



Multidimensional Uncertainty and Herd Behavior in Financial Markets

Christopher Avery; Peter Zemsky

The American Economic Review, Vol. 88, No. 4 (Sep., 1998), 724-748.

Stable URL:

<http://links.jstor.org/sici?sici=0002-8282%28199809%2988%3A4%3C724%3AMUAHBI%3E2.0.CO%3B2-%23>

The American Economic Review is currently published by American Economic Association.

Your use of the JSTOR archive indicates your acceptance of JSTOR's Terms and Conditions of Use, available at <http://www.jstor.org/about/terms.html>. JSTOR's Terms and Conditions of Use provides, in part, that unless you have obtained prior permission, you may not download an entire issue of a journal or multiple copies of articles, and you may use content in the JSTOR archive only for your personal, non-commercial use.

Please contact the publisher regarding any further use of this work. Publisher contact information may be obtained at <http://www.jstor.org/journals/aea.html>.

Each copy of any part of a JSTOR transmission must contain the same copyright notice that appears on the screen or printed page of such transmission.

JSTOR is an independent not-for-profit organization dedicated to creating and preserving a digital archive of scholarly journals. For more information regarding JSTOR, please contact jstor-info@umich.edu.

Multidimensional Uncertainty and Herd Behavior in Financial Markets

By CHRISTOPHER AVERY AND PETER ZEMSKY *

We study the relationship between asset prices and herd behavior, which occurs when traders follow the trend in past trades. When traders have private information on only a single dimension of uncertainty (the effect of a shock to the asset value), price adjustments prevent herd behavior. Herding arises when there are two dimensions of uncertainty (the existence and effect of a shock), but it need not distort prices because the market discounts the informativeness of trades during herding. With a third dimension of uncertainty (the quality of traders' information), herd behavior can lead to a significant, short-run mispricing. (JEL G12, G14, D83, D84)

In standard models of general equilibrium, the simultaneous execution of a large number of trades produces efficient outcomes, presuming that the Walrasian auctioneer has set prices correctly. As a challenge to these results and the associated view that decentralized markets tend to be efficient, an explosion of papers in the last several years argue that imitative or herd-like behavior can impede the flow of information in an economy when consumers act sequentially rather than concurrently. [Abhijit Banerjee (1992), Sushil Bikhchandani et al. (1992), Christophe Chamley and Douglas Gale (1992), Andrew Caplin and John Leahy (1993, 1994), and Jeremy Bulow and Paul Klemperer (1994)]. With sequential actions, the earliest decisions can have a disproportionate effect over long-run outcomes in the economy. A slight preponderance of public information is sufficient to induce all agents to follow the lead of the market, completely ignoring their private information. Bikhchandani et al. (BHW) describe that situation as an “informational cascade.” In BHW and Banerjee’s models, an informa-

tional cascade occurs in finite time with probability 1. That is, social learning completely breaks down as all consumers from some time forward make the same choice and reveal no new information. Because that choice is wrong with strictly positive probability, the equilibrium of these sequential market games is inefficient, even in the long run.

The herding literature recalls a once-prominent view of asset markets as driven by “animal spirits,” where investors behave like imitative lemmings. While the rational actor approach has largely driven this view from mainstream research in financial economics, it is far from gone. Both market participants and financial economists reportedly still believe that imitative behavior is widespread in financial markets (Andrea Devenow and Ivo Welch, 1996). This has led some researchers to assert that market participants engage in nonrational herd behavior (e.g., Andrei Shleifer and Lawrence H. Summers [1990]; Alan Kirman [1993]).

We investigate the relationship between rational herd behavior and asset prices. Past work on rational herding is not well suited to address this relationship because, in almost all cases, herding models fix the price for taking an action *ex ante*, retaining that price inflexibly under all circumstances.¹ We address the fol-

* Avery: Kennedy School of Government, Harvard University, Cambridge, MA 02138; Zemsky: INSEAD, Boulevard de Constance, 77305 Fontainebleau, Cedex, France. We are grateful for comments and advice from Anat Admati, Darrell Duffie, Glenn Ellison, Drew Fudenberg, Jim Hines, Allan Kleidon, Paul Pfleiderer, Lones Smith, Peter Sorensen, Xavier Vives, Richard Zeckhauser, and especially Paul Milgrom.

¹ An exception is Bulow and Klemperer’s model, but it still fixes the price after each purchase for a sufficient period to produce herding.

lowing questions: Can there be informational cascades in financial markets? Can herd behavior lead to the long-run mispricing of assets? Does it produce bubbles and crashes? Might it offer an explanation for excess volatility? We begin our analysis with an example which motivates a final question.

I. A Simple Example

Our model retains the basic features of the simplest model considered by BHW (p. 996), with the notable addition of a price mechanism. It is useful to review that model and to consider what happens when prices are allowed to vary over time in response to agents' actions.² In BHW, agents face a choice of whether or not to adopt a new technology, and the cost of adoption is fixed at $c = 1/2$. The value of the new technology, denoted V , is either 1 or 0. Each agent gets an independent, imperfect signal about V , denoted $x \in \{0, 1\}$, where $P(x = V) = p > 1/2$. Agents act sequentially and observe H_t , the history of actions up until time t . Let $\pi'_1 = P(V = 1 | H_t)$. The choice made by an agent depends on whether the expected value of adopting is greater than c . Consider the expected value of an agent with bad news (the value for an agent with good news is similar):

$$\begin{aligned} V'(x = 0) &= E[V | x = 0, H_t] \\ &= \frac{1 - p}{(1 - p)\pi'_1 + p(1 - \pi'_1)} \pi'_1. \end{aligned}$$

Assuming all prior agents have acted in accord with their signal, π'_1 increases with the difference between the number of prior agents who adopted and those who did not. Indeed, whenever there are two more adopters than non-adopters, it is the case that $V'(x = 1) > V'(x = 0) > 1/2$. Then agents at time t adopt regardless of their signal and an informational cascade begins.

