NBER WORKING PAPERS SERIES

TRADING, COMMUNICATION AND THE RESPONSE OF PRICE TO NEW INFORMATION

James Dow

Gary Gorton

Working Paper No. 3687

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
April 1991

We are grateful to Andy Abel, Franklin Allen, Fischer Black, Colin Camerer, Julian Franks, Douglas Gale, Meg Meyer, John Moore, Anthony Neuberger, Sam Orez, Krishna Ramaswamy, and Ailsa Roell, for helpful comments. The second author thanks the NSF (#SES-8618130) and the Geewax-Terker Research Program in Financial Instruments for financial support. A previous version of this paper was entitled "Self-Generating Trade and Rational Fads: The Response of Price to New Information." This paper is part of NBER's research program in Financial Markets and Monetary Economics. Any opinions expressed are those of the authors and not those of the National Bureau of Economic Research.

NBER Working Paper #3687
April 1991

TRADING, COMMUNICATION AND THE RESPONSE OF PRICE TO NEW INFORMATION

ABSTRACT

The dynamic behavior of security prices is studied in a setting where two agents trade strategically and learn over time from market prices. The model introduces an information structure which is intended to capture the notion that information is difficult to interpret. Strategic interaction and the complexity of the information result in a protracted price response. Indeed, equilibrium price paths of the model may display reversals in which the two 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.

James Dow
Institute of Finance
London Business School
Sussex Place, Regent's Park
London NW1 4SA
United Kingdom

Gary Gorton
Department of Finance
The Wharton School
University of
Pennsylvania
Philadelphia, PA 19104

I Introduction

In this paper we explore the link between price change and the arrival of news about underlying fundamental value, under the assumption that the adjustment is not instantaneous. Keynes' beauty contest, and the commentary on the Crash of 1987, illustrate a general disquiet in the economics profession that stock price movements appear to bear little relationship to changes in fundamentals. Our model is designed to emphasize that the price change following the arrival of news need not be a quick, smooth, response. Indeed, the pattern of the price response over time may be so complicated that there is no apparent relationship between the arrival of new information and the price. One implication of this is that a sharp price movement does not necessarily reflect the contemporaneous arrival of significant new information about the fundamental.

The assumption that the incorporation of new information into the price is not instantaneous is suggested by recent empirical work. For example, Oldfield and Rogalski (1980), French and Roll (1986), and Barclay, Litzenberger, and Warner (1990) find that stock return variances are higher during trading hours than during nontrading hours. Barclay, Litzenberger, and Warner (1990) also find an increase in volume associated with increased trading hours. Meese (1986) found that the return variance during trading hours for a sample of daily dollar-yen exchange rates is three times higher than during nontrading hours. It appears that the trading process itself is the mechanism for information exchange.

We allow for learning via trading and consider the response of price over time when there are multiple informed traders and an information structure which is intended to capture the notion that information is difficult to interpret. The potentially complex price response is caused by the complexity of the information that the traders receive. In order to illustrate this, we give an example of an information structure with the property that news may arrive which causes the price to fall initially and then rise. This reversal of the price path corresponds to traders initially believing that the news is bad and then revising their valuations in light of the other traders' reactions. Information is often modelled with some special structure–for example, each agent receiving a noisy signal on the truth–which precludes this type of reversal. However, in general, an agent's inference, based on his own signal, may be very different from the inference based on the combined signals of all agents. Information may be difficult to interpret, and the significance of some information may not be understood without other information. For example, suppose the agent hears that a company's president is about to move to a different firm, or decides that the industry will soon be affected by a new technology, or believes that trade policy will be affected by current international political developments. It is not clear whether these pieces of information are good or bad for the company's stock. Unlike the traditional 'truth plus noise' paradigm, the significance of this information will only be known when combined with many other pieces of information, some of which may already be known by other agents, while others may become known in the future. While the price reversal may be an extreme example of a complex price response to new information, the general point is that the immediate reaction may be significantly different from the eventual adjustment. This result holds, unless adjustment to news is instantaneous, whenever the implication drawn from individual signals is much weaker than the implication drawn from pooling the signals.

