

Herd on the Street: Informational Inefficiencies in a Market with Short-Term Speculation

KENNETH A. FROOT, DAVID S. SCHARFSTEIN, and
JEREMY C. STEIN*

ABSTRACT

Standard models of informed speculation suggest that traders try to learn information that others do not have. This result implicitly relies on the assumption that speculators have long horizons, i.e., can hold the asset forever. By contrast, we show that if speculators have short horizons, they may herd on the same information, trying to learn what other informed traders also know. There can be multiple herding equilibria, and herding speculators may even choose to study information that is completely unrelated to fundamentals.

HOW DO SPECULATORS' TRADING horizons affect the nature of asset prices? Does a market with numerous short-horizon traders perform less efficiently than one in which traders buy and hold? The classical response is that if speculators are rational, trading horizons should not affect asset prices. Even if a trader plans to sell his stock in five minutes, he cares about the expected price at that time. That price, in turn, depends on the expected price five minutes hence, and so on. Simple backwards induction then assures that even very short-horizon traders behave as if they were speculating on long-run fundamentals.

This traditional reasoning seems at odds with the way professional traders describe their jobs. Traders often emphasize that their objective is to predict near-term changes in asset prices. Rationally, they focus on learning anything that will help them do this more effectively. Often, it is claimed, such information has little to do with fundamentals. For example, according to one foreign exchange trader:¹

Ninety percent of what we do is based on perception. It doesn't matter if that perception is right or wrong or real. It only matters that other

* From Harvard University and NBER, Massachusetts Institute of Technology (MIT) and NBER, and MIT and NBER respectively. We thank Mike Fishman, Greg Mankiw, Stew Myers, Andre Perold, Julio Rotemberg, Andrei Shleifer, and seminar participants at numerous institutions for helpful comments. We are also grateful for research support from the Olin and Ford Foundations, MIT's International Financial Services Research Center, Batterymarch Financial Management, and the Division of Research at Harvard Business School.

¹ Quotation from the head of foreign exchange operations at Manufacturers Hanover Trust, "Making Book on the Buck," Mossberg, *Wall Street Journal* September 23, 1988.

people in the market believe it. I may know it's crazy, I may think it's wrong. But I lose my shirt by ignoring it. This business turns on decisions made in seconds. If you wait a minute to reflect on things, you're lost. I can't afford to be five steps ahead of everybody else in the market. That's suicide.

This account corresponds closely to the skeptical view of short-term trading offered by Keynes (1936) in *The General Theory*:

The actual, private object of most skilled investment today is to "beat the gun..." This battle of wits to anticipate the basis of conventional valuation a few months hence, rather than the prospective yield of an investment over a long term of years, does not even require gulls amongst the public to feed the maws of the professional; it can be played by professionals amongst themselves.

Keynes goes on to compare professional investors to beauty-contest judges who vote on the basis of contestants' expected popularity with other judges rather than on the basis of their absolute beauty.

In this paper, we develop a model of short-term trading that accords closely with these informal descriptions. We start with the assumption that there are at least some speculators who prefer to trade over short horizons. While we could explicitly model the rational behavior that gives rise to this assumption, in this paper we choose to take speculative horizons as given, and focus instead on the *implications* of short-term trading.²

We then show that the existence of short-term speculators can lead to a particular type of informational inefficiency. This occurs even though our model features fully rational agents. To see how inefficiencies can arise, consider an informed trader who plans to liquidate his position in the near future, before any public news arrives. He can profit on his information only if it is subsequently impounded into the price by the trades of similarly informed speculators. The trader therefore is made better off if there are others in the market *acting on the same information* that he is.

Positive informational spillovers of this sort are evident in the quotes above. In Keynes's beauty contest, the judges would be better off if they could coordinate their choices, even if they coordinate on somebody who is less than beautiful. Likewise, short-horizon traders would be better off if they could coordinate their research efforts on the same piece of information, even if that information is less revealing about the asset's long-run value. This is in

² There are at least two reasons why it might be rational for speculators to choose to trade over short horizons. First, some speculators such as money managers may need to prove to their clients or bosses that they are skilled investors. Promises of gains ten years hence would hardly justify a high current salary or the authority to continue managing a large portfolio (see Narayanan (1985) and Holmström and Ricart i Costa (1986)). Second, speculators who face imperfections in the capital market may find it relatively costly to finance long-horizon investment strategies. In particular, if speculators tie up their money in long-horizon investments, and at some point become credit constrained, they will not be able to take advantage of investment opportunities that arise in the future (see Shleifer and Vishny (1990)).

sharp contrast with most information-based asset pricing models (which implicitly assume a long horizon).³ In these models the information spillover is *negative*: a given trader is made better off if *nobody else* is trading on his information.

As will become clear, the negative spillovers that arise when traders have long horizons lead to contrarian research behavior. To take a concrete example, suppose that two variables, a and b , provide equally useful information about the value of a given security, and that an individual trader has the capacity to learn about either a or b , but not both. Efficient allocation of research effort requires that half of the traders study a , and the other half study b . And this is exactly what happens if traders have long horizons. If more than half are studying a , then a is more heavily impounded in the price than b . This negative spillover in a reduces the profits to those who study a , and so leads some investors to study b .

However, with short horizons, the outcome can be very different. Suppose everybody decides to study variable a . This can be an equilibrium, since there is no incentive to study b : even though b will affect the value of the asset when it is eventually liquidated, it will not be in its price in the near term, as nobody is trading on b information.

Thus, one sort of inefficiency created by short-horizon speculation is that traders may all tend to focus on one source of information, rather than on a diverse set of data. Moreover, these informational spillovers can be so powerful that groups of traders may choose to focus on very poor quality data, *or even on completely extraneous variables* that bear no relation at all to fundamentals.

It should be emphasized at the outset that the inefficiencies that we discuss in this paper pertain to traders' information acquisition strategies—the key implication of our model is that short horizons induce traders to allocate their research efforts in a less-than-socially-optimal fashion. However, as will become clear, our model assumes that the market is efficient at the *pricing* stage—once a (potentially quite narrow) set of information has been acquired, market prices incorporate this information in a Bayesian fashion. Thus the benchmark of “informational efficiency” that we define below reflects a different standard of market performance than is often studied in other works. In our usage, a market is “informationally efficient” only if *both* pricing is rational and research effort is allocated correctly.

Although its mechanism is very different, our model is not the first that attempts to capture Keynes's beauty-contest insight about the distinction between short and long trading horizons. Bubble models (Tirole (1982), Blanchard and Watson (1982)) address the same basic phenomenon. So do the models of noise trading and positive feedback trading by DeLong, Shleifer, Summers, and Waldman (1990a, 1990b). We will comment on these and other related work in what follows.

³ See, for example, Grossman (1976), Hellwig (1980), and Verrecchia (1982).

The paper is organized as follows. In Section I we lay out our basic model. Section II derives propositions relating to the informational efficiency of asset prices when traders have long and short horizons. We show that with short horizons, there can be “herding” on particular sources of information. In Section III we show that with sufficiently short horizons, speculators will trade on, and actually choose to study, completely extraneous noise. Section IV then relates our model to other work on the inefficiency of asset markets. Section IV also discusses extensions and implications of our model, stressing the connection between short speculative horizons and short-term corporate behavior.

I. The Model

We consider trade in a market for a single asset which is in fixed supply. The asset's only payout is a liquidating dividend of v , which is the sum of two normally distributed random variables:

$$v = a + b, \quad (1)$$

where a and b have means of zero and variances of σ_a^2 and σ_b^2 , respectively.

A. Types of Traders

The analysis takes as its starting point Kyle's (1985) model of informed trading. There are three types of traders in this model, all of whom are risk neutral and none of whom can observe v perfectly. The first type of traders are *market makers* who fill the orders of the two other types of traders: *informed speculators* and *liquidity traders*. Speculators and liquidity traders submit (“market”) orders to buy (sell) the asset from (to) long-lived market makers. Market makers cannot distinguish speculators' informed orders from liquidity traders' uninformed orders; they can observe only the total “order flow.” Because they are risk neutral and competitive, they earn zero profits. Thus, the market clearing price is the market makers' expectation of v , conditional on what they learn about v from the overall order flow. Since market makers do not have private information and are willing to hold until liquidation, they are best thought of as an uninformed fringe of long-term traders.