Now suppose that the agents are traders in a financial market and that their choice is

whether to buy or sell a unit of an asset where the true value of the asset is given by V . Further, suppose that the financial market is informationally efficient in that the cost of a unit of the asset reflects all *publicly* available information:

$$\hat{c} = E[V | H_t] = P(V = 1 | H_t) = \pi'_1.$$

The key observation from this simple exercise is that

$$V'(x = 1) > \hat{c} > V'(x = 0).$$

The asset price adjusts precisely so that there is no herding and agents always trade in accord with their signal! Reflecting on Adam Smith's invisible hand, it is not too surprising that an arbitrary fixed price leads to herd behavior and the persistence of inefficient decisions in an economy. We conclude that whether or not herd behavior affects asset prices, asset prices can certainly affect herd behavior. In this example, they completely eliminate it. Given the reported prevalence of herd behavior in financial markets, this raises the important question of whether herd behavior is consistent with a market composed of rational traders.

II. Overview of the Paper

In Section III we describe a general model and define terms. Of particular importance, we define herd behavior as a trade by an informed agent which follows the trend in past trades even though that trend is counter to his initial information about the asset value.

In Section IV, we show that there are limits to the distortions that can arise in a financial market where informed traders are rational actors and prices incorporate all publicly available information. We show that informational cascades are impossible: at any point in time there is always the possibility that new information reaches the market. Consistent with this steady flow of information, prices always converge to the true value. Hence, herd behavior can cause no long-run mispricing of assets. We show that the *ex ante* expected volatility in prices is determined by fundamentals, which means that herd behavior can not be the source of excess volatility. Finally,

² We use notation consistent with the rest of this paper. In comparing our results to those in BHW, our reference is to the simple model described here.

we generalize the example from Section I by stating a monotonicity condition on private signals such that herding is impossible. With monotonic signals, there is only a single dimension of uncertainty confronting the market, which we term *value uncertainty*.

In Section V we exhibit a plausible information structure in which herding does occur, by adding *event uncertainty* to value uncertainty. With event uncertainty, the market is uncertain as to whether the value of the asset has changed from its initial expected value. We show that any amount of event uncertainty produces the possibility of herd behavior. As event uncertainty becomes extreme (i.e., the probability that the asset has not changed value goes to 1), there is an arbitrarily long period of herd behavior when the asset value changes. This herd behavior is similar to the informational cascade of BHW in that the market does not learn about whether the asset value is high or low as all informed traders either buy or sell. Surprisingly, this extreme herd behavior has little effect on asset prices. We show that the movement in the asset price is bounded and that this bound can be small. Finally, we argue that herd behavior is not clearly at odds with optimal social learning in financial markets.

Given the above results, one might expect that we find no connection between herd behavior and market crashes. However, this is not the case. In Section VI we investigate the combination of event uncertainty with what we term *composition uncertainty*, which means that there is uncertainty as to the average accuracy of traders' information. We are then able to identify certain (highly unlikely) states of the world in which herd behavior can lead to a price bubble and crash. In these states, market participants have a mistaken, but rational, belief that most traders possess very accurate information. Then, market participants have trouble differentiating between a market composed of well-informed traders and one with poorly informed traders who are herding: in each case, there is a preponderance of activity on one side of the market. The resulting confusion allows uninformative herd behavior to have dramatic effects on prices. Our theory of price bubbles resembles the explanation advanced by Sanford J. Grossman

(1988) and Charles J. Jacklin et al. (1992) for the stock market crash of 1987: traders underestimated the prevalence of noninformative computer-based insurance trading.

Based on these results, we conclude that despite the significant constraints imposed by a rational financial market, herd behavior is robust to the operation of the price mechanism. In particular, as the number of dimensions of uncertainty with which the price mechanism must contend increase, herding becomes prevalent and extreme effects based on herd behavior occur in identifiable (but unlikely) states of the world.

In Section VII we consider the converse of herding, "contrarian behavior," where agents ignore their private information about value uncertainty to trade *against* the trend in past trades. We show that composition uncertainty can give rise to such behavior. The existence of herd and contrarian behavior rationalizes the observed practice of price charting. Section VIII concludes.

III. The General Model

A. Description

We begin by specifying a general model; we will add further assumptions in later sections. The market is for a single asset with true value V , which is restricted to be in $[0, 1]$. Prices are set by a competitive market maker who interacts with an infinite sequence of individuals chosen from a continuum of traders. Each trader is risk neutral and has the option to buy or sell one unit of stock or to refrain from trading. The sequence of traders is indexed by $t = 0, 1, 2, \dots$. We denote by H_t the publicly observable history of trades up until time t .

There are two broad classes of traders. Informed traders receive private information and maximize expected profit at the market maker's expense, while noise traders act for exogenous motives and without regard for expected profit.³ Let $\mu < 1$ be the probability of an informed trader arriving in any given pe-

³ Without the presence of noise traders, the no-trade theorem of Paul Milgrom and Nancy Stokey (1982) applies and the market breaks down.

riod; $1 - \mu$ is the probability of a noise trader arriving. For convenience, we assume that noise traders buy, sell, and do not trade with equal probability $\gamma = (1 - \mu)/3$.