Market efficiency implies zero serial correlation in asset prices. Otherwise, there is an obvious and simple arbitrage strategy: buy low, sell high. We do not allow this type of market inefficiency. Thus, the reversal does not take place *on average*, but only conditional on certain realizations of the news, and corresponding eventual asset prices. While sometimes good news will initially be interpreted as bad news (negative serial correlation), other news will be associated with positive serial correlation (in our model, a small initial positive response is followed by a larger positive movement). To emphasize the point that serial correlation may be nonzero, conditional on the true information and final price, we focus on the price reversal.

In our model traders each receive private signals. 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. Even a single trader receiving a private signal will spread trade over all periods much like the monopoly owner of a durable resource choosing an intertemporal pricing profile.

With multiple informed traders there are two additional effects. First, each trader, knowing that the other informed traders are acting on their information, will make inferences about the value of the asset from past market prices and combine these inferences with his own signal. Secondly, this learning process creates strategic possibilities. Since other traders are learning from prices, each trader will want to conceal his own information, and further, will actually want to mislead his rivals into drawing the wrong inferences. This effect conflicts with the trader's desire to profit immediately by trading on his information and, in equilibrium, he optimally balances these two effects.

The main point about multiple traders concerns the interpretation of information over time. Traders receive different signals and in time will infer each others' signals from the price. Because information is difficult to interpret, equilibrium price paths 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, by everybody, may be revealed, through trading to be good news.

These apparently unusual price paths may occur when the information structure differs from the standard 'truth plus noise' paradigm in which each agent receives a private signal equal to the value of the asset plus an error term. In such a setting, the exchange of information leads to averaging the individual valuations. The 'truth plus noise' structure leads to a price response in which, on average, the initial reaction is followed by successively smaller responses in the same direction. In a more general setting this is not necessarily true, as our information structure shows. Consider the following illustration.

Cray Research is a leading manufacturer of supercomputers which had remained closely identified with its founder Seymour Cray. In 1989 Mr. Cray was engaged in a large research project to develop a new computer technology based partly on the use of gallium arsenide chips. This was an ambitious project which, if successful, was likely to lead to a new generation of supercomputers, giving the company a strong competitive advantage. But it also seemed risky. If it failed it would not only have lost the company a large amount, but would have diverted resources away from development of the existing technologies. Because the project was radically different, failure would have no beneficial consequences for the existing technology. In the Spring of 1989 the company had not decided whether to continue research on the project, and Mr. Cray was considering leaving to found a separate company.[1]

This situation can be adapted to provide an illustration of our information structure. Consider two analysts, one of whom learns whether Mr. Cray is leaving or not, while the other has superior information about whether the new technology will succeed. If Mr. Cray stays and the technology is successful, the shares will have a high value since the company will have the most advanced technology. If he stays and the project fails, the shares will have a low value since the company will have wasted resources and neglected the development of its existing technology. If Mr. Cray leaves and the technology proves 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 manufacturer.

Suppose that analysts' prior beliefs are that Mr. Cray will stay and that the project will succeed. What happens if one analyst learns that Mr. Cray is likely to leave, and the other analyst studies the technology and decides that it is more likley to fail than is generally

[1] *Business Week*, April 30, 1990, describes what happened to Seymour Cray.

recognized in the market? After receiving his private signal, each analyst believes the shares are of low value. If each sells on his information, the price will fall. Realizing that the low price reflects sales by the other trader, each infers the signal received by the other trader and revises his valuation upwards. Then buying causes the price to rise again. There is a price reversal as the information is revealed through trading.

Since Grossman and Stiglitz (1976, 1980) a great deal of attention has been focused on the details of the price formation process underlying rational expectations models. For example, Kyle (1985) and Glosten and Milgrom (1985) set out models of the relationship between a market making institution and an informed trader. It is clear that if agents learn instantaneously, or if the market making institution does all the learning, then there can be no exchange of information via trading in real time. Since our focus is not on the price formation process, but rather on the effect on the price path of information exchange between multiple informed traders, we adopt a simple exogenous price formation process. The main requirement for our conclusion is a price formation institution which raises the price when people buy and lowers the price when people sell. Modelling price formation institutions has turned out to be a difficult theoretical problem. There are a number of price formation models which are less ad hoc than the one we use here, but they still have defects, and they are sufficiently complicated that they would be intractable in the setting here.

The paper proceeds as follows. Section II introduces the model. The equilibrium is derived in Section III. Section IV describes the evolution of prices and beliefs in equilibrium. For illustrative purposes the focus is on the instance of a price reversal. Section V discusses the robustness of the results and concludes.