There are n speculators, n_a of whom have observed a and n_b of whom have observed b . Below, we allow speculators to choose which piece of information to become informed about, thereby endogenizing n_a and n_b . We assume that each speculator can costlessly observe a or b , but not both. This is intended to capture the idea that there are limits to how much any one trader can learn over short periods of time.⁴

⁴ It would also be easy to endogenize the total number of speculators, n . We could, for example, assume that there are costs of becoming informed about a or b , and that these costs differ depending on how readily available the information is. If there is free entry into speculation, traders will then enter until their profits net of information-acquisition costs are driven to zero.

As in the Kyle model, speculators are large enough to affect the market price, and they take this into account when formulating their demands. If speculators did not anticipate their effect on price, they would want to take infinite long (short) positions when their price forecast is below (above) their forecast of value. While the assumption is an attractive simplifying feature of the model, it is by no means crucial. We could also assume that speculators are risk averse and behave competitively (although such a model is computationally more burdensome). Thus, the reader should not be misled into thinking that our results come from some form of market manipulation by a large strategic trader.

Liquidity traders, in contrast to speculators, have inelastic demands for the asset: they wish to buy or sell a fixed quantity regardless of its price. Liquidity traders play an important role in essentially all models of information acquisition (see, for example, Grossman and Stiglitz (1980), and Kyle (1985)). In their absence, prices would reveal all the information in the economy, so there would be no return to becoming informed. In our model, as in others, liquidity trades result in prices that are noisy indicators of v , thus creating returns to information.

B. Timing of Trade

At an initial date 0, the n speculators choose whether to become informed about a or about b . Following this decision there are three trading periods. At date 1, speculators submit their asset demands. We assume that half of them have their orders executed at date 1 and that the other half have their orders executed at date 2. At the time they submit their orders, speculators do not know at what date their orders will be executed. The assumption of staggered execution implies that speculators' information is only gradually incorporated into prices. As we will see below, trades that are executed at date 1 can be profitable because more (of the same) information arrives at date 2. Speculators that "beat the gun" are therefore able to profit.

There are a couple of interpretations of these assumptions. The most literal one is simply that there are slight order-processing uncertainties which create the possibility that orders placed at the same time will not be executed simultaneously. One problem with this interpretation is that in reality, transactional risks of this type are usually small.

A second, perhaps more plausible interpretation is that speculators set out to acquire information about a or b at date 0, but receive it with a variable lag—one half of the speculators learn the information earlier than the other half. Under this interpretation, our specification implies that speculators must submit their demands without knowing whether they are "early" or "late" learners. We believe that similar results would emerge (although at a cost of much greater complexity) if we were to adopt the more realistic assumption that speculators could adjust their demands *after* observing current price and using it to make a noisy inference about the probability that they received the information relatively early or late.

After trading at date 1 or date 2, all speculators close out their positions at date 3. As we explain below, this means that they may have short horizons in that they may unwind their position before information is publicly released. What matters here is that the price at which informed traders get out of the market, even when they have short horizons, may contain more information than the price at which they get in.

Liquidity traders have date-1 and date-2 demands of ε_1 and ε_2 , respectively. At time 3 they also unwind their positions, so that $\varepsilon_3 = -(\varepsilon_1 + \varepsilon_2)$. We assume that ε_1 and ε_2 are normally distributed with mean zero and variance σ_ε^2 .

Given these assumptions, the order flow at date t , F_t , for $t = 1, 2$, can be written:

$$F_t = \frac{n_a}{2}q_a + \frac{n_b}{2}q_b + \varepsilon_t, \quad (2)$$

where q_a and q_b are the equilibrium demands of speculators informed about a and b , respectively.

Because the order flow at date 3 is just the negative of the cumulative order flows at dates 1 and 2, trade at date 3 contains no new information about v . All traders are simply closing out their positions from the previous two periods. This assumption simplifies the exposition greatly, but is of no qualitative importance. If, for example, we were to assume that liquidity traders at dates 1 and 2 did not close out their positions and that ε_3 was, like earlier realizations, drawn independently, there would be additional confirming evidence about v in date-3 orders. As a result, informed speculators whose trades were executed at date 2 would have positive expected profits when they closed out their positions. All of our results continue to hold (at least qualitatively) under this alternative assumption about date-3 trade.

Our main objective is to consider the effects of short versus long speculative horizons. To do this in a simple way, we assume that with probability α , the dividend is publicly announced at date 3, so that the date-3 trading price is v . With probability $(1 - \alpha)$, however, v does not become public until date 4. The parameter α is thus a simple measure of the extent to which speculators have long (α near 1) versus short (α near zero) horizons.

When speculators and liquidity traders close out their positions at date 3, the risk-neutral market makers reabsorb the supply of the asset and hold it until v is paid. In the case where information is not made public at date 3, the date-3 price is then equal to the market makers' conditional expectation of v . This is in turn equal to the date-2 price since no new information is contained in the date-3 order flow.

It should be emphasized that when speculators close out their positions at date 3 (and v is not publicly announced), they transact *only* with the uninformed market makers—not with a new set of informed short-term traders. If a new group of informed traders were to enter at date 3, they would wish to learn about both components of v , because they would be holding until liquidation. This would cause the date-3 price to reflect some

information about both of these components, much as if there had been a noisy public release of news at date 3. Thus, in the current formulation of the model, assuming that a new batch of informed traders enters at date 3, has an effect that is similar to assuming longer horizons for the first group of informed traders, and our results may be overturned.

At first glance, this casts some doubt on the general applicability of the results. However, this second batch of informed traders would have such a strong effect only if they were to hold the asset until its liquidation with certainty. In Section IV.B, we argue that a more realistic (although more complex) steady-state version of the model would likely yield results similar to those we present below, without restrictions on the entry of new generations of informed traders.

C. Market-Maker Pricing Rules

Based on the observed order flows and their conjectures about the trading strategies of the speculators, market makers form beliefs about the expected value of the asset. Since q_a and q_b depend on the realized values of a and b , the order flows provide information about v . Given that market makers' priors are normally distributed around a mean of zero, their posterior belief having seen the date-1 order flow, F_1 , is just F_1 multiplied by some constant, λ_1 . This constant is equal to the coefficient in a regression of v on F_1 . Thus, the price at date 1, p_1 , equals $\lambda_1 F_1$, where

$$\lambda_1 = \frac{\text{cov}[v, F_1]}{\text{var}[F_1]} = \frac{\text{cov}\left[a + b, \frac{n_a}{2}q_a + \frac{n_b}{2}q_b + \varepsilon_1\right]}{\text{var}\left[\frac{n_a}{2}q_a + \frac{n_b}{2}q_b + \varepsilon_1\right]}. \tag{3}$$

Similarly, the date-2 order flow provides information about the value of the asset. Given that the component of the order flow due to speculators' demands is the same at dates 1 and 2, and that the variances of ε_t are the same for the two periods, market makers put equal weight on these two order flows in forming their expectations about the asset's value at date 2. Thus, the market makers' conditional expectation of v is a function of the average order flow: $p_2 = \lambda_2(F_1 + F_2)/2$, where

$$\lambda_2 = \frac{\text{cov}\left[v, \frac{F_1 + F_2}{2}\right]}{\text{var}\left[\frac{F_1 + F_2}{2}\right]} = \frac{\text{cov}\left[a + b, \frac{n_a}{2}q_a + \frac{n_b}{2}q_b + \frac{1}{2}(\varepsilon_1 + \varepsilon_2)\right]}{\text{var}\left[\frac{n_a}{2}q_a + \frac{n_b}{2}q_b + \frac{1}{2}(\varepsilon_1 + \varepsilon_2)\right]}. \tag{4}$$

Since the information about v in the combined order flow is more precise than the information in F_1 alone, $\lambda_2 > \lambda_1$. Of course, the equilibrium value of λ_t depends on the trading strategies of the informed speculators, and these trading strategies, in turn, depend on the way in which the market maker sets prices.