Informed traders receive private information $x_\theta \in [0, 1]$, where x_θ is drawn from the distribution $f_\theta(x_\theta|V)$ and θ is a trader's type.⁴ We assume a finite number of possible types and we denote the probability that a trader of type θ arrives by $\mu_\theta > 0$. The expected value of an informed trader is denoted $V'_\theta(x) = E[V|H_t, x_\theta = x]$. The market maker's expected value for the asset given public information is denoted $V'_m = E[V|H_t]$, which we shall sometimes refer to as the price.⁵

We assume that there is always a minimal amount of "useful" information in the market. That is, as long as past trading does not identify the value perfectly, then there is strictly positive probability that some trader has an assessed value that differs from the market maker's (by a nontrivial amount). More precisely, we assume that if there does not exist a v such that $P(V = v|H_t) = 1$, then there exists at least one θ and set of signal realizations $R \subset [0, 1]$ with $P(x_\theta \in R|H_t) > 0$ such that $V'_\theta(x_\theta) \neq V'_m$ for $x_\theta \in R$. Moreover, if $|V'_m - V| = \delta > 0$ then for some $\epsilon(\delta) > 0$, $|V'_\theta(x_\theta) - V'_m| > \epsilon(\delta)$.

The market maker allows for adverse selection by setting a (bid-ask) spread between the prices at which he will sell and buy a unit of stock. Perfect competition among market makers restricts the market maker to zero profits at both the bid and ask prices. That is, the trader who arrives in period t faces a bid, B^t , and an ask, A^t , which satisfy:

$$B^t = E[V|h_t = S, H_t],$$

$$A^t = E[V|h_t = B, H_t],$$

⁴ Thus, a trader potentially has *two* pieces of private information to trade on—the value of $x_\theta \in [0, 1]$ and his type θ . A trader's type constitutes private information if there is uncertainty about the composition of the market. See Section VI for details.

⁵ We do this when we want to abstract from the existence of the bid-ask spread in interpreting our results.

where h_t is the action taken by the trader who arrives in period t , with $h_t = B$ indicating a buy, $h_t = S$ indicating a sell, and $h_t = NT$ indicating no trade. Finally, we define the market maker's assessed distribution for the possible values as $\pi'_v = P(V = v|H_t)$. By Bayes' theorem, these priors respond to trade as follows:

$$(1) \quad \pi'_v{}^{t+1} = \pi'_v{}^t \frac{P(h_t|V = v, H_t)}{P(h_t, H_t)},$$

where $P(h_t, H_t) = \sum_v \pi'_v P(h_t|V = v, H_t)$.

Our model is a special case of the model developed by Lawrence Glosten and Milgrom (1985) with the notable simplification that our noise traders have completely inelastic demand. Because our noise traders are willing to absorb any amount of losses, the market never breaks down due to adverse selection and zero profit equilibrium prices always exist.

PROPOSITION 1: *In each period t there exist unique bid and ask prices which satisfy $B^t \leq V'_m \leq A^t$. V'_m and π'_v are martingales with respect to H_t .*

PROOF:

See Appendix.

The market maker accounts for the information which is contained in buy and sell orders in setting prices. Thus, $A^t \geq V'_m$ and $B^t \leq V'_m$. V'_m and π'_v are expectations based on all of the information contained in the prior history of trade, H_t . Therefore, they are martingales with respect to H_t ; if this were not the case, then the market maker's assessment of V'_m and π'_v would be systematically mistaken in a manner which should be predictable to him.

B. The Definition of Herd Behavior

We differentiate between an informational cascade and herd behavior. In the example of Section I, herd behavior always implies an informational cascade. With the simple information structure used there, no information reaches the market when traders with bad signals ($x = 0$) and those with good signals ($x = 1$) are taking the same action. However, in a

more general model, imitative behavior need not imply an informational cascade.

Definition 1: An *informational cascade* occurs in period t when

$$P(h_t|V, H_t) = P(h_t|H_t) \quad \forall V, h_t.$$

In an informational cascade, no new information reaches the market because the distribution over the observable actions is independent of the state of the world. In particular, this happens when the actions of all informed traders are independent of their private information, such as when they are all buying (in all states of the world).

Definition 2: A trader with private information x_θ engages in *herd behavior* at time t if he buys when $V_\theta^0(x_\theta) < V_m^0 < V_m^t$ or if he sells when $V_\theta^0(x_\theta) > V_m^0 > V_m^t$; and buying (or selling) is strictly preferred to other actions.

Herd behavior by a trader satisfies three properties, which we discuss for the case of herd buying. First, it must be that initially (before the start of trade) a trader's information leads him to be pessimistic about the value of the asset so that he is inclined to sell: $V_\theta^0(x_\theta) < V_m^0$. Second, the history of trading must be positive: $V_m^0 < V_m^t$. Finally, the trader must want to buy given this positive history and his signal, which implies that $V_m^t \leq A^t < V_\theta^t(x_\theta)$. These three properties demonstrate the extreme nature of herd behavior. Initially, the trader's signal constitutes negative information, causing him to reduce his assessment of the asset's value. Yet, after observing the trading history, the signal constitutes positive information, causing him to increase his assessment of the asset's value from V_m^t .

In our definition, herd behavior occurs when agents imitate the prior actions (buying or selling) of others. An alternative approach is to define herding as a socially inefficient reliance on public information (see Xavier Vives, 1997).⁶ In contrast, we start with a behavioral

⁶ In settings where agents learn from the actions of others, an informational externality naturally arises in that future agents benefit when earlier agents take actions that reveal their private information. Hence the connection be-

between efficiency and herd behavior, which can obscure private information. There are two drawbacks to using an efficiency based definition of herd behavior for studying financial markets. First, it requires a welfare benchmark, which is generally lacking in asymmetric information models of asset markets. Second, traders can place a very high weight on public information without exhibiting the sort of strongly imitative behavior studied here (see the work of Vives, 1995, 1997).

