strategic traders. Unlike the strategic traders, when making this inference, they cannot correct for one of the period 1 quantities. They must simultaneously estimate both of the quantities. In this section we describe both of these types of learning. We start by describing how a strategic trader learns.

Prior to trade in period 1 the trader's valuation of the asset depends only on the signal received. Since the asset is worth 1 if both traders receive the same signal, and is worthless otherwise, the expected value of the asset is the probability that the other trader has received the same signal. Let β^S be the trader's belief that the other trader has the same signal, where S represents the trader's own signal. Thus, S = A or B and $\beta^A = \alpha$ and $\beta^B = 1 - \alpha$. We also define $\bar{S}$ by $\bar{S}$ = A when S = B and $\bar{S}$ = B when S = A.

After period 1 each trader revises his belief about the value of the asset by drawing inferences from the period 1 price. In doing so the trader must correct for the amount of the asset he put on the market in period 1. Let $\beta^S(P_1, Q_1^j)$ be the conditional belief of trader j after he has traded quantity Q_1^j and observed price P_1. By Bayes' Rule:

$$\beta^S(P_1,\ Q_1^j) = \text{Prob}(S|P_1,\ Q_1^j)$$

$$= \frac{\lambda(P_1|Q_1^j,\ g^S)\beta^S}{\lambda(P_1|\ Q_1^j)}$$

where the notation $\lambda(P_1|x)$ represents the likelihood of price P_1 conditional on one of the traders choosing quantity x, while $\lambda(P_1|x,\ y)$ is conditional on one of the traders choosing x and the other trader choosing y. Making use of the fact that:

$$\lambda(P_1|Q_1^j) = \lambda(P_1|Q_1^j,\ g^S)\beta^S + \lambda(P_1|Q_1^j,\ g^{\bar{S}})(1 - \beta^S),$$

we have:

$$\beta^A(P_1,\ Q_1^j) = \text{Prob}(A|P_1,\ Q_1^j) = \frac{\lambda(P_1|Q_1^j,\ g^A)\alpha}{\lambda(P_1|Q_1^j,\ g^A)\alpha + \lambda(P_1|Q_1^j,\ g^B)(1 - \alpha)}$$

$$= \frac{\phi\left(\frac{P_1 - \pi(g^A + Q_1^j)}{\sigma}\right) \alpha}{\phi\left(\frac{P_1 - \pi(g^A + Q_1^j)}{\sigma}\right) \alpha + \phi\left(\frac{P_1 - \pi(g^B + Q_1^j)}{\sigma}\right) (1 - \alpha)} \quad (5)$$

where ϕ represents the density for a standard normal random variable.

We now characterize more precisely how a trader who changes his period 1 quantity must correct to allow for the change in the distribution of price. Since the demand function is a straight line, a shift in quantity simply results in a fixed change in price regardless of the value of the random shock. The correction is therefore straightforward, as we now show:[8]

<u>Lemma 1:</u> $\beta^S(P_1, Q_1^j) = \beta^S(P_1 + bQ_1^j, 0)$. (6)

<u>Proof:</u> Since π is linear with slope -b:

$$\lambda(P_1 + bQ_1^j | g^S) = \phi\left(\frac{P_1 + bQ_1^j - \pi(g^S)}{\sigma}\right)$$

$$= \phi\left(\frac{P_1 - \pi(Q_1^j + g^S)}{\sigma}\right) = \lambda(P_1 | Q_1^j, g^S) \quad ||$$

Next, we consider the inferences of the chartists. From the chartists' point of view there are four possible outcomes: both traders received A signals, both received B signals, or the two traders received different signals. The price can be used to assign posterior probabilities to these four outcomes. Let γ_{AA} be the posterior probability of the event that the two traders both received signal A. By Bayes' Rule,

$$\gamma_{AA}(P_1) = \frac{\lambda(P_1 | 2g^A)\ \alpha^2}{\lambda(P_1 | g^A, g^A)\alpha^2 + \lambda(P_1 | g^B, g^B)(1 - \alpha)^2 + \lambda(P_1 | g^A, g^B) 2\alpha(1 - \alpha)} \quad (7)$$

$$= \frac{\phi\left(\frac{P_1 - \pi(2g^A)}{\sigma}\right)\alpha^2}{}$$

We similarly define $\gamma_{BB}(P_1)$ and $\gamma_{AB}(P_1) = \gamma_{BA}(P_1)$.

The value of the asset to the chartists is the posterior probability that the period 1 total quantity traded was $2g^A$ plus the posterior probability that quantity was $2g^B$. Call this number $\gamma(P_1)$. Then:

$$\gamma(P_1) = \gamma_{AA}(P_1) + \gamma_{BB}(P_1)$$

The chartists must also predict the quantity of the asset which will be sold in period 2 by the strategic traders. Let $\tilde{h}(P_1)$ be this expected quantity. Then:

$$\tilde{h}(P_1) = \gamma_{AA}(P_1)(2h^A(P_1)) + \gamma_{BB}(P_1)(2h^B(P_1)) + 2\gamma_{AB}(P_1)(h^A(P_1) + h^B(P_1)) \quad (8)$$

The above descriptions of learning by the strategic traders and the chartists will now be embedded in the description of the equilibrium trading process.