D. Speculators' Demands

Speculators' demands depend on the information they observe. In forming their demands they take as given the number of speculators who are informed about a and b , the trading strategies of these speculators, and the pricing strategy of the market maker.

Consider then the decision facing speculator i who has observed a . First suppose that the dividend is to be announced at date 3. Since the speculator's order of q_a^i is equally likely to be executed at date 1 as at date 2, he or she expects to purchase at an average price of $\frac{p_1 + p_2}{2}$. Expected profits on each unit purchased are then $E[v - \frac{1}{2}(p_1 + p_2)|a]$. Next suppose that no announcement is made at date 3. If the order is executed at date 1, the speculator earns $E[p_2 - p_1|a]$, whereas if the order is executed at date 2 the speculator earns nothing since he or she buys at p_2 and sells at date 3 at a price of p_2 . Thus, in the case where the dividend is not announced until date 4, the speculator's expected profits are $\frac{1}{2}E[p_2 - p_1|a]$. Since the dividend is announced at date 3 with probability α and at date 4 with probability $1 - \alpha$, the expected utility of speculator i conditional on the realization of a is:

$$\begin{aligned} U_a^i &= q_a^i E \left[\alpha \left(v - \frac{p_1 + p_2}{2} \right) + (1 - \alpha) \frac{p_2 - p_1}{2} \middle| a \right] \\ &= q_a^i E \left[\alpha v - \frac{p_1}{2} + \frac{p_2}{2} (1 - 2\alpha) \right], \end{aligned} \quad (5)$$

where q_a^i is speculator i 's demand.

The expectation of v for a speculator who has observed a is just a . The observed value of a also enables the speculator to forecast prices at dates 1 and 2, since he or she knows the realization of a and the trading strategies of the other speculators who have observed a . However, the speculator knows nothing of the order flow generated either by liquidity traders or by speculators who have observed b . These flows have zero mean conditional on a . Thus, if a speculator's order is executed at time 1, his or her expectation of the price at that time is:

$$E[p_1|a] = \lambda_1 E[F_1|a] = \lambda_1 \left(q_a^i + \left(\frac{n_a}{2} - 1 \right) \bar{q}_a \right), \quad (6)$$

where \bar{q}_a denotes the conjectured demands of the $\frac{n_a}{2} - 1$ other speculators who are informed about a and have their orders executed at date 1.⁵ In contrast, the speculator's expectation of the date-2 price (which is independ-

⁵ Note that we assume from the outset that the demands of all other speculators informed about a are equal. Thus, we are focusing on symmetric equilibria.

ent of whether the order is executed at date 1 or date 2) is given by

$$E[p_2|\alpha] = \lambda_2 E\left[\frac{F_1 + F_2}{2} \middle| \alpha\right] = \lambda_2 \left(\frac{q_a^i + (n_a - 1)\bar{q}_a}{2}\right) \tag{7}$$

Given these expectations, speculator i who has observed a chooses q_a^i to maximize:

$$U_a^i = q_a^i \left\{ \alpha a - \frac{\lambda_1}{2} \left(q_a^i + \left(\frac{n_a}{2} - 1 \right) \bar{q}_a \right) + \frac{\lambda_2}{4} (1 - 2\alpha) (q_a^i + (n_a - 1)\bar{q}_a) \right\} \tag{8}$$

This expression shows clearly how trade by other speculators who observe a has spillover effects on the utility of speculator i who also observes a . If informed traders have long-term horizons (i.e., $\alpha = 1$), everything else being equal, more trade by other similarly informed speculators lowers speculator i 's expected utility: $\frac{dU_a^i}{dq_a} < 0$. Negative spillovers like these are standard in most information-based asset pricing models. Each agent expects to gain only to the extent that he or she can trade on information that is not already incorporated into price.

By contrast, if speculators have short horizons, and therefore liquidate their holdings before v is realized, spillovers are *positive*. In this case ($\alpha = 0$), $\frac{dU_a^i}{dq_a} > 0$. To see intuitively why this is so, consider the extreme case in which speculator i is the only one who trades on his or her information about a . Such a speculator cannot hope to earn a profit since there is no way for a to get impounded further into the price before his or her position must be unwound.

If, however, other speculators who are informed about a trade aggressively, speculator i will earn profits if his or her order is executed at date 1. This occurs because a great deal of additional information about a is later impounded in the date-2 price, and speculator i will sell at this price. Thus, if the speculator cannot hold the asset until it is liquidated, expected profits increase in the amount of trade by similarly informed traders. Taking the demands of other “ a -speculators” and the market depth parameters λ_1 and λ_2 as given, the first-order condition for q_a^i is:

$$q_a^i = \frac{4\alpha a - 2\lambda_1 \left(\frac{n_a}{2} - 1 \right) \bar{q}_a + \lambda_2 (1 - 2\alpha) (n_a - 1) \bar{q}_a}{4\lambda_1 - 2\lambda_2 (1 - 2\alpha)}. \tag{9}$$

We can derive an analogous expression for speculators who are informed about b , replacing a with b throughout.

Equation (9) shows that if a -speculators hold the asset until liquidation ($\alpha = 1$), their demands are “strategic substitutes” in the terminology of

Bulow, Geanakoplos, and Klemperer (1985). As other traders become more aggressive, not only do i 's expected profits fall (due to the negative spillover), but i also trades *less*: holding all else constant, $\frac{dq_a^i}{d\bar{q}_a} < 0$. This derives from the fact that more information about a is already in the price and so the marginal returns from trading on a are lower.⁶

The more interesting case is when speculators liquidate their holdings before the dividend is known ($\alpha = 0$). In that case, demands are "strategic complements." When rival speculators trade more aggressively, each speculator wishes to trade more aggressively as well: $\frac{dq_a^i}{d\bar{q}_a} > 0$. The marginal return from trading increases because more news about a will be in the price when speculators *sell*, and thus they stand to gain more at that time. In general, trading by similarly informed speculators is more likely to be a strategic complement the smaller is α . Strategic complementarity of this sort is the crucial feature of our model and it gives rise to the herding equilibria that we focus on below.⁷

II. Equilibrium

In order to solve for the equilibrium of this game, we first focus on the trading subgame which takes n_a and n_b as given. Then we move to the earlier research stage of the game and allow traders to choose which source of information to study. The solution to this research game endogenizes n_a and n_b . In an interior research equilibrium, the expected utilities of a - and b -speculators must be equalized. By itself, however, this condition turns out to be too weak to support an equilibrium. Thus, we need to impose an

⁶ Many economic games exhibit strategic substitutability, most notably the Cournot model of product-market competition. When an industry rival increases its production, all firms reduce their production because the market price is lower and hence the marginal returns from production are lower.