IV. Bounds on the Effect of Herd Behavior

Asset prices have a profound effect on herd behavior. As suggested by our earlier example, the price mechanism eliminates the possibility of informational cascades.

PROPOSITION 2: An *informational cascade never occurs in market equilibrium*.⁷

PROOF:

See Appendix.

Our assumption of minimal useful information implies that there is always private information in the economy. As long as private information exists, some traders must base their trading strategy on that information, but this assures that the distribution over observed actions is not independent of the state. Hence, an informational cascade is impossible. Like several of our results, Proposition 2 relies on a basic intuition about our model: informed trade is driven by information asymmetries between traders and the market maker.

Proposition 2 requires limited frictions in the market. Otherwise, trade and the flow of information can stop. In Ho Lee (1995) shows that informational cascades arise if there are transaction costs. Over time, the expected profit of informed traders declines to zero as the asset price becomes more accurate. If there are transaction costs to trading, informed traders will (almost surely) stop trading at some point. Then, no new information reaches the market. Similarly, in the original Glosten and

⁷ We are very grateful to an anonymous referee for suggesting this result.

Milgrom paper, informational cascades arise if the market breaks down due to adverse selection.

Consider the following restriction on the private information in the economy.

Definition 3: A signal x_θ is *monotonic* if there exists a function $v(x_\theta)$ such that $V'_\theta(x_\theta)$ is always (weakly) between $v(x_\theta)$ and V'_m for all trading histories H_t .

Monotonic signals are particularly well behaved. Given any public information, they always move a trader's expected value towards some fixed valuation, $v(x_\theta)$. Monotonic signals are pervasive in the literature on asymmetric information in financial markets. For instance, the signals in the example of Section I, which are often used in Glosten-Milgrom style models, are monotonic because $V'(x) \in [x, V'_m]$. In addition, noisy rational expectations models (e.g. Grossman and Joseph E. Stiglitz, 1980) *require* monotonic signals for tractability. We now show that it is the ubiquitous assumption of monotonic signals that explains the absence of herd behavior in the received literature on the microstructure of financial markets.

PROPOSITION 3: *A trader with a monotonic signal never engages in herd behavior.*

PROOF:

Suppose a trader with a monotonic signal x_θ engages in herd buying at time t . Then $V'_\theta(x_\theta) > A^t \geq V'_m$. Since the signal is monotonic, this implies that $v(x_\theta) > V'_m$. But then $V'_\theta(x_\theta) > V'_m$ and the trader was not originally pessimistic. This is a contradiction. Similarly, herd selling never occurs.

With monotonic signals, a trader who wants to buy when the price has risen must also want to buy initially, which assures that any buying is not herding.⁸ If we abstract from the exist-

⁸ Vives (1995) develops a dynamic noisy rational expectations model which complements our analysis. Consistent with Proposition 3 and his use of monotonic signals, traders in Vives' model never engage in herd be-

havior as defined here. They always buy if the value of their signal is above the price and sell if it is below. The amount that they buy or sell does change over time as public information accumulates.

tence of a bid-ask spread (as when μ is small), agents with monotonic signals have particularly simple trading strategies. They buy if $v(x_\theta)$ is above the price, V'_m , and sell if it is below that. Then, traders need not concern themselves with the trading history at all! This rules out herd behavior, since it leaves no room for the trend in the trading history to influence trading. When traders have monotonic signals, we say that there is only a single dimension of uncertainty in the market. Our motivation is that a scalar, V'_m , can summarize for traders all the information they need to extract from the trading history. We label this single dimension of uncertainty as *value uncertainty*, as it relates directly to the underlying value of the asset.⁹

In Section V we show that there exist plausible nonmonotonic signals which produce herd behavior. However, we now show that the effect of herd behavior on prices must be limited. The impossibility of informational cascades implies that each period of trade reveals some information even if there is herd behavior. Since there is a continual flow of information, it is natural that the trading price must converge to the true asset value.

PROPOSITION 4: *The bid and ask prices converge almost surely to the true value V .*

PROOF:

See Appendix.

Glosten and Milgrom (1985) show that the bid and the ask prices must converge together as long as the market does not break down. Hence, all private information becomes public

⁹ We do not make precise our notion of "dimensions" of uncertainty, but leave it as an intuitive construct that we find useful for interpreting our results. We shall speak of the asset price as having a single dimension in our model, even though technically there is both a bid and an ask price. We think of multidimensional prices as arising when there are derivative securities (such as options) that are traded.

over time. Convergence is then a direct consequence of our assumption that a nontrivial amount of private information always exists so long as the true value is not yet identified by the trading history.

While this convergence result is not new, it has significant implications for the applicability of results from the recent herding literature to financial markets. Price convergence directly rules out the sort of long-run inefficiencies found in earlier herding papers. Further, when coupled with the martingale property of prices, convergence provides a bound on the volatility of prices. We denote the change in the market maker's expectations from one period to the next as $\Delta V_m^t = V_m^t - V_m^{t-1}$.

COROLLARY 1: *The variance of price paths is bounded as follows.*¹⁰

$$\sum_{t=1}^{t=T} \text{Var}(\Delta V_m^t) \leq \text{Var}(V).$$

Additionally, for a fixed t ,

$$\text{Var}(V - V_m^t) = \text{Var}(V) - \text{Var}(V_m^t).$$

PROOF:

See Appendix.

The first part of Corollary 1 states that the expected volatility is bounded by the fundamental uncertainty over V . Hence, it is not possible to explain volatility in excess of fundamentals in our general model, whether or not there is herd behavior. In addition, as time passes, the "remaining variance" in the price change process diminishes, so that V_m^t must be more and more accurate over time, as implied by the second part of the corollary. That is a natural property with an important implication: any set of volatile price paths must either converge quickly to the true value or (as the next corollary emphasizes) they can only occur with small probability.