II The Model

We consider a single asset market characterized as a downward sloping noisy linear demand curve, described below. Two strategic traders interact with each other and with this demand curve. 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 period, the model is analogous to standard Cournot duopoly where sellers of a good have private information about the common, unknown, production cost. However, in our model the demand curve will adjust to reflect past information as we describe below.

Within each trading period 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 agents 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.

In each trading period, agents' orders are aggregated and executed at the market clearing price. Nash equilibrium of the model requires that each trader choose a strategy that maximizes his payoff, given the the strategies of the other agents in the model. Below, we characterize the properties of the equilibrium expected price path for given realizations of the private signals.

A Strategic Traders

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 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. The expected value of the asset conditional on receiving an A signal is α, but its value will actually be either 1 in case the other agent received an A signal, or 0 if the other agent received a B signal. A is good news because the valuation given A, α, exceeds the unconditional valuation, $\alpha^2 + (1-\alpha)^2$. Conversely, the expected value of the asset conditional on receiving a B signal is $(1-\alpha)$. In other words, a B signal is "bad news" before the other trader's signal is known.

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 the asset value is *not* publicly revealed by period 2. The characterization of the equilibrium in period 2, given below, will be conditional on the information not being revealed. Otherwise, the period 2 equilibrium would be analogous to the period 0 equilibrium in which the strategic traders would not trade.

We will use the following notation. Q_t^j will denote the quantity sold by trader j in period t. 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. We describe the symmetric Nash equilibrium of the game.

[2]

Trader 1's strategy is a rule for choosing Q_t^1. We represent the strategy as 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

[2] We assume there are no restrictions on borrowing and short sales.

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. Therefore, 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.

B The Trading Institution

The strategic traders face a noisy downward sloping demand curve. The expected inverse (residual) demand curve is given by:

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

where Q is the total (net) quantity sold by the strategic traders. The slope, -b, is a constant while the intercept, a, changes as a function of the past history, as we describe in subsection C. The realized price, P, will differ from the expected price due to fluctuations in the noise term: $P = \pi(Q) + \epsilon$ where ϵ is normally distributed with zero mean and variance σ^2. Subscripts will denote the different periods: P_0 is the period 0 price, etc.

The interpretation of equation (1) is that the downward sloping linear demand curve represents the preferences of liquidity traders. Suppose that 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 an undervalued financial asset in order to invest in a superior real alternative. In other words, these private projects require owner financing (daughter's wedding, childrens' college, family business, etc.) and the owners are liquidity constrained. They are willing to pay a premium to satisfy their 'liquidity needs,' but they are not willing to pay any price.[3]

[3]In the context of a trading system based on specialists, the model corresponds to a situation in which

We also assume that the downward sloping demand curve is noisy. We assume this noise is infinitesimal, i.e., we consider limiting equilibria as the variance, σ^2, tends to zero.[4]

The assumption that the demand curve is noisy insures the existence of equilibrium in pure strategies in our model and may be interpreted randomness in exogenous liquidity needs that affect the whole population.[5]

C Market Efficiency

Above we interpreted the downward slope of the residual demand curve as being due to liquidity traders. Note that the intercept represents the expected price of the asset if the strategic traders did not submit orders. We assume a market efficiency condition which requires that the intercept adjust so that the expected price equals the expected value of the asset, conditional on available public information.

The rationale for market efficiency is as follows. If informed strategic traders can learn from past prices, so can outside observers. Such technical analysts, or chartists, would be prepared to enter the market and trade on the basis of their inferences. If there is free entry of such technical analysts into the model, they will make zero profits in equilibrium.[6] For ease of exposition we describe the market efficiency condition in terms of inferences made by such technical analysts.

Nash equilibrium for the technical analysts is equivalent to a market efficiency condition.

there are an infinite number of liquidity traders that submit limit orders every period with the rule that all orders are cleared at the marginal price.

[4]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.

[5]Noise in the price creates an arbitrage opportunity since it causes the price to deviate from the fundamental, even in the absence of private information, since it causes negative serial correlation. In our model no agents have strategies spaces which would allow them to make such an arbitrage. With infinitesimally small noise, however, transactions costs make this arbitrage infeasible. We are primarily interested in the limiting equilibria as the noise goes to zero which provides some justification for this assumption.