⁷ Strategic complementarities are present in numerous other models including the product-market model of Bertrand competition with differentiated goods. In that model, firms lower their prices in response to rivals' price decreases: in contrast to the Cournot model, a firm becomes more aggressive in response to increased aggressiveness by rivals. But note that while our model has positive spillovers, the Bertrand model has negative spillovers in that a rival's more aggressive pricing strategy lowers the firm's expected profits. In this sense, our model is closest to the technology adoption models of Farrell and Saloner (1985) and Katz and Shapiro (1985) which feature both strategic complementarities and positive spillovers: firms are made better off when others adopt the same technology and this leads them to coordinate their technology choices. Another example is Scharfstein and Stein (1990) who show that reputational concerns in the labor market can generate positive spillovers in investment and generate herd behavior among corporate managers. Spillovers and strategic complementarities in financial markets have also been explored. For example, Admati and Pfleiderer (1988) present a model in which liquidity traders prefer to trade at the same time as other liquidity traders while Pagano (1989) shows that they wish to trade in the same market. These agglomeration effects mitigate the adverse selection problem that liquidity traders typically face.

additional condition: that, in equilibrium, no speculator can be made better off by deviating and studying the other source of information. Operationally, this means that at the point where expected utilities are equalized, the expected utility of a given a -speculator must be decreasing in n_a (and similarly for the expected utility of b -speculators).⁸

Starting with the trading subgame where n_a and n_b are fixed, we note that in a symmetric equilibrium, $q_k^i = \bar{q}_k$ for $k = a, b$. Thus, solving (9) for an equilibrium q_a , we have:

$$q_a = \frac{4\alpha a}{\lambda_1(n_a + 2) - \lambda_2(n_a + 1)(1 - 2\alpha)} \equiv \delta_a a, \tag{10}$$

where δ_a is defined by the equation. Similarly, in equilibrium,

$$q_b = \frac{4\alpha b}{\lambda_1(n_b + 2) - \lambda_2(n_b + 1)(1 - 2\alpha)} \equiv \delta_b b. \tag{11}$$

The variables δ_a and δ_b measure the aggressiveness with which a - and b -speculators trade.

These equations only tell us speculators' demands given their conjectures about the values of λ_1 and λ_2 chosen by the market makers. But, the chosen λ_1 and λ_2 themselves depend on speculators' trading strategies. As discussed above, λ_1 is just the regression coefficient of v on F_1 :

$$\lambda_1 = \frac{2(n_a \delta_a \sigma_a^2 + n_b \delta_b \sigma_b^2)}{n_a^2 \delta_a^2 \sigma_a^2 + n_b^2 \delta_b^2 \sigma_b^2 + 4\sigma_\varepsilon^2}. \tag{12}$$

Recall that δ_a and δ_b depend on both λ_1 and λ_2 so that this equation alone does not determine λ_1 . A similar expression holds for λ_2 :

$$\lambda_2 = \frac{2(n_a \delta_a \sigma_a^2 + n_b \delta_b \sigma_b^2)}{n_a^2 \delta_a^2 \sigma_a^2 + n_b^2 \delta_b^2 \sigma_b^2 + 2\sigma_\varepsilon^2}. \tag{13}$$

We have not been able to derive closed form expressions for the endogenous variables λ_1 , λ_2 , δ_a , and δ_b . However, it is worth noting from equations (12) and (13) that λ_2 is greater than λ_1 as we claimed earlier.

Provided the four equations, (10)–(13), have a solution, we can calculate the expected utilities of a - and b -speculators for any fixed n_a and n_b . Denote these expected utilities by EU_a and EU_b , respectively. Note that these expected utilities are calculated before a and b are realized and are not to be confused with a speculator's expected utility conditional on observing a .

⁸ This refers to a total derivative of expected utility with respect to the number of traders: implicitly we are assuming that n_a and n_b become public after the research game has been played and before the trading game commences. Therefore, when a single speculator deviates and changes research strategies, all other players take this into account in playing the trading game.

Given the expected utilities that follow from an arbitrary n_a and n_b , we wish to determine the values of n_a and n_b that are consistent with equilibrium in the earlier research stage of the game. As noted above, an interior equilibrium requires both that 1) $EU_a = EU_b$; and 2) $\frac{dEU_a}{dn_a} < 0$ and $\frac{dEU_b}{dn_b} < 0$.

In order to evaluate any such research equilibrium, we need a benchmark for informational efficiency. It is easiest to come up with an unambiguous benchmark in the case where $\sigma_a^2 = \sigma_b^2$, so for the remainder of this section we restrict our attention to this case. Intuitively, in this symmetric case, the informationally efficient outcome should involve $n_a = n_b$. Indeed, it is straightforward to show that with $\sigma_a^2 = \sigma_b^2$, the allocation $n_a = n_b$ optimizes a number of sensible measures of price informativeness. For example, we prove in the Appendix that the allocation $n_a = n_b$ would be chosen by a social planner seeking to minimize the average variance of prices about true value, where this average variance is given by:

$$\frac{1}{2}E(v - p_1)^2 + \frac{1}{2}E(v - p_2)^2 \quad (14)$$

where the expectation is taken over all realizations of a , b , ε_1 , and ε_2 . In this problem, the social planner chooses n_a and n_b to minimize the expectation in (14), given the λ_1 , λ_2 , F_1 , and F_2 that follow from this choice.⁹

With this benchmark in mind, we are now able to prove the following propositions for the case where $\sigma_a^2 = \sigma_b^2$. (Again, see the Appendix for details.)

Proposition 1: If speculators have "long" horizons (α sufficiently close to 1), the informationally efficient outcome is an equilibrium of the research game.

Proposition 2: If speculators have "short" horizons (α sufficiently close to zero), the informationally efficient outcome is not an equilibrium of the research game—i.e., any research equilibrium must involve a degree of "herding."

Figures 1 and 2 help to understand the intuition behind the propositions. They are constructed for an example in which $n = 20$, $\sigma_a^2 = \sigma_b^2 = \sigma_\varepsilon^2 = 1$. In Figure 1, $\alpha = 0.25$; in Figure 2, $\alpha = 0.03$.¹⁰ On the figures' vertical axes are the levels of expected utility for a - and b -speculators. On the horizontal axes are the number of speculators informed about a , n_a , holding the total number of speculators n fixed.

Figure 1 shows the expected utility levels of a - and b -speculators for a "large" value of α . Expected utility is clearly decreasing in the number of

⁹ Note that our definition of informational efficiency is ad hoc in that it is not based on an explicit measure of social welfare.

¹⁰ None of the qualitative properties of the figures appear to depend on the specific parameters used.

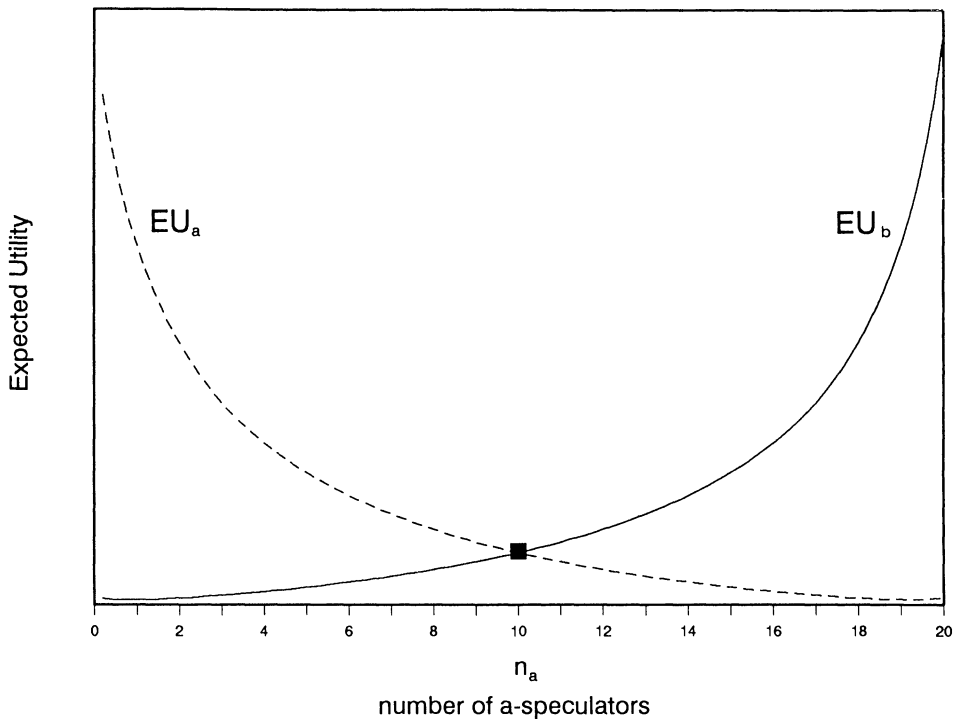


Figure 1. Equilibrium when traders have relatively long horizons. EU_a and EU_b are the expected utilities of a - and b -speculators, respectively. These expected utilities are graphed as a function of the number of a -speculators, n_a , holding the total number of speculators fixed. In the construction of this figure the number of speculators is 20; the variances of a , b , and ε are unity; and α , the probability that the liquidating dividend is announced at date 3, is 0.25.

similarly informed traders; for large values of α the usual “contrarian” effects in research dominate. The only point where $EU_a = EU_b$ occurs at $n_a = n/2$. To see intuitively why this allocation is an equilibrium, suppose that an extra speculator chooses to study a , so that $n_a > n_b$. The figure shows that this deviation leads to $EU_a < EU_b$, thus making the deviator worse off. Therefore, an allocation in which $n_a > n_b$ cannot be an equilibrium, nor can an allocation in which $n_b > n_a$. In sum, even though α is less than one, the negative spillovers and strategic substitutability effects can dominate, yielding an equilibrium that is similar to those in other information-based asset pricing models.