COROLLARY 2: *Consider some $a < V_m^t$. Then $P(V < a | H_t) \leq (1 - V_m^t)/(1 - a)$.*

¹⁰ We are grateful to Paul Milgrom for suggesting this result.

PROOF:

See Appendix.

This result implies that high prices occur only rarely when asset values are low. Consider a market where $V \in \{0, 1\}$ as in BHW. Corollary 2 implies that $P(V = 0) = 1 - V_m^t$ (i.e., as the price V_m^t goes to one, the probability that $V = 0$ goes to zero). This result limits the probability of a price bubble, which is a situation where the asset value moves far away from its true value. In general, there is an inverse relationship between the magnitude of a price bubble and the probability with which it occurs. In particular, extreme price bubbles, where the asset incorrectly attains its maximum possible value, are zero probability events.

V. Event Uncertainty and Herd Behavior

A. Existence of Herd Behavior

Proposition 3 poses a puzzle. How do we reconcile the reported prevalence of herd behavior with its absence in a rational financial market with monotonic signals? A closely related puzzle is the existence of price "charting," where traders use detailed charts of price histories in their trading strategies (David Brown and Robert Jennings, 1989). Charting is puzzling because the trading history plays at most a limited role in a trader's strategy when he has a monotonic signal. For any history, the set of potential buyers (traders who are more optimistic than the market maker) is given by the condition $v(x_\theta) > A^t$, while the set of potential sellers is given by the condition $v(x_\theta) < B^t$.¹¹

While it is certainly not difficult to specify nonmonotonic signals, the more interesting question is whether such signals are likely to be common in financial markets. Consider

¹¹ Hence, the only role of the trading history is in helping traders to assess whether their private information is sufficiently strong to justify trading given the bid-ask spread. We do not find this weak history dependence to be a satisfactory theory of price charting. We seek a rationalization based on strong history dependence, where the trading history drives a trader from buying to selling, as occurs under our definitions of herd and contrarian behavior.

then, that many shocks to an asset's value are not publicly known, at least initially. For example, a trader might learn from a contact who works at a company that there will be a change in management, that a new product has been developed, or that a merger is being considered—all before a public announcement. Then, the trader has private information about two “dimensions” of uncertainty. In addition to information related to value uncertainty—is it a good or a bad merger?—the trader has private information that there has been a shock to the underlying value of the asset. We follow the finance literature and refer to this second dimension as *event uncertainty* (David Easley and Maureen O'Hara, 1992). We offer the following formalization of event uncertainty.

Definition 4: There is *event uncertainty* when $1 > P(V = V_m^0) > 0$.¹²

We now extend the simple BHW information structure used in our introductory example to incorporate event uncertainty. Traders are informed if there is an information event (i.e., $V \neq V_m^0$) and if there is, then traders have signals as in BHW. Formally, *information structure I* (IS I) is defined as follows. The true value of the asset is $V \in \{0, \frac{1}{2}, 1\}$ with initial priors satisfying $\pi_{1/2}^0 > 0$ and $\pi_1^0 = \pi_0^0 > 0$. Then, $V_m^0 = \frac{1}{2}$ and there is event uncertainty. There is a single type of informed trader with signal x , where

$$P(x = \frac{1}{2} | V) = \begin{cases} 1 & \text{if } V = \frac{1}{2}, \\ 0 & \text{if } V \neq \frac{1}{2}, \end{cases}$$

$$P(x = 1 | V) = \begin{cases} p & \text{if } V = 1, \\ 1 - p & \text{if } V = 0, \end{cases}$$

$$P(x = 0 | V) = \begin{cases} p & \text{if } V = 0, \\ 1 - p & \text{if } V = 1, \end{cases}$$

$1 \geq p > \frac{1}{2}$. If there were no event uncertainty

¹² The event uncertainty that we study satisfies the additional property that informed traders know whether or not an information event has occurred. That is, $P(V = V_m^0 | x_t) \in \{0, 1\}$.

(i.e., $\pi_{1/2}^0 = 0$), then the above signals are monotonic. The addition of event uncertainty in this way makes signals nonmonotonic and herd behavior possible.¹³

PROPOSITION 5: *Under IS I, price paths with herd behavior occur with positive probability for $p < 1$. They do not occur for $p = 1$. Herd behavior is misdirected with positive probability.*¹⁴

PROOF:

See Appendix.

In BHW, a preponderance of one action chosen by earlier agents leads others to believe that the action is a good one, regardless of their own private information. Similarly, a sufficient excess of buys over sells in our model leads a trader to believe that the asset value is more likely to have gone up rather than down, regardless of his signal. However, with informationally efficient prices, rational individuals only act based on information asymmetries between themselves and the market maker. The history of trades can only be the source of asymmetric information if it is interpreted differently by the market maker and the informed traders. With only value uncertainty, there is only a single interpretation of the history of trade and hence herd behavior is impossible.

With the addition of event uncertainty, informed traders know that an information event has occurred, while the market maker does not. This information asymmetry gives the traders an advantage in interpreting the history of trades. They are quicker to adjust their valuation to the trend in past trades than the

¹³ It is possible to have event uncertainty while preserving the monotonicity of signals. For example, if $V \in \{0, \frac{1}{2}, 1\}$, $P(x = V) = p$, $P(x = 1 | V = \frac{1}{2}) = P(x = 0 | V = \frac{1}{2}) = (1 - p)/2$, $P(x = \frac{1}{2} | V = 1) = P(x = \frac{1}{2} | V = 0) = q$, and $p > q \geq (1 - p)/2$, then x is monotonic. Note that such a signal precludes informed traders from knowing with certainty whether or not an information event has occurred, which we find to be a natural feature of event uncertainty.