[6]An alternative interpretation of the chartists' effect on the demand curve is that the liquidity 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 separate groups.

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.

D Summary

In the model a trader who receives private information has an incentive to spread his trades over both remaining periods because the demand curve slopes down. He will trade some amount in the first period and, in doing so, will provide information about the signal received to the other strategic trader. Symmetrically, he will be able to make inferences from the period 1 price about the signal received by the other trader. There is, therefore, an additional incentive for traders to restrict their period 1 quantities, namely, a desire to deceive other traders. Chartists, observing the period 1 price, update their valuation of the asset. In period 2, the strategic traders again trade, facing a demand curve with an intercept representing the new inferences of the chartists.

III The Equilibrium

To introduce the basic structure of the model we start by considering equilibrium in period 0. Because information has not yet arrived this period is essentially separate from periods 1 and 2 and will serve as a benchmark for describing the equilibrium in later periods. In subsequent periods agents will learn from the price and so we describe how they learn. Equilibrium in periods 1 and 2 is determined recursively.

A Equilibrium in Period 0

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

$$[\pi(Q_0^1 + Q_0^2) - [\alpha^2 + (1-\alpha)^2]]Q_0^1 \,. \tag{2}$$

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:

$$\begin{aligned} \alpha^2 + (1-\alpha)^2 &= \pi(Q_0^1 + Q_0^2) + Q_0^1 \pi'(Q_0^1 + Q_0^2) \\ &= a_0 - b(Q_0^1 + Q_0^2) - Q_0^1 b \,. \end{aligned} \tag{3}$$

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

$$g = \frac{[a_0 - \alpha^2 - (1-\alpha)^2]}{3b} \,. \tag{4}$$

The market efficiency condition requires that, in period 0, $a_0 = \alpha^2 + (1-\alpha)^2$. The average 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.

B Learning from the Price

In equilibrium, agents learn from the period 1 price. Because the traders' period 1 quantities depend on the signals they received, the period 1 price will reveal information about these signals. This information affects the period 2 equilibrium in two ways. First, each strategic trader uses this information to update his beliefs about the other's signal, 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 takes this into account in inferring the quantity sold by the other trader. Second, chartists use the price information to update their beliefs. Their inference is less precise. They receive no signal of their own, so they infer 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 simultaneously infer both of the quantities.

Prior to trade in period 1 each trader's valuation depends only on his signal. Since the asset is worth 1 if both receive the same signal, and is worthless otherwise, the expected value is the probability that the other trader 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. So $S = A$ or B and $\beta^A = \alpha$ and $\beta^B = 1 - \alpha$. We also define $\overline{S}$ by $\overline{S} = A$ when $S = B$ and $\overline{S} = B$ when $S = A$.

After period 1 each trader revises his belief. 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:

$$\begin{aligned}\beta_S(P_1, Q_1^j) &= \text{Prob}(S|P_1, Q_1^j) \\ &= \frac{\lambda(P_1|Q_1^j, g^S)\beta^S}{\lambda(P_1|Q_1^j)}\end{aligned}$$

where the $\lambda(P_1|x)$ is 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:

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

we have:

$$\begin{aligned}\beta^A(P_1, Q_1^j) = \text{Prob}(A|P_1, Q_1^j) &= \frac{\lambda(P_1|Q_1^j, g^A)\alpha}{\lambda(P_1|Q_1^j, g^A)\alpha + \lambda(P_1|Q_1^j, g^B)(1-\alpha)} \\ &= \frac{\phi\left(\frac{P_1 - \pi(g^A + Q_1^j)}{\sigma}\right)\alpha}{\phi\left(\frac{P_1 - \pi(g^A + Q_1^j)}{\sigma}\right)\alpha + \phi\left(\frac{P_1 - \pi(g^B + Q_1^j)}{\sigma}\right)(1-\alpha)} \qquad (5)\end{aligned}$$

where ϕ is 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 realization of the random shock. The correction is straightforward, as we now show:[7]