Figure 2 depicts the corresponding levels of utility when α is small. Note that in this case expected utility is no longer monotonic in the number of similarly informed traders. Indeed, it is increasing when the allocation of information is approximately symmetric; positive spillovers and strategic complementarities dominate.

There are now three points at which $EU_a = EU_b$. The efficient symmetric allocation (which involves equal numbers of a - and b -speculators) is one such point, but it is no longer an equilibrium. This is because EU_a is increasing in

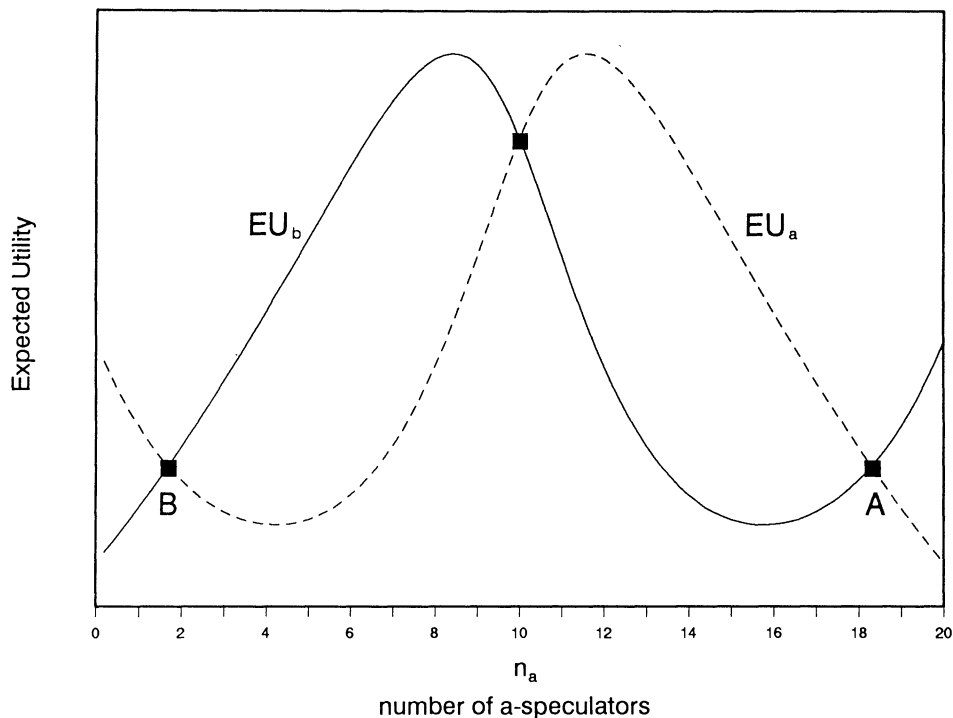


Figure 2. Equilibrium when traders have relatively short horizons. EU_a and EU_b are the expected utilities of a - and b -speculators, respectively. These expected utilities are graphed as a function of the number of a -speculators, n_a , holding the total number of speculators fixed. In the construction of this figure the total number of speculators is 20; the variances of a , b , and ε are unity; and α , the probability that the liquidating dividend is announced at date 3, is 0.03.

the number of a -speculators at this point. A b -speculator would thus be made better off by deviating and studying a —positive spillovers dominate at the symmetric allocation. In contrast, the two other intersection points in Figure 2, designated as A and B , are equilibria. To see this, consider the A equilibrium. If an additional b -speculator deviated to study a , he or she would be made worse off since EU_a is downward sloping in the neighborhood of A . Similarly, if an a -speculator deviated to study b , he or she would be worse off.

These herding equilibria are clearly inefficient. As we decrease α below the level seen in Figure 2, then the herding equilibria become more extreme, eventually reaching the corner solutions where $n_a = n$ and $n_b = n$.

III. Trading on Noise

In the discussion above, we assumed that a and b are components of v —each piece of information is actually helpful in predicting fundamental value. We now relax this assumption. We ask whether the informational

spillovers are strong enough to make possible herding on information *that is completely unrelated to fundamentals*.

Suppose that n_v traders know v and that $n_c = n - n_v$ traders know a variable c which is independent of fundamentals. The utility of v -speculators is essentially as discussed in earlier sections: it is given by (5), with a replaced everywhere by v . Similarly, the i th v -speculator's—or "fundamentalists"—demand is given by an expression analogous to (10), which in the symmetric case $q_v^i = \bar{q}_v$ can be written as:

$$q_v = \frac{4\alpha v}{\lambda_1(n_v + 2) - \lambda_2(n_v + 1)(1 - 2\alpha)} \equiv \delta_v v. \tag{15}$$

The comparable expression for the "chartist" trader who learns c is slightly different.¹¹ Because c is uncorrelated with v , the i th chartist's expected utility conditional on observing c is given by:

$$U_c^i = q_c^i \left\{ -\frac{\lambda_1}{2} \left(q_c^i + \left(\frac{n_c}{2} - 1 \right) \bar{q}_c \right) + \frac{\lambda_2}{4} (1 - 2\alpha) (q_c^i + (n_c - 1) \bar{q}_c) \right\}. \tag{16}$$

It is clear from (16) that chartist traders will not want to trade if α is sufficiently near one. If there is a high probability that speculators will sell out at a price equal to fundamentals, v , then chartists—who cannot forecast any component of v —would consistently lose money if they were to trade. Thus, chartists can participate in a trading subgame only if there is a sufficiently high probability that they will be selling out before all information becomes public.

Assuming a symmetric equilibrium in the trading subgame ($q_c^i = \bar{q}_c$), the first-order condition for the i th chartist implies that if chartists are to trade, the following must hold:

$$n_c = \frac{2\lambda_1 - \lambda_2(1 - 2\alpha)}{\lambda_2(1 - 2\alpha) - \lambda_1}. \tag{17}$$

The market makers' problem is slightly changed, because v covaries only with the component of the order flow attributable to fundamentalists. Thus, market makers now set market depth parameters, λ_1 and λ_2 , according to:

$$\lambda_1 = \frac{\text{cov}[v, F_1]}{\text{var}[F_1]} = \frac{2\delta_v n_v \sigma_v^2}{\delta_v^2 n_v^2 \sigma_v^2 + \delta_c^2 n_c^2 \sigma_c^2 + 4\sigma_\epsilon^2}, \tag{18}$$

$$\lambda_2 = \frac{\text{cov}\left[v, \frac{F_1 + F_2}{2}\right]}{\text{var}\left[\frac{F_1 + F_2}{2}\right]} = \frac{2\delta_v n_v \sigma_v^2}{\delta_v^2 n_v^2 \sigma_v^2 + \delta_c^2 n_c^2 \sigma_c^2 + 2\sigma_\epsilon^2}. \tag{19}$$

Of course, as before, we have that with informed trading $\lambda_2 > \lambda_1 > 0$.

¹¹ Chartism is one example of trading on information unrelated to underlying value, or "noise." For a different model of the interaction between chartists and fundamentalists see Frankel and Froot (1989).