¹⁴ Easley and O'Hara (1992) study IS I for the special case where $p = 1$. Their focus is on how the market maker learns that an information event has occurred.

market maker, who must consider the possibility that there has been no change in the underlying value of the asset and the trend is due to noise traders. Thus, event uncertainty dulls price adjustment in the short run.

PROPOSITION 6: Consider IS I and some trading history H_t that results in priors $\pi'_v = P(V = v|H_t)$. For $\pi'_1 = \pi'_0$, there is no herd behavior. For $\pi'_1 \neq \pi'_0$, there exists a critical value for the precision of traders' signals $\bar{p}(\mu, \pi'_0, \pi'_1)$ such that traders engage in herd behavior in period t if and only if $p < \bar{p}$. This \bar{p} decreases with μ and increases with $\pi'_{1/2}$ (holding π'_1/π'_0 constant). If $\pi'_1 > \pi'_0$, then any herd behavior involves buying and \bar{p} increases with π'_1/π'_0 (holding $\pi'_{1/2}$ constant). If $\pi'_1 < \pi'_0$, then any herd behavior involves selling and \bar{p} increases with π'_0/π'_1 (holding $\pi'_{1/2}$ constant).

PROOF:

See Appendix.

Proposition 6 identifies the three forces that produce herd behavior. First, herd behavior results when the weight of information in the history of trade overwhelms an individual's private information about value uncertainty. A reduction in p reduces the information contained in a private signal about value uncertainty, while an increase in $|\pi'_1 - \pi'_0|$ increases the amount of information contained in the history. Either change makes it easier for the trading history to overwhelm the information about value uncertainty in private signals and thus makes it easier for herding to arise. Second, herding occurs when prices become sufficiently unresponsive to the trading history. As the probability of an information event decreases (i.e., $\pi'_{1/2}$ increases), prices respond less to the trading history and thus more of the information in the trading history is private.¹⁵ Third, herding requires that the bid-ask spread not deter potential herders from

trading. A decrease in μ reduces adverse selection and leads to a tighter bid-ask spread.

B. Existence of Pronounced Herd Behavior

Proposition 5 shows that herding is possible for any $p < 1$ and $\pi_{1/2} > 0$. We now show that as event uncertainty becomes extreme, herd behavior becomes pronounced, resembling the cascades of BHW.

PROPOSITION 7: Consider IS I with $p < 1$ and suppose that an information event occurs. In the limit as the probability of an information event becomes arbitrarily small (i.e., $\pi_{1/2}^0 \rightarrow 1$), the probability that there is some herd behavior in the trading history goes to 1. Moreover, the trading history almost surely takes the following form: (i) a finite, initial period of trading during which herd behavior does not occur, (ii) an arbitrarily long period of herd behavior of one type (i.e., always buy or always sell). This herd behavior is in the wrong direction with a strictly positive probability,

$$\lambda \in \left[\frac{(1-p)^2}{p^2 + (1-p)^2}, 1-p \right].$$

In the limit as $\mu \rightarrow 0$, the probability of herd behavior in the wrong direction goes to $1-p$.

PROOF:

See Appendix.

When an information event is very surprising (i.e., $\pi_{1/2}$ close to 1), the market maker discounts almost completely the informativeness of trading. The price remains fixed at the initial expected asset value of $1/2$ for an arbitrarily long period of time. With a fixed price, our model almost recreates the BHW model and hence it is not surprising that cascade-like behavior arises. The only difference with BHW is the existence of noise traders.¹⁶

¹⁵ Note that the information gleaned by informed traders from observing the trading history is fixed by π'_1/π'_0 regardless of the value of $\pi'_{1/2}$ because they know whether or not an information event has occurred.

¹⁶ The main effect of noise traders is to increase somewhat the probability of herding in the wrong direction. In BHW it takes two more adopters than nonadopters to start an informational cascade. Here the imbalance between

There is an important distinction between the herding of Proposition 7 and an informational cascade, which cannot occur in equilibrium. In an informational cascade, no new information reaches the market. Under the conditions of Proposition 7, behavior resembles a cascade when there has been an information event. For a very long interval, all informed traders act as buyers (or sellers). However, new information still reaches the market. Had there been no information event, informed traders would not be trading and the volume of trade and the imbalance in trade would both be less. Hence, the market maker is learning that there was an information event. After a sufficient period of time, the market maker learns enough about event uncertainty that prices adjust, which ends herding.

The addition of event uncertainty makes herd behavior possible and even extreme. However, herd behavior in IS I—no matter how extreme—does not distort the asset price. Herding keeps information about the new asset value from entering the market and rational actors (including the market maker) account for this.

PROPOSITION 8: *Consider IS I. During any interval of trading in which there is herd behavior, the movement in the asset price is less than*

$$\Delta = \frac{3\mu(p - \frac{1}{2})}{2 + \mu}.$$

In the limit as either $\mu \rightarrow 0$ or $p \rightarrow \frac{1}{2}$, $\Delta = 0$.

buys and sells necessary to start herd behavior depends on the probability of a noise trader. The greater the probability of noise traders, the greater the imbalance required. However, the key statistic for a cascade—the probability that it is in the wrong direction—is very similar. In BHW, the probability of a wrong direction cascade is the probability that two agents have the wrong signal $(1 - p)^2 / (p^2 + (1 - p)^2)$. Here λ is between the probability that two traders have the wrong signal and the probability that just one trader has the wrong signal.

PROOF:

See Appendix.