Lemma 1

$$\beta^S(P_1, Q_1^j) = \beta^S(P_1 + bQ_1^j, 0)\ . \tag{6}$$

PROOF: Since π is linear with slope $-b$:

$$\begin{aligned}\lambda(P_1 + bQ_1^j|0, g^S) &= \phi\left(\frac{P_1 + bQ_1^j - \pi(g^S)}{\sigma}\right) \\ &= \phi\left(\frac{P_1 + \pi(Q_1^j + g^S)}{\sigma}\right) = \lambda(P_1|Q_1^j, g^S)\end{aligned}$$

||

The market efficiency condition requires that chartists make zero profits in period 2, conditional on their inferences from the period 1 prices. 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. Let $\gamma_{AA}(P_1)$ be the posterior probability of the event that the two traders both received signal A. This is the same as the probability that the quantity sold was $2g^A$. By Bayes' Rule,

$$\begin{aligned}\gamma_{AA}(P_1) &= \frac{\lambda(P_1|2g^A)\alpha^2}{\lambda(P_1|g^A, g^A)\alpha^2 + \lambda(P_1|g^B, g^B)(1-\alpha)^2 + \lambda(P_1|g^A, g^B)2\alpha(1-\alpha)} \qquad (7) \\ &= \frac{\phi\left(\frac{P_1 - \pi(2g^A)}{\sigma}\right)\alpha^2}{\phi\left(\frac{P_1 - \pi(2g^A)}{\sigma}\right)\alpha^2 + \phi\left(\frac{P_1 - \pi(2g^B)}{\sigma}\right)(1-\alpha)^2 + \phi\left(\frac{P_1 - \pi(g^A + g^B)}{\sigma}\right)2\alpha(1-\alpha)}\ .\end{aligned}$$

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

Let the value of the asset be $\gamma(P_1) = \gamma_{AA}(P_1) + \gamma_{BB}(P_1)$. The chartists will also predict the quantity of the asset which will be sold in period 2 by the strategic traders:

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

[7] If the demand function is not linear, some parts of the demand function are more informative than others. A trader can design experiments with different informativeness by choosing different quantitites. Thus, in choosing his period 1 quantity he has an additional factor to take into account, namely, the amount of information that will be revealed by the period 1 price. Including this effect makes the problem intractable and hence we have assumed a straight line demand function.

C Equilibrium in Period 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. We proceed recursively by first fixing arbitrary period 1 strategies, and deriving period 2 *equilibrium* strategies. The period 1 equilibrium strategies are determined given the dependence of the period 2 equilibrium on period 1. This is a standard backward induction procedure for computing the equilibrium.

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 $\overline{S}$ then $h^{\overline{S}}(P_1, g^{\overline{S}})$ is the quantity. Assuming that Trader 2 plays his equilibrium strategy, $\{h^A, h^B\}$, in period 2, this predicted quantity is:

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

Trader 1's period 2 objective is:

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

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

$$\begin{aligned}\beta^S(P_1, Q_1^1) = \beta^S(P_1, Q_1^1)[\pi(Q_2^1 + h^S(P_1, g^S)) + Q_2^1\pi'(Q_2^1 + h^S(P_1, g^S))] +\\ (1 - \beta^S(P_1, Q_1^1))[\pi(Q_2^1 + h^{\overline{S}}(P_1, g^{\overline{S}})) + Q_2^1\pi'(Q_2^1 + h^{\overline{S}}(P_1, g^{\overline{S}}))] .\end{aligned}$$

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:

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

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:

$$h^A(P_1, Q_1^1) =$$
$$\frac{a_2(P_1)}{2b} - \beta^A(P_1, Q_1^1)\left[\frac{h^A(P_1, g^A)}{2} + \frac{1}{2b}\right] - (1 - \beta^A(P_1, Q_1^1))\left[\frac{h^B(P_1, g^B)}{2}\right] \tag{11}$$
$$h^B(P_1, Q_1^1) =$$
$$\frac{a_2(P_1)}{2b} - (1 - \beta^A(P_1, Q_1^1))\left[\frac{h^B(P_1, g^B)}{2} + \frac{1}{2b}\right] - \beta^A(P_1, Q_1^1))\left[\frac{h^A(P_1, g^A)}{2}\right] . \tag{12}$$