C. Equilibrium in Period 2

Now let us consider equilibrium in periods 1 and 2. While period 0 can be solved separately, periods 1 and 2 are connected because the period 1 price will influence the period 2 trades. Our analysis proceeds recursively by first fixing arbitrary period 1 strategies, and deriving period 2 equilibrium strategies. Then, the period 1 equilibrium strategies are determined given the dependence of the period 2 equilibrium on period 1. This is the standard backward induction procedure for finitely repeated games.

In period 2 equilibrium Trader 1 chooses a quantity as a function of his own signal, S, his beliefs concerning the value of the asset, $\beta^S(P_1,Q_1^1)$, and his prediction of the other trader's period 2 quantity. If Trader 2 received signal S, then $h^S(P_1, g^S)$ is Trader 2's quantity in period 2. If Trader 2 received signal $\bar{S}$, then $h^{\bar{S}}(P_1, g^{\bar{S}})$ is the quantity. Assuming that Trader 2 is playing his equilibrium strategy, $\{h^A, h^B\}$, in period 2, this predicted quantity is:

$$\beta^S(P_1,Q_1^1)\, h^S(P_1,g^S) + (1 - \beta^S(P_1,Q_1^1))\, h^{\bar{S}}(P_1,g^{\bar{S}})$$

Thus, Trader 1's period 2 objective is to:

$$\max_{Q_2^1}\{\beta^S(P_1,Q_1^1)[\pi(Q_2^1+ h^S(P_1,g^S) - 1] +(1 - \beta^S(P_1,Q_1^1))[\pi(Q_2^1+ h^{\bar{S}}(P_1,g^{\bar{S}})]\}Q_2^1 \qquad (9)$$

Notice that expression (9) is simply the difference between the expected price and the expected value, multiplied by the quantity. The first order condition for Problem (9) is given by:

$$\beta^S(P_1,\ Q_1^1) = \beta^S(P_1,\ Q_1^1)[\pi(Q_2^1 + h^S(P_1,g^S) + Q_2^1\ \pi'(Q_2^1 + h^S(P_1,g^S))] +$$

$$(1 - \beta^S(P_1,\ Q_1^1))[\pi(Q_2^1 + h^{\bar{S}}(P_1,g^{\bar{S}}) + Q_2^1\ \pi'(Q_2^1 + h^{\bar{S}}(P_1,g^{\bar{S}}))]$$

This is the standard condition that the marginal cost equal the marginal revenue. Since $\pi(Q_2^1 + Q_2^2) = a_2 - bQ_2^1 - bQ_2^2$, the reaction function can be written:

$$Q_2^1 = [a_2 - \beta^S(P_1,Q_1^1)(bh^S(P_1,g^S) + 1) - (1 - \beta^S(P_1,Q_1^1,S))(bh^{\bar{S}}(P_1,g^{\bar{S}}))]/2b \qquad (10)$$

Equation (10) defines two reaction functions for Trader 1: one if signal A was received, and the other if signal B was received. A symmetric equilibrium strategy in period 2, $\{h^A, h^B\}$, is implicitly defined by the Cournot-Nash equilibrium condition that it should be the solution to Problem (9) against itself. This defines two equations for Trader 1 which depend on Trader 2's period 1 strategies, g^A and g^B, as well as his own period 1 quantity, Q_1^1:

$$h^A(P_1,Q_1^1) = \frac{a_2(P_1)}{2b} - \beta^A(P_1,Q_1^1)[\frac{h^A(P_1,g^A)}{2} + \frac{1}{2b}] - (1 - \beta^A(P_1, Q_1^1))[\frac{h^B(P_1, g^B)}{2}] \quad (11)$$

$$h^B(P_1,Q_1^1) = \frac{a_2(P_1)}{2b} - (1 - \beta^A(P_1,Q_1^1))[\frac{h^B(P_1,g^B)}{2} + \frac{1}{2b}] - \beta(P_1,Q_1^1,A)[\frac{h^A(P_1, g^A)}{2}] \quad (12)$$

There are two similar equations for Trader 2 which depend on Trader 1's period 1 strategies and on Trader 2's period 1 quantity. Notice that the agents may be pursuing different period 1 strategies, because these are not necessarily equilibrium strategies. Thus, there are four equations which define the period 2 symmetric equilibrium strategies.

To simplify the computation of the solution we first assume that both traders use the same period 1 strategies, g^A and g^B.[9] Then these four equations are replaced by two, which depend on g^A and g^B:

$$h^A(P_1,g^A) = \frac{a_2(P_1)}{2b} - \beta^A(P_1,g^A)[\frac{h^A(P_1,g^A)}{2} + \frac{1}{2b}] - (1 - \beta^A(P_1, g^A))[\frac{h^B(P_1, g^B)}{2}] \quad (13)$$

$$h^B(P_1,g^B) = \frac{a_2(P_1)}{2b} - (1 - \beta^A(P_1,g^B))[\frac{h^B(P_1,g^B)}{2} + \frac{1}{2b}]$$

$$- \beta(P_1,g^B,A)[\frac{h^A(P_1, g^A)}{2}] \qquad (14)$$

These equations can be solved simultaneously to obtain period 2 equilibrium strategies:

$$h^A(P_1,g^A) = \frac{a_2(P_1)}{3b} + [\frac{1 - \beta^A(P_1,g^B) - 4\beta^A(P_1,g^A) + 2\beta^A(P_1,g^A)\beta(P_1,g^B)}{3b\ (2 - \beta^A(P_1,g^B) + \beta^A(P_1,g^A)\ )}] \qquad (15)$$

$$h^B(P_1,g^B) = [\frac{2a_2(P_1) + \beta^A(P_1,g^B) + \beta^A(P_1,g^B)\beta^A(P_1,g^A) + [\beta^A(P_1,g^A)]^2}{3b\ [2 + \beta^A(P_1,g^A) - \beta^A(P_1,g^B)]}] - \frac{1}{3b} \qquad (16)$$