Herding is triggered when the information contained in a trade pushes the value of an informed trader past the bid or the ask. However, during the entire interval of herding that follows, the valuations of traders are fixed because they realize that no new information about value uncertainty is being revealed. Therefore, as soon as the price moves by more than the information contained in the last trade prior to herding, herd behavior stops. In some circumstances, the information contained in a single trade is small. In particular, for a small probability μ of an informed trader arriving or for low-precision signals, each trade conveys little information and the price movement during herding is small. In conclusion, no one is fooled by herd behavior in IS I. Any price rise during periods with herding results only from information about the new value which was contained in trading prior to herding. All that is learned during herding is that an information event has occurred.

C. Information Aggregation and Herd Behavior

We now explore the effect of herd behavior on the aggregation of dispersed private information in a financial market. A useful benchmark is Vives (1997), who studies the effect of imitative behavior on the efficiency of information revelation in a fixed-price setting. He considers an economy where a sequence of agents make the same, irreversible decision and have monotonic private signals about the optimal decision. Vives proposes as a welfare benchmark a team solution which assigns agents decision rules that minimize the mean decision error. He finds that in the decentralized economy, agents put too much weight on the decisions of others relative to the welfare benchmark. Agents do not internalize the negative externality on later agents caused by their imitative behavior, which obscures their private information.

Vives' approach does not transfer well to financial markets. In an asset market, the natural team solution would be to maximize the profits of the informed traders. However,

trading profits based on asymmetric information have no clear link to social welfare. What would seem more important is the information revealed in the trading history—especially the information reflected in the asset price. It is such information which is likely to affect decision-making in the real economy. With event uncertainty, there are potentially two pieces of information to be revealed: first, whether an information event has occurred, and then if it has, whether it is a positive or negative event.

PROPOSITION 9: *Consider IS I. A period of herding reveals more information about the existence of an information event than a period in which agents trade based on their information about value uncertainty. Specifically, herding is more effective at decreasing $E[\pi'_{1/2} | V \neq 1/2, H_t]$.*¹⁷

PROOF:

See Appendix.

While herding is costly in that it obscures information about value uncertainty, it has a benefit. It is more effective at revealing the existence of an information event. By focusing the trade of the informed on a single action when there is an event, herding reduces the effect of noise trading. For example, with herd buying, a sell order must come from a noise trader and a buy order becomes a strong signal that there has been an information event. We now consider how well a decentralized market trades off these costs and benefits of herd behavior.

PROPOSITION 10: *Consider IS I and suppose that an information event occurs. In the limit as $\mu \rightarrow 0$, the choice of informed traders between herding and trading based on their information about value uncertainty minimizes*

¹⁷ Note that since $E[\pi'_{1/2} | H_t] = P(V = 1/2)E[\pi'_{1/2} | V = 1/2, H_t] + P(V \neq 1/2)E[\pi'_{1/2} | V \neq 1/2, H_t]$, minimizing $E[\pi'_{1/2} | V \neq 1/2, H_t]$ is equivalent to maximizing $E[\pi'_{1/2} | V = 1/2, H_t]$. Hence Proposition 9 also establishes that herding is more effective at revealing that an information event has not occurred, should this be the case.

the deviation of the asset price from its new value. That is, the trading strategy of the informed approaches the trading strategy which minimizes $E[|V_m^{t+1} - V| | V \neq 1/2, H_t]$.

PROOF:

See Appendix.

The decentralized economy can come arbitrarily close to maximizing the movement of the asset price towards its new value.¹⁸ We conclude that the incentives of self-interested traders with private information do not diverge from social interests as much as the fixed-price herding literature suggests. We now reconsider the incentives for informed traders to herd, as identified in Proposition 6, to see why they are consistent with maximal information revelation (via the price). Traders herd when they have sufficiently low-precision signals, but then the information lost from herd behavior is small. They do not herd when $\pi'_1 = \pi'_0$, but then information that an event has occurred has no impact on prices since there is no public information as to whether the event is good or bad. Conversely, traders herd when there is already good information about the new asset value (i.e., $|\pi'_1 - \pi'_0|$ large), but this is when the additional information from more trading based on value uncertainty is small. They herd when there is little awareness of an information event ($\pi'_{1/2}$ large), but then price does not respond strongly to new information about value uncertainty.

Proposition 10 is a fitting end to Sections IV and V, which sing the praises of informationally efficient financial markets populated by rational traders.

VI. Herd Behavior and Price Bubbles

Herd behavior in a financial market is of particular interest because of the possibility that it might offer an explanation for price bubbles and excess volatility. Because price is a

¹⁸ However, Proposition 10 only addresses information revelation in one period of trade. An analysis of optimal information revelation over longer time horizons is beyond the scope of this paper.

martingale that converges to the true value, it is not possible to have volatility in excess of fundamentals in our general model (Corollary 1), nor can price bubbles be both likely to occur and extreme (Corollary 2). However, it still may be possible to identify (unlikely) circumstances which consistently produce highly volatile price paths. Here we investigate whether herd behavior can produce an unsustainable run-up in price that results in a crash. In the previous section, we saw that herd behavior need not distort prices at all. In IS I, herding produces an imbalance in trading, but market participants understand that this is due to herd behavior and hence prices and valuations do not respond. We now consider whether herding is always likely to be so transparent.

A. Uncertainty About the Composition of the Market

When a trader learns of an information event, his assessment of its impact on the asset value is sometimes precise and sometimes imprecise. For example, a trader may or may not be confident in his ability to predict the effect on profits of a change in a firm's product mix or of a merger decision, depending on whether he has complementary pieces of information, such as detailed information about the merger partner. For the market as a whole, some information events will have a high proportion of well-informed traders, while others will have only a few. If the market is uncertain *ex ante* about the proportion of different types of traders, we have a third dimension of uncertainty.

Definition 5: There is *composition uncertainty* when the probability of traders of different types, μ_θ , is not common knowledge.