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. The two traders may be pursuing different period 1 strategies because these are not necessarily equilibrium strategies. There are, therefore, 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.[8] These four equations are replaced by two, which depend on g^A and g^B:

$$h^A(P_1, g^A) =$$
$$\frac{a_2(P_1)}{2b} - \beta^A(P_1, g^A)\left[\frac{h^A(P_1, g^A)}{2} + \frac{1}{2b}\right] - (1 - \beta^A(P_1, g^A))\left[\frac{h^B(P_1, g^B)}{2}\right] \tag{13}$$
$$h^B(P_1, g^B) =$$
$$\frac{a_2(P_1)}{2b} - (1 - \beta^A(P_1, g^B))\left[\frac{h^B(P_1, g^B)}{2} + \frac{1}{2b}\right] - \beta^A(P_1, g^B))\left[\frac{h^A(P_1, g^A)}{2}\right] . \tag{14}$$

These equations are solved simultaneously to obtain period 2 equilibrium strategies:

$$h^A(P_1, g^A) =$$

[8]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.

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

$$h^B(P_1, g^B) =$$

$$\left[\frac{2a_2(P_1) + \beta^A(P_1, g^B)\beta^A(P_1, g^B) + \beta^A(P_1, g^A) + [\beta^A(P_1, g^A)]^2}{3b[2 + \beta^A(P_1, g^A) - \beta^A(P_1, g^B)]}\right] - \frac{1}{3b}\,. \tag{16}$$

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. The linear demand curve implies that the only effect of changing quantity in period 1 is to shift the price by a constant amount. It follows 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.

Equilibrium includes market efficiency which is imposed by requiring that the chartists earn zero expected profits. The chartists' expectation of the strategic traders' volume, conditional on P_1, must be such that the expected price equals the expected value of the asset:

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

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, the quantity he chose in period 1, and Trader 2's period 1 strategy. This value is given by:

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

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

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 $Q_1^1 = g^A$. (18) becomes:

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

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, (18) becomes:

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

It follows from Lemma 1 that $V^A(P_1, Q_1^1) = V^A(P_1 + b(Q_1^1 - g^A), g^A)$ and so:[9]

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

D Equilibrium in Period 1

In period 1 each trader's objective is to maximize the sum of expected period 1 profits, and 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 can 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:

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

[9]To interpret (21) consider a standard model of Cournot duopoly with linear inverse demand with intercept a and slope $-b$, and constant marginal cost, c. 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 an extension of the standard model to our framework.

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

$$\begin{aligned} V_2^S(Q_1^1) &= E_{P_1}\{V^S(P_1, Q_1^1)\} \\ &= \int_{-\infty}^{\infty} V^S(P_1, Q_1^1) f(P_1) dP_1 \end{aligned} \tag{22}$$

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:

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

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

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

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

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

Similarly, if Trader 1 receives the B signal, the first order condition for is:

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

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, 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 refer to this benchmark case below. Notice that the value functions in the covariance terms are, themselves, determined by both g_A and g_B, as shown in Section C above.

The market efficiency condition states that the expected period 1 price must equal the expected value:

$$\alpha^2 + (1-\alpha)^2 = \pi(2(\alpha g^A + (1-\alpha)g^B)) = a_1 - 2b(\alpha g^A + (1-\alpha)g^B) . \tag{27}$$

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 the solution to the three simultaneous equations (25), (26), and (27).

IV 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. We 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 because the period 1 price will be low, but will be followed by a high period 2 price.

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 they will sell the asset. Conversely, if they both receive good news (signal A) they would buy the asset. This implies that when both traders receive B signals, the price in period 1 will, on average, be lower than the price in period 0.

In period 2 both traders will probably have observed a low price in period 1. They will each 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, there will be a reversal: the average price in period 2 will be higher than the period 0 price.[10]

[10] A weaker type of reversal would occur if the expected period 2 price were above the expected period 1 price, but still below the period 0 price. This could happen if the noise in the system was too large, so that the market price does not communicate much information. By considering the effects of reduction in the noise we can focus on cases where the period 2 expected price does reflect the correct valuation relative to the period 0 price.

A From Period 0 to Period 1

In order for any learning to occur, traders must submit different quantities when they receive different signals.

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

PROOF: The proof is by contradiction. If $g^A = g^B$ in equilibrium, then no learning occurs, and a trader who deviated in period 1 would only affects his period 1 profits. There would be no deception effect on period 2 profits since the other trader would not change his inferences. We 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 is no inference from the period 1 price, for all $\epsilon < 0$, $V_2^A(\epsilon) = V_2^A(g^A)$. But, for small $\epsilon > 0$, $V_1^A(\epsilon) > 0 = V_1^A(g^A)$. 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). ||

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. 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 prefers a lower period 1 price, and 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 react 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, the price will rise.