As we have already mentioned, chartists cannot trade profitably when horizons are sufficiently long. As a result, when α is near one, there is a unique research equilibrium in which all speculators study fundamentals, $n_v = n$.

It is also clear that chartists cannot trade profitably if they are the only ones in the market—i.e., if there are no fundamentalists. This follows immediately from an inspection of (18) and (19). If nobody is trading on fundamentals, then the order flow is completely uninformative about v and market makers set $p_1 = p_2 = 0$.

If, by contrast, there are some active fundamentalists, it is possible to support a trading equilibrium that involves chartist trading, for a range of shorter horizons. The presence of fundamentalists in the trading game ensures that order flow contains some information about v , so that $\lambda_2 > \lambda_1 > 0$. This in turn creates room for chartists to trade profitably, provided that there are enough of them to move the price with c in the short run. We prove the following proposition in the Appendix:

Proposition 3: If $n_c > n_v$, and α is sufficiently small, then there exists an equilibrium in the trading sub-game in which both chartists and fundamentalists submit positive market orders and earn positive expected profits ($EU_c^i > 0$). In addition, there is always an equilibrium in which fundamentalists trade actively, but chartists do not.

The positive spillovers and strategic complementarities allow chartists in the aggregate to bootstrap their way into profitable trading. Given that other chartists are trading, each expects the price to move with c and therefore each trades actively.¹²

Proposition 3 suggests that if a large enough number of traders are endowed with information about c , they may trade on it and earn profits. It does not say, however, that speculators will actually *choose* to study c if they could instead learn v . Nevertheless, we have constructed numerical examples in which some speculators choose to study v and others choose to study c .¹³

Figures 3 and 4 illustrate the effects at work. Once again, the vertical axes measure traders' expected utility levels, EU_v and EU_c , and the horizontal axes measure the number of chartist speculators n_c , given n . As before, the figures are constructed for an example in which $n = 20$, $\sigma_v^2 = \sigma_c^2 = \sigma_\varepsilon^2 = 1$; in Figure 3, $\alpha = 0.01$, and in Figure 4, $\alpha = 0.001$. Note that we graph only the relevant range, $n_c > n/2$.

Figure 3 exhibits the case in which horizons are relatively long-term, i.e., α is relatively large. It is immediately clear that $EU_v > EU_c$, regardless of the number of traders informed about each. Chartists receive positive utility from trading and therefore would trade actively in the trading subgame. However,

¹² The potential for traders who reduce the informational efficiency of prices to "create their own space" for profitable activity is also seen in Stein (1987) and DeLong et al. (1990a).

¹³ Numerical examples were required because we have not been able to derive explicitly the roots of the polynomial expression given by $EU_c = EU_v$.

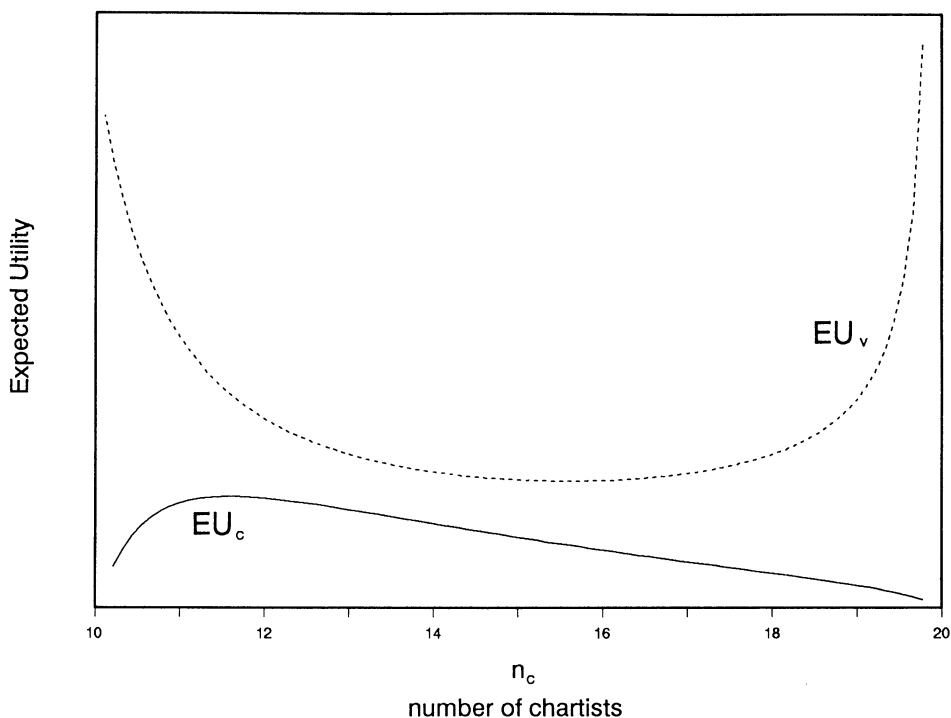


Figure 3. Equilibrium when fundamentalists and chartists have relatively long horizons. EU_v and EU_c are the expected utilities of fundamentalist speculators (those who observe the asset's true value, v) and chartist speculators (those who observe pure noise, c), respectively. These expected utilities are graphed as a function of the number of chartist speculators, n_c , holding the total number of speculators fixed. In the construction of this figure the total number of speculators is 20; the variances of v , c , and ε are unity; and α , the probability that the liquidating dividend is announced at date 3, is 0.01.

once we allow speculators to choose which source of information to study, none chooses c . The only research equilibrium is $n_c = 0$, where all traders choose to study v .

Figure 4 is a comparable graph for the case in which α is relatively small. Here it is unlikely that new outside information arrives before the current v and c traders sell. As a result, informational spillovers are a more important factor in determining expected utility levels. As before, there is an efficient research equilibrium (not shown on the graph) where all traders choose to study v ; $n_c = 0$. Note, however, that if speculators conjecture that a majority will become chartists, then there are two other points, shown in the graph, at which expected utilities are equalized. The point with fewer chartists is not an equilibrium since EU_c is upward sloping at this point; informational spillovers are so strong at this point that more speculators would want to study c . By contrast, the point labeled C in the figure is an equilibrium. Interestingly, at point C, n_c is much greater than n_v : if c is studied at all in

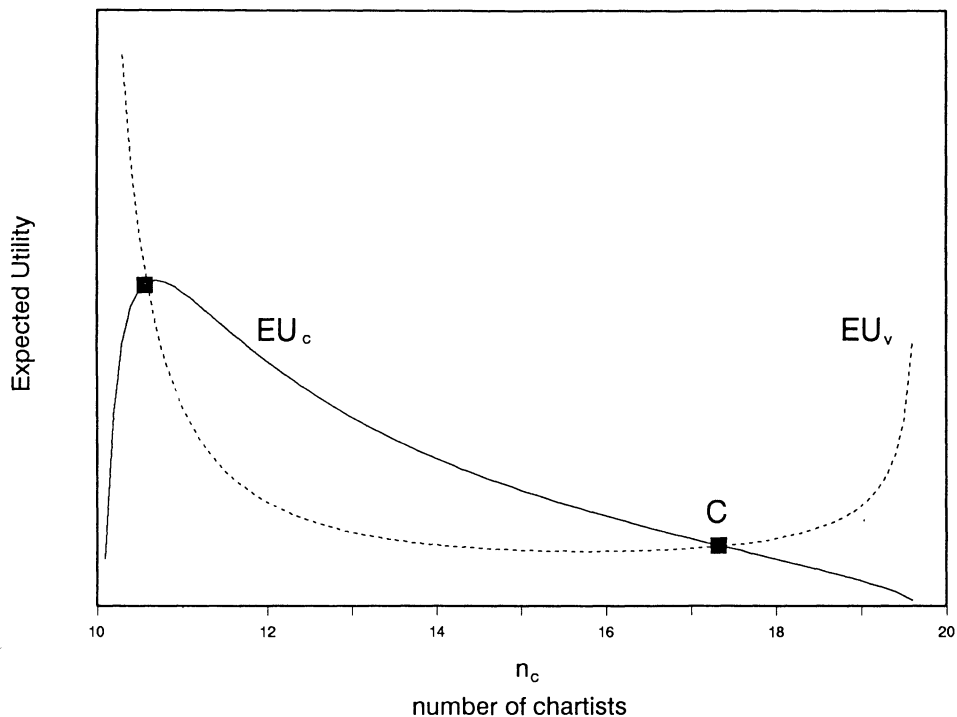


Figure 4. Equilibrium when fundamentalists and chartists have relatively short horizons. EU_v and EU_c are the expected utilities of fundamentalist speculators (those who observe the asset's true value, v) and chartist speculators (those who observe pure noise, c), respectively. These expected utilities are graphed as a function of the number of chartist speculators, n_c , holding the total number of speculators fixed. In the construction of this figure the total number of speculators is 20; the variances of v , c , and ε are unity; and α , the probability that the liquidating dividend is announced at date 3, is 0.001.

equilibrium, the majority of traders will want to study it, even though c is completely unrelated to fundamental value.