We have now derived the period 2 equilibrium strategies for the special case where both traders use the same period 1 strategies. If a trader actually uses a different strategy, $Q_1^1 \neq g^S$, the effect on his period 2 equilibrium strategy is straightforward. Recall that the linear demand curve implies that the only effect of changing quantity in period 1 is to shift the price by a constant amount. Therefore, it follows immediately from Lemma 1 that:

$$h^A(P_1, Q_1^1) = h^A(P_1 + b(Q_1^1 - g^A), g^A)$$

This completes the description of the period 2 equilibrium strategies for arbitrary period 1 quantities.

In addition to the above condition for the strategic traders, equilibrium also requires that the chartists earn zero expected profits. In choosing an amount to trade in period 2, chartists rely on inferences from the period 1 price. Thus, equilibrium implies that, conditional on P_1, the chartists' expectation of the strategic traders' volume be such that the expected price equals the expected value of the asset. The expected value, $\gamma(P_1)$, and the expected quantity, $\tilde{h}(P_1)$, were derived in Section IV.B, above. Market efficiency requires:

$$a_2(P_1) = \gamma(P_1) + b\tilde{h}(P_1) \tag{17}$$

This defines the intercept of the period 2 inverse demand curve.

At the beginning of period 2, Trader 1's expected equilibrium period 2 profits depend on the period 1 price and the quantity he chose in period 1, as well as Trader 2's period 1 strategy. This value is given by:

$$V^S(P_1, Q_1^1) \equiv \{\beta^S(P_1, Q_1^1)[\pi(h^S(P_1,g^S) + h^S(P_1, Q_1^1)) - 1] + (1 - \beta^S(P_1, Q_1^1))[\pi(h^S(P_1,Q_1^1) + h^{\bar{S}}(P_1,g^{\bar{S}}))]\}h^S(P_1,Q_1^1) \tag{18}$$

We can derive a more compact expression for the value. Again, we use Lemma 1 and start by supposing that the trader received signal A and chose $Q_1^1 = g^A$. Then (18) becomes:

$$V^A(P_1, g^A) = \{\beta^A(P_1, g^A)[a_2(P_1) - 2bh^A(P_1,g^A) - 1] + [1 - \beta^A(P_1, g^A)][a_2(P_1) - bh^A(P_1,g^A) - bh^B(P_1,g^B)]\}h^A(P_1,g^A)$$

which, substituting equations (15) and (16), equals:

$$[h^A(P_1,g^A)]^2 b \tag{19}$$

Similarly, if Trader 1 has received signal B, then (18) becomes:

$$[h^B(P_1,g^B)]^2 b \tag{20}$$

It follows from Lemma 1:

$$V^A(P_1,\ Q_1^1) = V^A(P_1 + b(Q_1^1 - g^A),\ g^A).$$

And therefore,

$$V^S(P_1,\ Q_1^1) = [h^S(P_1,\ Q_1^1)]^2 b. \tag{21}$$

To understand (21) consider a standard model of Cournot duopoly with linear inverse demand with intercept a and slope -b, and constant marginal cost, c. In this model equilibrium quantities are $Q = (a - c)/3b$, for each trader, and the equilibrium price is $P = c + (1/3)(a - c)$. Profits per unit are $(1/3)(a - c)$ and equilibrium profits are $(a - c)^2/9b = Q^2 b$ for each trader. Equation (21) is, therefore, an extension of the standard model to our framework.

D. Equilibrium in Period 1

In period 1 a trader's objective is to maximize the sum of expected period 1 profits, plus expected period 2 profits which are affected by period 1 price. Suppose Trader 2 uses strategies g^A and g^B in period 1. We then solve for Trader 1's best period 1 strategies in response to these. In symmetric Nash equilibrium g^A and g^B must be best responses to themselves.

The instantaneous payoff in period 1 is:

$$V_1^S(Q_1^1) = Q_1^1\{\beta^S[\pi(Q_1^1 + g^S) - 1] + (1 - \beta^S)\pi(Q_1^1 + g^{\bar{S}})\}$$

The expected value of period 2 profits, after choosing period 1 quantities, but before the price history is known, is:

$$V_2^S(Q_1^1) = E_{P_1}\{V^S(P_1, Q_1^1)\}$$

$$= \int_{-\infty}^{\infty} V^S(P_1, Q_1^1)]f(P_1)dP_1 \tag{22}$$

where $f(P_1) \equiv \phi\left(\frac{P_1 - a_1 + bQ_1^1 + bg^A}{\sigma}\right) \alpha + \phi\left(\frac{P_1 - a_1 + bQ_1^1 + bg^B}{\sigma}\right)(1 - \alpha)$.