B 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 each trader's expected valuation in period 2 is greater than the expected valuation in period 1.*

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

(iii) If the traders get different signals, then each 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. Case (i) 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.

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

We now show that if both traders receive B signals, a reversal occurs in which the expected period 2 price exceeds the initial price, while the period 1 price is lower than the initial price. This 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 there will not be adequate time for this learning.

We first define a benchmark for the amount of noise. We consider, as the benchmark, the case where $\delta = 0$, that is, the case where traders give no weight to outcomes in period 2 because the information will be public.[11] In this case, the traders do not act to conceal their information in period 1. The benchmark equilibrium period 1 quantities are given by solving equations (25), (26), and (27), ignoring the covariance terms:

$$
\begin{aligned}
g^A &= \left(\frac{-1}{2b}\right)(1-\alpha)(2\alpha-1) \\
g^B &= \left(\frac{1}{2b}\right)(\alpha(2\alpha-1))
\end{aligned}
$$

Note that the average quantity is zero, so $a_1 = a_0$. The notation '*' will indicate benchmark equilibrium quantities.

From the chartists' point of view there are four possible events: 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. 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 amount of noise was small, their expected

[11] Note that in the benchmark 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.

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 limiting equilibrium as σ^2 tends to zero.

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

$$\begin{aligned} E^*[\beta^B(P_1, g^B)|BB] &= 1 , \\ E^*[\gamma(P_1)|BB] = E^*[\gamma_{BB}(P_1)|BB] &= 1 . \end{aligned}$$

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. 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,

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

Period 2: For sufficiently small δ the strategic traders' strategies are close to the equilibrium strategies chosen in the benchmark equilibrium. The details are shown in the Appendix. But, in the limit 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. ||

V Concluding Remarks

The view that price movements reflect the contemporaneous arrival of new information presumes that learning is instantaneous. The empirical evidence on self-generating trade suggests that this is not the case. We have analyzed a model in which new information can arrive and the price does not instantaneously adjust. In fact, in adjusting to the new value, the price may initially move in the opposite direction. Such price movements are not anomalous since they are part of the learning process. The diffusion of information can be protracted because of strategic interaction.

The result that the initial price response to the arrival of new information can differ from the ultimate response does not require our specific information structure. Rather, the key requirement is that the marginal value of private signals be low, but that the revision implied by combining the signals be large. The more common 'truth plus noise' paradigm, in which a price reversal could not occur, is restrictive and probably unrealistic. It implies that initial price reactions to news are followed by smaller movements in the same direction.

Since we have chosen a specific model of the price formation process, it is natural to inquire as to whether the results would be forthcoming in alternative price formation models.[12]

Models of security price formation have concentrated on the specialist system, for example Glosten and Milgrom (1985) and Kyle (1985).

13

[12]Learning might also conceivably occur through the use of limit orders, i.e., orders to buy if the price falls below a specificed limit, and similarly for selling at a high price. Like other researchers in this area we restricted traders to using market orders and ruled out more complex strategies, such as limit orders or demand functions. In reality, limit orders have the disadvantage that if the price jumps, then outstanding limit orders are cleared at disadvantageous prices as the orders are crossed to get to the new market price. This can make limit orders undesirable. From a theoretical standpoint, we could consider a type of idealized limit order in which agents could condition on the current price, either by making their order subject to a price limit or even to the extent of submitting an entire demand function. A satisfactory treatment of these types of strategies has yet to be developed. As Dubey, Geanakoplos, and Shubik (1987) show, rational expectations models cannot, in general, be implemented as the equilibria of demand function submission games.

[13]A specialist is a regulated monopolist who aggregates orders and sets a price, taking up any imbalances

These models cannot be used in our setting because combining them with stratetgic interaction between multiple informed agents over time, and with our information structure, would lead to an intractable problem.[14] Nevertheless, we can informally ask whether our results would hold in similar models. In these models the specialist does all the learning. The key question is whether, in our setting, the specialist would observe two large sell orders, conclude that the two signals were both B signals and consequently set a high price. In this case there would be no reversal. This would appear to be possible in Kyle's model, but not in Glosten and Milgrom. In Glosten and Milgrom trade is sequential and so the reversal would seem likely. In Kyle, unless the trades arrived at exactly the same instant, it seems that the reversal would still occur.