IV. Discussion

A. Inefficiencies in Markets with Short-Term Trading

In typical models of informed trading, informational externalities are negative. In such models, which effectively feature speculators with long horizons, the returns to acquiring information fall as the number of other identically informed traders increases. Negative externalities of this sort encourage contrarian information acquisition.

In contrast, our results are driven by *positive* informational spillovers: as more speculators study a given piece of information, more of that information disseminates into the market, and therefore, the profits from learning that information early increase. This implies that profit-maximizing speculators

may choose to ignore some information about fundamentals. In equilibrium, speculators herd: they acquire “too much” of some types of information and “too little” of others.

There are other classes of models in which short trading horizons can lead to inefficiencies. The first—that of fads and noise trading—focuses on the implications of less-than-fully rational traders. DeLong, Shleifer, Summers, and Waldman (1990a) demonstrate that shorter horizons on the part of “smart money” traders allow the behavior of noise traders to have a greater impact on asset prices. They argue that “if sophisticated investors’ horizons are long...arbitrage becomes less risky and prices approach fundamental values” (p. 713). In related work, DeLong et al. (1990b) examine “positive-feedback” traders who predictably extrapolate past price trends. In this model, rational speculators can increase their overall profits by taking advantage of the short-horizon extrapolation of positive-feedback traders. In undertaking this trading strategy, they drive the asset price *away* from its fundamental value.¹⁴

A second class of models in which inefficiencies arise from short-term speculative horizons is that of rational bubbles. These models employ only rational speculators, but prices nevertheless exhibit extraneous fluctuations. Traders have short-term horizons in that they are not able to enforce infinite-horizon arbitrage conditions. As a result, prices may contain an extraneous component which grows at the discount rate. If this component is present, the market will be “stuck” on an inefficient path along which prices eventually explode. The efficient equilibrium is also possible: if the initial price is equal to its present value level, then the bubble can never get started.

One problem with this latter type of model is that it offers no mechanism for what drives the market away from efficiency. Indeed, in bubble models sensible candidates would, if anything, drive the economy toward the *efficient* allocation. The infinite-horizon transactions that are ruled out by assumption in such models become hugely profitable as the bubble—the wedge between prices and the present value of fundamentals—explodes. It is easy to believe that agents facing very large wedges would attempt such transactions, which by induction would eliminate bubble-type inefficiencies from the start. Our approach may be preferable in this regard, in that the positive spillovers drive the market away from the efficient outcome.

B. An Infinite-Horizon Extension

As noted in Section 2, there might be no herd behavior in the current formulation of the model if it were certain that a new group of informed speculators would enter at date 3. However, our model could likely be extended so as to handle overlapping generations of informed traders without

¹⁴ Frankel and Froot (1989) present a model in which optimizing portfolio managers must choose between the advice of rational fundamentals traders and chartists.

losing the principal results. One possibility is an infinite-horizon, steady-state approach which we describe briefly.

Suppose that at the beginning of each period there are k pieces of information that speculators can study. At the end of the period, one of these pieces of information will be publicly announced, although it is not known initially which it will be. At the beginning of the next period a new piece of information, which was previously impossible to learn about, is then available to be studied. For example, suppose that a company is always engaged in k R & D projects, about which speculators may learn. In each period, one project reaches a conclusion and its results are revealed publicly, although speculators cannot predict in advance which project it will be. In the next period a new project is begun in its place.

Under these circumstances, herding equilibria like those described above may arise. Suppose that each generation herds on a single piece of information. At the time they make this choice, they are uncertain about what the next generation will happen to herd on. To see that this could be an equilibrium, consider an individual speculator's incentive to deviate from the herd by studying a different piece of information. The speculator can profit from the deviation only if the piece of information that he or she alone studies is publicly revealed or if the next generation herds on it. If k is large, neither outcome is likely, and the incentive to deviate is small.¹⁵

C. Empirical Implications

Because the mechanism driving our results is different from that in related models, it has different empirical implications. First, unlike the noise trading and bubble models discussed above, our model implies that prices will follow a random walk: no publicly available information will help in predicting future price changes. (Of course, informed traders can partially predict future price changes because their information has not been impounded fully into prices.)

Second, the model can help to make sense of the often puzzling behavior of many market participants. In practice, short-term traders often use forecasting methods that appear at best tangentially related to fundamental values. Chartism is one example of such a method. Economists and even traders seem to agree that there are better methods of determining long-run value. Yet, the very fact that a large number of traders use chartist models may be enough to generate positive profits for those traders who already know how to chart. Even stronger, when such methods are popular, it is optimal for speculators to *choose* to chart. They rationally ignore opportunities to learn about v , the realization of which is a distant "five steps ahead." Such an equilibrium can persist even if chartist methods contain no relevant long-term information.

¹⁵ Note that even in this infinite-horizon model, the herding equilibria are in no sense bubbles — the price is always equal to the expectation of present value (conditional on some information set) and no transversality conditions are violated.

The herding equilibria also suggest that traders may focus on different variables at different times. For example, in the infinite-horizon model sketched above, each new generation of speculators switches to studying an entirely different source of information. This kind of behavior sounds reminiscent of markets which track certain variables closely for short periods of time. Of course, if the underlying valuation model is changing, one would expect this type of behavior anyway, but it seems to us that the market's romance with individual variables is often extremely brief and only tenuously connected with underlying fundamentals.

D. The Welfare Effects of Short Speculative Horizons

Short speculative horizons can affect social welfare through two distinct channels. First, short-term trading can, as we have demonstrated, have a direct negative impact on the informational quality of asset prices. This in turn can lead to less-informed allocational decisions if agents look to asset prices to guide production decisions.

Second, there may be a connection between short-term trading on the part of market participants and short-sighted behavior by corporations with respect to certain investments. Recent work (see, e.g., Stein (1989)) has shown that if an investment cannot be directly observed by the stock market, there will be a tendency for managers concerned with stock-price maximization to skimp on this investment. For example, managers may be reluctant to cut prices aggressively to gain market share, because the accompanying reduction in current profits may be interpreted as bad news about product demand, rather than as a program of increased investment in market share.

However, this work on corporate "short-termism" typically takes the information asymmetry between managers and the market as *exogenous*. Continuing with the above example, it may simply be assumed that the market cannot observe whether or not managers have adopted a legitimate strategy of cutting prices to gain market share.

Our results on the consequences of short-term trading suggest that such informational asymmetries can be endogenized within a broader model. In particular, short-term trading in the capital market might imply that even if some traders *could* potentially learn more about managers' pricing strategies, they might rationally choose not to, focusing instead on, say, predicting next month's earnings announcement. Thus, short-term trading might help create the informational problems that are a necessary precondition for short-sighted investment behavior by corporations.