Hence, the objective in period 1 is to:

$$\max_{Q_1^1} \; V_1^S(Q_1^1) + \delta V_2^S(Q_1^1) \tag{23}$$

where δ is the probability that the value of the asset will <u>not</u> be publicly revealed between periods 1 and 2.

Suppose that Trader 1 received the A signal. Then the first order condition for Problem (23) is given by:

$$a_1 - b\left(\alpha g^A + (1 - \alpha)g^B\right) - 2bQ_1^1 - \alpha + \delta\frac{\partial V_2^A(Q_1^1)}{\partial Q_1^1} = 0 \tag{24}$$

The derivative in equation (24) may be written as a covariance, as shown in the Appendix, giving:

$$Q_1^1 = \left(\frac{1}{2b}\right)\{a_1 - \left(\alpha g^A + (1-\alpha)g^B\right) - \alpha - \frac{\delta b}{\sigma^2}\operatorname{COV}\left(V^A(P_1, Q_1^1), P_1\right)\} \tag{25}$$

Similarly, if Trader 1 receives the B signal, the first order condition for Problem (23) can be written as:

$$Q_1^1 = \left(\frac{1}{2b}\right)\{a_1 - \left(\alpha g^A + (1-\alpha)g^B\right) - (1 - \alpha) - \frac{\delta b}{\sigma^2}\operatorname{COV}\left(V^B(P_1, Q_1^1), P_1\right)\} \tag{26}$$

The covariance terms in equations (25) and (26) represent the desire on the part of a strategic trader to deceive the other trader. If the covariance terms were zero, then each trader would simply trade on the basis of his information without attempting to conceal his information in order to profit in period 2. (We will later refer to this as "myopia.")

As before, equilibrium requires that the strategic traders play symmetric Nash equilibrium strategies and that the chartists earn zero expected profits. The first requirement means that g^A and g^B solve equations (25) and (26), respectively. Notice, however, that the value function is itself determined by both g^A and g^B, as shown in Section C above. The zero expected profit condition for the chartists states that the expected period 1 price must equal the expected value. In other words,

$$\alpha^2 + (1 - \alpha)^2 = \pi(g^A + g^B) = a_1 - 2b(\alpha g^A + (1 - \alpha)g^B) \qquad (27)$$

Thus, $a_1 = a_0 + 2b(\alpha g^A + (1 - \alpha)g^B)$. (Recall that $a_0 = \alpha^2 + (1 - \alpha)^2$.)

In equilibrium, the best quantity for a trader with an A signal, as given by equation (25), must equal the equilibrium quantity g^A. Similarly, the quantity given by equation (26) must be g^B. Equilibrium is, therefore, the solution to the three simultaneous equations (25), (26), and (27).

V. Equilibrium Belief and Price Reversals

We now turn to qualitative characterizations of the equilibrium calculated above. Our main goal is to describe the equilibrium price paths as a function of the different constellations of private signals. In particular we will focus on the case when the two traders each receive a B signal. The equilibrium price path in this case graphically illustrates the effect of learning with our information structure. In the next section we will show that this case is also capable of generating time series data which are qualitatively similar to those found by the empirical researchers referred to in the introduction.

If both traders receive bad news initially, i.e., both receive B signals, they will tend to revise their beliefs in period 2 and think the asset is worth a lot. Therefore, the expected price, having dropped in period 1, will rise in period 2. Thus, while each trader received bad news initially, trading reveals that, in fact, the news is good. Beliefs are reversed because of learning. Notice that all traders are revising their beliefs concerning the value of the asset in the same direction. Because of the particular information structure in this model, this fad-like behavior is rational.

To demonstrate the possibility of reversal we proceed in two stages. Suppose that both traders receive signal B. In period 1 both traders believe that the asset is most likely worth zero. We first show that in period 1 these traders, who received bad news (signal B) will tend to sell the asset. Conversely, if they had both received good news (signal A) they would have bought the asset. This shows that when both traders receive B signals, the price in period 1 will, on average, be lower than the price in period 0. Recall, however, that the volume of trade will be restricted because of the desire to deceive the other agents.

In period 2 both traders will on average have observed a low price in period 1. Thus, they will each tend to infer that the other trader also got a B signal. This implies that the asset is valuable, so they will both revise their beliefs upwards. The chartists, having seen the low period 1 price, will also revise their beliefs upwards. Then the price in period 2 will tend to rise. So long as the amount of noise in the system is sufficiently low, the average price in period 2 will be higher than the period 0 price. Thus prices reverse.

A. From Period 0 to Period 1

In order for any learning to occur, traders must submit different quantities when they receive different signals. We first show that this must be the case:

Proposition 1: $g^A \neq g^B$.

Proof: The proof is by contradiction. If $g^A = g^B$ in equilibrium, then no learning occurs. In other words, a trader who deviates in period 1 would only affect his period 1 profits. There would be no deception effect on period 2 profits since the other trader would not change his inferences. We will therefore show that if $g^A = g^B$, such a period 1 deviation is profitable.

There are three possible cases.