The model here suggests that alternative information structures, combined with learning, can potentially explain some apparently anomalous asset price behavior. The link between asset prices and fundamentals may be more complicated than previously understood. The fundamental is the value of the asset conditional on the 'available' information. In our model information becomes 'available' as agents trade and learn. As our information structure illustrates, when information aggregation is not instantaneous the aggregation process may seem strange from the vantage point of the final asset value.

---

in supply and demand with his own inventory. Specialists are used on the New York Stock Exchange and American regional stock exchanges. Other financial markets use a variety of different price formation systems: many are more decentralized (e.g., the government bond and foreign exchange markets, the competing market makers on the London Stock Exchange and the American NASDAQ over-the-counter stock market), while others concentrate trade through specific channels (e.g., the Tokyo Stock Exchange Saitori system, the Toronto Stock Exchange CATS system which is also used in Paris and on other European exchanges) which also act as price-setting mechanisms. The theoretical models of specialist behavior assume that the specialist sets the price so as to earn zero expected profits on a trade by trade basis. In fact, while specialists are required to earn normal profits over some horizon, their exact objective function is unclear and probably more complicated.

[14]Glosten and Milgrom (1985) take traders' actions as exogenous functions of private signals while Kyle (1985) has only one informed trader.

Appendix

Derivations of equations (25) and (26)

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

$$\frac{\partial V_2^S(Q_1^1)}{\partial Q_1^1} =$$

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

where:

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

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

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

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

Note also that by Lemma 1:

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

And therefore:

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

Equation (A1) becomes:

$$-\frac{b}{\sigma^2}\int_{-\infty}^{\infty} V^A(P_1, Q_1^1)\epsilon f(P_1) dP_1 = -\frac{b}{\sigma^2}\text{cov}(V^A(P_1, Q_1^1), P_1) .$$

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:

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

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

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

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 Benchmark Equilibrium Strategies

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

$$\begin{aligned} 2bg^B &= a_0 + b[\alpha g^A + (1-\alpha)g^B] - (1-\alpha) - \left(\frac{\delta b}{\sigma^2}\right)\text{cov}^B \\ 2bg^A &= a_0 + b[\alpha g^A + (1-\alpha)g^B] - \alpha - \left(\frac{\delta b}{\sigma^2}\right)\text{cov}^A . \end{aligned}$$

These imply that:

$$2bg^B = \alpha(2\alpha - 1) - \left(\frac{\delta b}{\sigma^2}\right)[(2-\alpha)\text{cov}^B + \alpha\text{cov}^A] . \qquad \text{(A2)}$$

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:

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

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:

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

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:

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

Proceeding in a similar fashion:

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

In other words, for sufficiently small δ the strategic traders behave as if period 2 does not matter. But, in the limit as σ^2 tends to zero, the chartists value the asset at 1, and the strategic traders are inactive because they have no information advantage. Therefore, the price converges to 1. ||

References

Barclay, Michael, Robert Litzenberger, and Jerold Warner, 1990, Private information, trading volume, and stock return variances, *Review of Financial Studies*, 3(2), 233–254.

Canning, David, 1991, Asset price variability, mean reversion and quasi-rational expectations, Cambridge University Working Paper.

Dubey, Pradeep, John Geanakoplos, 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.

Glosten, Lawrence R. and Paul R. Milgrom, 1985, Bid, ask and transaction prices in a specialist market with heterogeneously informed traders, *Journal of Financial Economics*, 14, 71-100.

Grossman, Sanford and Joseph Stiglitz, 1976, Information and competitive price systems, *American Economic Review*, 66, 246–53.

Grossman, Sanford and Joseph Stiglitz, 1980, On the impossibility of informationally efficient markets, *American Economic Review*, 70, 393–408.

Kyle, Albert, 1985, Continuous auctions and insider trading, *Econometrica*, 53, 1315–35.

Meese, Richard, 1986, Testing for bubbles in exchange markets: A case for sparkling rates?, *Journal of Political Economy*, 94(2), 345–73.

Oldfield, George and Richard Rogalski, 1980, A theory of common stock returns over trading and non-trading periods, *Journal of Finance*, 35, 729–751.