Appendix

Proof of Proposition 1: We start by showing that $n_a = n_b$ is an equilibrium. First, given the expression for an a -speculator's utility in (8) and the analog for a b -speculator, inspection reveals that $EU_a = EU_b$ at $n_a = n_b$. We next need to check that a b -speculator could not be made better off by deviating

and studying a . This amounts to showing that EU_a is decreasing in n_a at $n_a = n_b$ (which by symmetry implies that EU_b is decreasing in n_b). We start with two observations: i) $\frac{d\lambda_1}{dn_a} = \frac{d\lambda_2}{dn_a} = 0$ at $n_a = n_b$ (by the symmetry of equations (12) and (13) with respect to n_a and n_b); and ii) $\frac{dEU_a}{d\delta_a} = 0$ and similarly for b -speculators (by the envelope theorem). Using these facts, we can differentiate the expectation of equation (8) with respect to n_a and show that $\frac{dEU_a}{dn_a}$ is equal to a positive number times the quantity $-\lambda_1 + \lambda_2(1 - 2\alpha)$, which is negative for α sufficiently near one. Thus, $n_a = n_b$ satisfies our equilibrium conditions when there are long horizons.

Second, we show that the allocation $n_a = n_b$ is the solution to the social planner's problem, which we have defined in the text as:

$$\min_{n_a, n_b} \frac{1}{2} E(v - p_1)^2 + \frac{1}{2} E(u - p_2)^2. \quad (A1)$$

The choice of n_a and n_b affects λ_1 , λ_2 , F_1 , and F_2 . Thus, the derivative of expression (A1) with respect to n_a (recognizing that $n_a = n - n_b$) is given by:

$$E\left[(v - p_1) \left(-\frac{d\lambda_1}{dn_a} F_1 - \frac{1}{2}(q_a - q_b)\right)\right] + E\left[(v - p_2) \left(-\frac{d\lambda_2}{dn_a} \frac{F_1 + F_2}{2} - \frac{1}{2}(q_a - q_b)\right)\right], \quad (A2)$$

where we have used the fact that at $n_a = n_b$, $\frac{n_a}{2} \frac{\partial q_a}{\partial n_a} = \frac{n_b}{2} \frac{\partial q_b}{\partial n_b}$.

The fact that market makers set prices in a Bayesian fashion implies that the forecast errors, $v - p_1$ and $v - p_2$, are independent of all information including the order flows:

$$E[(v - p_1)F_1] = E\left[(v - p_2) \frac{F_1 + F_2}{2}\right] = 0.$$

Using this observation, (A2) becomes:

$$-\frac{1}{2} E[((v - p_1) + (v - p_2))q_a] + \frac{1}{2} E[((v - p_1) + (v - p_2))q_b]. \quad (A3)$$

The first term is the negative of expected utility from learning a in the case where $\alpha = 1$. The second term is the expected utility from learning b . As shown above, these utilities are equalized at $n_a = n_b$. Therefore, the derivative of (A1) with respect to n_a is equal to zero (and similarly for n_b), and thus the social planner's optimization conditions are satisfied at this point.

Proof of Proposition 2: As in Proposition 1, $EU_a = EU_b$ at $n_a = n_b$. Again, we note that $\frac{dEU_a}{dn_a}$ is equal to a positive number times the quantity $-\lambda_1 + \lambda_2(1 - 2\alpha)$. This quantity is positive for α sufficiently near zero. Thus, for sufficiently small α , $n_a = n_b$ no longer satisfies our conditions for an equilibrium.

Proof of Proposition 3: Using equations (15), (17), (18), and (19) we can solve directly for the four endogenous variables δ_v , δ_c , λ_1 , and λ_2 as functions of α , n_v , and n_c . Algebra yields:

$$\lambda_1 = \left(\frac{j\sigma_v^2 n_v (k - 1)}{\sigma_\varepsilon^2 k} \right)^{1/2}, \tag{A4}$$

$$\lambda_2 = k \lambda_1, \tag{A5}$$

$$\delta_v = \frac{j}{\lambda_1}, \tag{A6}$$

$$\delta_c = \left(\frac{\sigma_\varepsilon^2 (4 - 2k - jn_v k)}{\sigma_c^2 n_c^2 (k - 1)} \right)^{1/2}, \tag{A7}$$

where $k = \frac{2 + n_c}{(n_c + 1)(1 - 2\alpha)}$ and $j = \frac{4\alpha(n_c + 1)}{(n_c - n_v)}$.

In order to have an interior solution, with all four variables positive and finite, we require $n_c > n_v$ and α sufficiently small. (If these conditions are not satisfied, one can see that it is impossible to have $\delta_c^2 > 0$.)

Note that even if these conditions are satisfied, there is always an equilibrium in which chartists do not trade at all. In this case, (17) need not hold. That the no-chartist equilibrium is self-sustaining is evident from (16)—if no chartists are trading, $\bar{q}_c = 0$ and U_c^i is negative, so that any single chartist would lose money by being the only one to trade.

REFERENCES

Admati, Anat and Paul Pfleiderer, 1988, A theory of intraday trading patterns, *Review of Financial Studies* 1, 3-40.
 Blanchard, Olivier J. and Mark W. Watson, 1982, Bubbles, rational expectations, and financial markets, in Paul Wachtel, ed.: *Crises in the Economic and Financial Structure* (Lexington Books, Lexington, MA).
 Bulow, Jeremy, John Geanakoplos, and Paul Klemperer, 1985, Multimarket oligopoly: Strategic substitutes and complements, *Journal of Political Economy* 93, 488-511.
 DeLong, J. B., Andrei Shleifer, Lawrence Summers, and Robert Waldman, 1990a, The economic consequences of noise traders, *Journal of Political Economy* 98, 703-738.
 ———, Andrei Shleifer, Lawrence Summers, and Robert Waldman, 1990b, Positive feedback investment strategies and destabilizing rational speculation, *Journal of Finance* 45, 379-395.
 Farrell, Joseph and Garth Saloner, 1985, Standardization, compatibility, and innovation, *RAND Journal of Economics* 16, 70-83.

- Frankel, Jeffrey A. and Kenneth A. Froot, 1989, Chartists, fundamentalists, and the demand for dollars, in Anthony Courakis and Mark Taylor, eds.: *Policy Issues for Interdependent Economies* (MacMillan, London).
- Grossman, Sanford J., 1976, On the efficiency of competitive stock markets where traders have diverse information, *Journal of Finance* 31, 573–585.
- and Joseph E. Stiglitz, 1980, On the impossibility of informationally efficient markets, *American Economic Review* 70, 393–408.
- Hellwig, Martin F., 1980, On the aggregation of information in competitive markets, *Journal of Economic Theory* 22, 477–498.
- Holmström, Bengt and Joan Ricart i Costa, 1986, Managerial incentives and capital management, *Quarterly Journal of Economics* 101, 835–860.
- Katz, Michael and Carl Shapiro, 1985, Network externalities, competition, and compatibility, *American Economic Review* 75, 429–440.
- Keynes, John M., 1936, *The General Theory of Employment, Interest and Money* (Macmillan, London).
- Kyle, Albert S., 1985, Continuous auctions and insider trading, *Econometrica* 53, 1315–1336.
- Narayanan, M. P., 1985, Observability and the payback criterion, *Journal of Business* 58, 309–323.
- Pagano, Marco, 1989, Endogenous market thinness and stock price volatility, *Review of Economic Studies* 56, 269–288.
- Scharfstein, David S. and Jeremy C. Stein, 1990, Herd behavior and investment, *American Economic Review* 80, 465–479.
- Shleifer, Andrei and Robert Vishny, 1990, Equilibrium short horizons of investors and firms, *American Economic Review* 80, 148–153.
- Stein, Jeremy C., 1987, Informational externalities and welfare-reducing speculation, *Journal of Political Economy* 95, 1123–1145.
- , 1989, Efficient capital markets, inefficient firms: A model of myopic corporate behavior, *Quarterly Journal of Economics* 104, 655–670.
- Tirole, Jean, 1982, On the possibility of speculation under rational expectations, *Econometrica* 50, 1163–1182.
- Verrecchia, Robert, 1982, Information acquisition in a noisy rational expectations economy, *Econometrica* 50, 1415–1430.