Case (i): $g^A = g^B = 0$. Since there will be no inference from the period 1 price, for all $\varepsilon < 0$, $V_2^A(\varepsilon) = V_2^A(g^A)$. But, for small $\varepsilon > 0$, $V_1^A(\varepsilon) > 0 = V_1^A(g^A)$. In other words, A is good news so the trader can profit in period 1 by buying the asset.

Case (ii): $g^A = g^B < 0$. Again, since there is no inference from the period 1 price, $V_2^B(0) = V_2^B(g^B)$. But, $V_1^B(0) = 0 > V_1^B(g^B)$. In other words, a trader who receives a B signal can eliminate the loss from buying in period 1.

Case (iii): $g^A = g^B > 0$. This case is analogous to Case (ii). ||

Having established that learning does occur from the period 1 price, we now address the direction of the inferences from the period 1 price. We first show that the strategic traders sell on bad news and buy on good news. Thus, a high price in period 1 is associated with the traders receiving A signals.

Proposition 2: $g^A < g^B$.

Proof: Again the proof is by contradiction. Suppose that $g^A > g^B$, so that the probability that a trader received a B signal is strictly increasing in the period 1 price. In other words, a high period 1 price signals a B.

First, note that each strategic trader wants to conceal his information from the others, because expected profits in the second period are higher if the second period price is more favorable. Therefore, under the hypothesis $g^A > g^B$, a trader with a B signal always prefers a lower period 1 price. And a trader with an A signal prefers a higher period 1 price. Thus, $V_2^A(Q_1^1)$ is decreasing in Q_1^1, and $V_2^B(Q_1^1)$ is increasing in Q_1^1.

Second, observe that $g^A > g^B$ implies that at least one of the following must hold: (i) $g^A > 0$; or (ii) $g^B < 0$.

Case (i): $g^A > 0$. Since $V_2^A(Q_1^1)$ is decreasing in Q_1^1, if the trader deviates from equilibrium by choosing quantity 0 in period 1, instead of g^A, then period 2 value rises: $V_2^A(0) > V_2^A(g^A)$. Now consider the period 1 portion of the objective, $V_1^A(g^A)$. But $g^A > 0$, which implies $V_1^A(g^A) < 0$. Consequently, $V_1^A(0) = 0 > V_1^A(g^A)$.

Case (ii): $g^B < 0$. This is symmetric to Case 1. ||

The intuition underlying Proposition 2 is straightforward. If the traders reacted to good news initially by selling the asset, then each trader could unambiguously benefit by deviating: he could simultaneously increase profits in period 1 and increase the beneficial deception of the other strategic trader and the chartists. Equilibrium requires that the immediate benefits of trading in period 1 be balanced against the subsequent costs of revealing information. This can only happen if $g^A < g^B$. An immediate consequence of Proposition 2 is that if both traders receive B signals, the price in period 1 will fall compared to the period 0 price. If they both receive A signals, then the price will rise.

B. <u>From Period 1 to Period 2</u>

In order to draw inferences from the period 1 price, agents must know how the strategic traders choose quantities depending on the signal received. The above propositions demonstrate that the traders will buy if they receive good news, and sell if they receive bad news. Therefore, agents will learn from period 1 to period 2.

How do strategic traders' beliefs evolve? Since learning does take place after period 1, their beliefs about the other trader's signal and the asset value will, on average, be updated in the right direction.

Proposition 3: (i) If both traders get B signals, then a trader's expected valuation in period 2 is greater than the valuation in period 1.

(ii) If both traders get A signals, then a trader's expected valuation, in period 2 is greater than in period 1.

(iii) If the traders get different signals, then a trader's expected valuation in period 2 is lower than in period 1.

Proof: See Appendix.||

Proposition 3 characterizes how traders learn in response to the period 1 price. In particular, Case (i) of Proposition 2 shows that it is possible that a piece of information can initially be interpreted by the traders as bad news, and subsequently be viewed as good news by the traders. The reversal is not due to noise: in order to make the comparison, we average out the noise in the system.

The chartists will also observe the period 1 price and draw their own inferences about the signals the strategic traders received. The chartists' inference problem, however, is more complicated since the only information they have is the period 1 price. The complication is that the chartists must make inferences about both signals, while each trader need only infer the other trader's signal. The counterpart of Proposition 3 is therefore:

Proposition 4: (i) $E[\gamma_{BB}|BB] > (1 - \alpha)^2$.

(ii) $E[\gamma_{AA}|AA] > \alpha^2$.

(iii) $E[\gamma_{AB}|AB] = E[\gamma_{BA}|BA] > \alpha(1 - \alpha)$.

The proof is similar to that of Proposition 3 and is omitted.

C. From Period 0 to Period 2

The price path corresponding to the event BB is of particular interest for our results. Our main proposition, Proposition 5, explains how a price reversal may occur in this event. This characterization depends on the amount of noise in the asset demand. The reason is that a reversal of the price path will only occur if sufficient learning takes place. Since our model has only two periods after the information arrives, if the demand function is too noisy then there will not be adequate time for this learning.

We first define a benchmark for the amount of noise in terms of 'myopic' behavior. If the traders do not act to conceal their information we say that they are behaving myopically. The myopic equilibrium period 1 quantitites are given by solving equations (25), (26), and (27), ignoring the covariance terms. We refer to this as the 'myopic equilibrium,' but stress that it is not an actual equilibrium of the model.[10] It is a hypothetical device for isolating the effect of strategic deception on the price. The subscript 'M' will indicate myopic equilibrium outcomes.

Considering the event BB, notice that, due to the noise in demand, the strategic traders, while learning, may not revise their valuation very much. In other words, although Proposition 3 shows that

$$E[\beta^B(P_1, g^B)|BB] > 1 - \alpha,$$

the amount of the revision may not be very large.

The chartists' inference problem is more complex. While they learn about the pair of signals received, their ultimate concern is with the value of the asset. There are four possible events in terms of the signals, AA, AB, BA, and BB. These correspond to three possible period 1 quantities traded by the strategic traders, hence three possible expected prices. The actual price is a noisy signal, with variance σ^2, on the expected price. Therefore, the ability of chartists to distinguish between the three possible expected prices will depend on this variance. In the event of two B signals being received, it is possible that while their assessment of the event BB increases, their valuation of the asset falls. The reason for this is that their posterior likelihood of the event AA may fall by more than the rise in their posterior likelihood of the event BB. If the the amount of noise was small, then their expected valuation would be revised in the correct direction. There are similar characterizations of equilibrium price paths for the events AB and AA. These are omitted.

We first consider the myopic equilibrium as σ^2 tends to zero.

Lemma 2: In the limit as σ^2 tends to zero,

$$E_M[\beta^B(P_1,\ g^B)|BB] = 1,$$

$$E_M[\gamma_{BB}(P_1)|BB] + E_M[\gamma_{AA}(P_1)|BB] = E_M[\gamma_{BB}(P_1)|BB] = 1.$$

The proof is immediate.

Proposition 5: Consider the limit of the equilibrium as σ^2 goes to zero, in the event that both traders receive B signals. Then the expected period 1 price is less than the expected period 0 price. For sufficiently small δ, the expected period 2 price is greater than the expected period

0 price.

Proof: Period 1: If both traders receive signal B, they each choose g^B. By Proposition 1, $g^B > \alpha g^A + (1 - \alpha)g^B$. Thus,

$$E[P_1|BB] = a_1 - b(2g^B) < a_1 - b(2)[\alpha g^A + (1 - \alpha)g^B] = E[P_0] = a_0.$$

Period 2: For sufficiently small δ the strategic traders behave myopically. The details are shown in the Appendix. But, in the limit of the myopic equilibrium as σ^2 tends to zero, the chartists value the asset at 1, and the strategic traders are inactive because they have no informational advantage. Therefore, the price converges to 1.||

D. Time Series Properties of the Price

We have described the equilibrium price paths for the three different constellations of signals averaged over the noise in the demand curve. Now consider the equilibrium outcomes given the price in period 1, averaged across the different possible signals. The martingale property of asset prices would imply that expected period 2 price should equal the period 1 price. In our model, however, the information revealed by the period 1 price has quite different implications for the period 2 price. In our model prices do not follow a martingale. The martingale property of prices implies, among other restrictions, that returns are serially uncorrelated. Since in our model prices need not follow a martingale, it follows that returns need not be serially uncorrelated.

Another way of seeing this point is to consider the strategies of the technical analysts conditional on observing the period 1 price. The chartists trade on the basis of the following rule: buy on a price change; sell if there is little price movement. The reason this is their best strategy is that a high period 1 price implies that there was heavy demand from the informed traders. Hence, they both received A signals and the asset is worth a lot. Conversely, a low period 1 price is associated with the informed traders receving B signals which also implies that the asset is valuable. Finally, if the price does not move much in period 1, then it is likely that the informed traders received different signals. Thus, for example, a low price in period 1 implies a high price in period 2 which is a clear violation of the martingale property of asset prices.

VII. Concluding Remarks

We conclude by discussing two points. The first concerns limit orders, and the second concerns trading volume.

In our model the strategic traders and the chartists are restricted to using market orders. If they were able to submit 'limit orders' in which the trade in a given period is conditioned on the price being formed in that period, then the model would become more complex, and the martingale property might reappear. In fact, the limit orders which are allowed in some financial markets are not exactly analogous to such strategies. If the asset price jumps or moves quickly, outstanding limit orders will be cleared at disadvantageous prices as the orders are crossed to get to the new market price. Thus, submitting a limit order, rather than a market order, is not a dominant strategy. The integration of traders' optimal choice between limit orders and market orders into a similar trading model is the subject of our current research. Furthermore, we stress the Dubey, Geanakoplos, and Shubik (1987) result that a demand schedule submission game does not uniquely implement rational expectations equilibrium. In other words, it is not immediately obvious that the lack of the martingale property is due to the restriction of traders' strategies to market orders.

In this model we ignored budget constraints and short sale restrictions. Therefore, we could not jointly determine equilibrium price-volume paths. This seems to be the most promising direction for testable implications of the model, and is the subject of further research.

Appendix

Derivations of equations (25) and (26)

In equation (24) of the main text note that:

$$\frac{\partial V_2^S(Q_1^1)}{\partial Q_1^1} = \int_{-\infty}^{\infty} \{V^S(P_1,Q_1^1) \frac{\partial f(P_1)}{\partial Q_1^1} + [\frac{\partial V^S(P_1,Q_1^1)}{\partial Q_1^1} + \frac{\partial V^S(P_1,Q_1^1)}{\partial P_1} \frac{\partial P_1}{\partial Q_1^1}]f(P_1)\}dP_1 \quad (A1)$$

where:

$$\frac{\partial f(P_1)}{\partial Q_1^1} = \alpha\phi(\frac{P_1 - a + bQ_1^1 + bg^A}{\sigma})(\frac{P_1 - a + bQ_1^1 + bg^A}{\sigma})(\frac{b}{\sigma})(-1) +$$

$$(1 - \alpha)\phi(\frac{P_1 - a + bQ_1^1 + bg^B}{\sigma})(\frac{P_1 - a + bQ_1^1 + bg^B}{\sigma})(\frac{b}{\sigma})(-1)$$

Define $\varepsilon = P_1 - a_1 + bQ_1^1 + bQ_1^2$. In other words, ε is the deviation from the period 1 expected price. Then:

$$\frac{\partial f(P_1)}{\partial Q_1^1} = - \frac{b\varepsilon}{\sigma^2} f(P_1).$$

Note also that by Lemma 1:

$$V^S(P_1, Q_1^1) = V^S(P_1 + bQ_1^1, 0)$$

And therefore:

$$\frac{\partial V^S(P_1, Q_1^1)}{\partial Q_1^1} - \frac{\partial V^S(P_1, Q_1^1)}{\partial P_1} b = 0.$$

Thus, equation (A1) becomes:

$$-\frac{b}{\sigma^2}\int_{-\infty}^{\infty} V^A(P_1, Q_1^1)\varepsilon f(P_1)dP_1 = -\frac{b}{\sigma^2}\,\mathrm{COV}(V^A(P_1, Q_1^1), P_1)$$

__Proof of Proposition 3:__

Case (i): Consider Trader 1's valuation, based on the period 1 price, P_1. Define ε to be the deviation of this price from the expected price:

$$\varepsilon = P_1 - a_1 + 2bg^B.$$

Trader 1's belief, $\beta(P_1, g^B, B)$ that the asset is worth a lot, is a function of two independent random variables, namely, Trader 2's signal, and the realization of the net liquidity trader noise, ε. Consider $E[\beta(P_1, g^B, B)|\varepsilon, B]$, the expectation of Trader 1's belief, conditional on both the realization of ε, and the signal received by Trader 2. Since these two variables completely determine the belief, this is simply the expectation of a degenerate random variable. By Proposition 2,

$$E[\beta(P_1, g^B, B)|\varepsilon, B] > E[\beta(P_1, g^B, B)|\varepsilon]$$

since ε is independent of Trader 2's signal. Taking expectations over ε, the result follows immediately:

$$E[\beta(P_1, g^B, B)|\, B] > \beta(B).$$

The other cases are similar.||

Proof of Proposition 5: Convergence to Myopic Equilibrium Strategies

Substituting the market efficiency condition from equation (27) into the other equilibrium equations, (25) and (26), gives:

$$2bg^B = a_0 + b[\alpha g^A + (1 - \alpha)g^B] - (1 - \alpha) -(\frac{\delta b}{\sigma^2})COV^B$$

$$2bg^A = a_0 + b[\alpha g^A + (1 - \alpha)g^B] - \alpha - (\frac{\delta b}{\sigma^2})COV^A$$

These imply that:

$$2bg^B = \alpha(2\alpha - 1) - (\frac{\delta b}{\sigma^2})[(2 - \alpha)COV^B + \alpha COV^A] \tag{A2}$$

Define ρ^A to be the correlation coefficient between the expected profits in the case of an A signal with the period 1 price, and σ^A to be the standard deviation of expected second period profits in case of an A signal. Analogously define ρ^B and σ^B. Equation (A2) can therefore be written as:

$$2bg^B = \frac{\alpha(2\alpha - 1)}{2b} - (\frac{\delta}{2\sigma^2})[(2 - \alpha)\rho^B\sigma^B + \sigma\rho^A\sigma^A]$$

Since $-1 < \rho^A < 0$ and $0 < \rho^B < 1$, by Proposition 1, the term in square brackets will be at most $(2 - \alpha)\sigma^B$, and at least $(-\alpha)\sigma^A$. Therefore:

$$(\frac{\alpha}{2b})(2\alpha - 1) - (\frac{\delta}{2\sigma})(2 - \alpha)\sigma^B(\delta) < g^B(\delta) < (\frac{\alpha}{2b})(2\alpha - 1) + (\frac{\delta}{2\sigma})\alpha\sigma^A(\delta)$$

where we have emphasized the dependence of g^B, σ^A and σ^B on δ.

Expected profits are bounded by 0 and 1/4b. The latter is the amount of money a trader would expect to make if he knew the asset was worth 1, the chartists thought it was worth nothing, and the other trader was inactive. Consequently, σ^A and σ^B, although they depend on δ, are bounded independently of δ. Therefore:

$$\lim_{\delta \to 0} g^B(\delta) = (\frac{\alpha}{2b})(2\alpha - 1) = g_M^B.$$

Proceeding in a similar fashion:

$$\lim_{\delta \to 0} g^A(\delta) = \left(\frac{1}{2b}\right)(2\alpha - 1)(\alpha - 1) = g_M^A.$$

In other words, for sufficiently small δ the strategic traders behave myopically. But, in the limit of the myopic equilibrium as σ^2 tends to zero, the chartists value the asset at 1, and the strategic traders are inactive because they have no informational advantage. Therefore, the price converges to 1.||

Footnotes

[1]Limit orders are intermediate strategies between fixed quantities and entire demand functions. Introducing limit orders complicates the strategies because agents choose two parameters instead of one. Allowing limit orders in the model would seem to be valuable, and is on our research agenda, but they are not considered in this paper. Limit orders are discussed further in the concluding section of the paper.

[2]The model here is different from Kyle (1985) in a number of other important respects, including the absence of a market maker, which will be made clear below.

[3]Notice that this assumes that the private knowledge or expertise cannot itself have been revealed already through trading, possibly in some other asset market. This is plausible when the expertise is sufficiently complex.

[4]Kyle (1985) originated the notion of liquidity traders. Unlike Kyle, however, our liquidity traders are not willing to pay any price to satisfy their preference for liquidity.

[5]Because the demand curve is linear, the random noise in liquidity trader volume could have been added to the quantity rather than to the price. In other words, the random shock to quantity is simply ε/b.

[6]An alternative interpretation of the chartists' effect on the demand curve is that the liqudity traders, understanding the model, adjust their demands on the basis of the publicly available information. For ease of exposition we prefer the interpretation given in the text of the paper which treats them as two seperate groups.

[7]Since it is common knowledge that in period 0 there is no private information, it seems reasonable to assume that the market efficiency condition holds without chartists trading. In other words, the informed traders and the chartists can only make 'super normal' profits when they have an informational advantage.

[8]If the demand function is not linear, then some parts of the demand function are more informative than others. In other words, a trader choosing his period 1 quantity has an additional factor to take into account, namely, the amount of information that will be revealed by the period 1 price. Including this 'confusion effect' makes the problem intractable and hence we have assumed a straight line demand function.

[9]The alternative is to solve all four equations simultaneously. This would be notationally cumbersome, but would naturally give the same solution. The special structure of the problem which allows us to proceed as we do is the fact that, because of linear demand, the influence of a trader's period 1 quantity on his inferences is straightforward as shown by Lemma 1.

[10]Note that in the myopic equilibrium the traders are still behaving as Cournot-Nash duopolists in period 1, although they ignore the effects of their actions on the period 2 outcome.

References

Akerlof, George A. (1970), "The Market for 'Lemons': Quality Uncertainty and the Market Mechanism," Quarterly Journal of Economics (August), 488-500.

Aumann, Robert (1976), "Agreeing to Disagree," Annals of Statistics 4, 1236-39.

Barclay, Michael, Robert Litzenberger, and Jerold Warner (1988), "Private Information, Trading Volume, and Stock Return Variances," Review of Financial Studies, forthcoming.

Bamber, Linda Smith (1986), "The Information Content of Annual Earnings Releases: A Trading Volume Approach," Journal of Accounting Research 24(1), 40-56.

Barwise, J. (1981), "Scenes and Other Situations," Journal of Philosophy LXXVIII, 369-97.

Dubey, Pradeep, John Geanakopolos, and Martin Shubik (1987), "The Revelation of Information in Strategic Market Games," Journal of Mathematical Economics 16, 105-37.

French, Ken and Richard Roll (1986), "Stock Return Variances: The Arrival of Information and the Reaction of Traders," Journal of Financial Economics 17, 5-26.

Geanakoplos, John (1988), "Common Knowledge, Bayesian Learning, and Market Speculation With Bounded Rationality," Yale University, working paper.

__________ and H. Polemarchakis (1982), "We Can't Disagree Forever," Journal of Economic Theory 28, 192-200.

Grossman, Sanford and Joseph Stiglitz (1980), "On the Impossibility of Informationally Efficient Markets," American Economic Review 70, 393-408.

____________________ (1976), "Information and Competitive Price Systems," American Economic Review 66, 246-53.

Keynes, John M. (1964), The General Theory of Employment, Interest, and Money (New York: Harcourt, Brace, and World; reprint of 1936 original).

Kyle, Albert (1985), "Continuous Auctions and Insider Trading," Econometrica 53, 1315-35.

Lehmann, Bruce (1988), "Fads, Martingales, and Market Efficiency," Columbia University, working paper.

Milgrom, Paul and Nancy Stokey (1982), "Information, Trade and Common Knowledge," Journal of Economic Theory 26, 17-27.

Meese, Richard (1986), "Testing for Bubbles in Exchange Markets: A Case for Sparkling Rates?," Journal of Political Economy 94(2), 345-73.

Moses, Yoram, Danny Dolev, and Joseph Halpern (1986), "Cheating Husbands and Other Stories: A Case Study of Knowledge, Action, and Communication," Distributed Computing 1, 167-76.

Rothschild, Michael and Joseph Stiglitz (1976), "Equilibrium in Competitive Insurance Markets: An Essay on the Economics of Imperfect Information," Quarterly Journal of Economics 90, 629-49.

Sebenius, J. and J. Geanakoplos (1983), "Don't Bet on It: Contingent Agreements With Asymmetric Information," Journal of the American Statistical Association 78, 424-26.

Tirole, Jean (1982), "On the Possibility of Speculation Under Rational Expectations," Econometrica 50, 1163-81.