Composition uncertainty complicates learning for market participants, especially in the presence of herd behavior. Note that trading patterns in a market with many poorly informed traders and herding mimic the trading patterns in a market with well-informed traders. In a poorly informed market, a sequence of buy orders is natural because of herding. In a well-informed market, a sequence of buy orders

is also natural because the agents tend to have the same (very informative) private signal. Without knowledge of the composition of the market, it can then become difficult to distinguish whether a sequence of buy orders reveals a large amount of information about value uncertainty (because the market is well informed) or almost none at all (because the market is poorly informed and informed traders are herding). We specify a new information structure in order to show that this confusion can lead to extreme short-run price effects due to herding.

Information structure II (IS II) adds composition uncertainty to IS I. The true value of the asset is still $V \in \{0, 1/2, 1\}$. The signals of informed traders take the same form, but now there are two types of trader with $\theta \in \{H, L\}$. The difference between the two types of traders is the precision of their information when there is an information event. In particular,

$$P(x_\theta = 1 | V) = \begin{cases} p_\theta & \text{if } V = 1, \\ 1 - p_\theta & \text{if } V = 0, \end{cases}$$

$$P(x_\theta = 0 | V) = \begin{cases} p_\theta & \text{if } V = 0, \\ 1 - p_\theta & \text{if } V = 1, \end{cases}$$

and $p_H = 1$ while $1 > p_L > 1/2$. Hence *H* types are perfectly informed (i.e., $E[V | x_H] = V$), while *L* types have noisy signals when the asset value changes.

The level of information in the market is indexed by $I \in \{W, P\}$. The difference between a well-informed market ($I = W$) and a poorly informed market ($I = P$) is in the proportions of each type of informed trader. Let μ'_θ be the probability of a type θ trader in a type *I* market. For example, μ_H^W is the probability of a high-precision trader in a well-informed market. We assume that there is a fixed probability of an informed trader ($\mu_H^I + \mu_L^I = \mu$) and that there are more *H* types in a well-informed market than in a poorly informed market ($\mu_H^W > \mu_H^P$). The state of the world is given by the combination of the asset's underlying value and the amount of information available: (V, I) . The market-maker's assessed probabilities conditional on

the trading history prior to time t are denoted $\pi_{V,I}^t$.

B. An Example of a Price Bubble

We now show how price bubbles can arise in IS II by means of a simulation.¹⁹ The initial priors for the simulation are $\pi_{1/2}^0 = 0.9999$, $\pi_{v,W}^0/\pi_{v,P}^0 = 99$. The true state is $(V, I) = (0, P)$. These prior probabilities strongly suggest that the market is well informed and that an information event is unlikely, but we assume that a negative information event occurs ($V = 0$) and that the market is poorly informed. We focus on extreme event uncertainty as in Section V, subsection B. The rest of the parameters are: $p_L = 0.51$, $\gamma = 0.25$, $\mu_H^W = 0.125$, $\mu_L^W = 0.125$, $\mu_H^P = 0$, and $\mu_L^P = 0.25$. With these parameters, a poorly informed economy has no well-informed traders ($\mu_H^P = 0$) and hence there is very little information about value uncertainty in any trade (since $p_L = 0.51$). Despite the lack of information in this economy, typical simulated price paths, such as the one shown in Figure 1, are highly volatile.²⁰ Figure 2 shows the 20-period moving average of the probability of a buy and of no trade (the probability of a sell is the residual).

With extreme event uncertainty, the price remains close to $1/2$ for the first 30 periods. During this initial interval, however, Figure 2 shows a large buildup of buy orders, which is due to herding. Three of the first five traders buy and this is enough to prompt L types to engage in herd buying. Herd buying lasts from period 5 to period 56. As in the case of event uncertainty alone, the market maker continually increases his prior on an information event as buy orders continue to arrive at a high rate. However, unlike the case of event uncertainty

alone, the price moves dramatically during the interval with herding (from 0.5 to 0.94) as the market maker concludes that there has been an information event. For comparison, the maximum price rise during an interval with herding if the composition of the market were known is $\Delta = 0.03$, from Proposition 8.

Since it is impossible to distinguish between well-informed and poorly informed economies during herding, both individual traders and the market maker must rely on their initial assessments. Because the initial assessments are that a well-informed economy is relatively more likely than a poorly informed one, the market maker increases the price and L types increase their valuations throughout the period of herding almost as if the market were well informed. Eventually, the market maker ends herding by increasing the price and the spread sufficiently in period 57. Figure 2 shows the effect of the end of herding. There is a fall in buying and a rise in no trade as the bid/ask spread forces L types out of the market. This drop in trading volume signals (over time) that previous actions were due to herding rather than to trading by H types. Note the similarity in the two flat spots in the price path. In periods 1–40, the market maker takes time to learn that there has been a change in fundamentals, while in periods 55–100, he takes time to learn that the market is poorly informed. In each case, the market maker is slow to respond because he has extreme beliefs.

Once it becomes apparent that the market is poorly informed, the price naturally has to drop, for there simply is not as much information in previous trading as had been assumed. This brings us back to the case of event uncertainty alone: the price should only have adjusted according to the information content of one poorly informed signal rather than to that of many well-informed trades. As a result, the price falls to near $1/2$ before any further informed trading takes place. Around period 220, the probability of no trade declines as L types reenter the market—this time trading on their information about the new asset value. The market is only learning about one dimension of uncertainty at a time. At first, the market learns that an information event has

¹⁹ The analysis of price paths is a nontrivial exercise. The stochastic process that generates prices is especially complex with herding. In any period the history takes one of three possible values and, depending on the history up until that period, there can be any one of six different distributions over those values. This is why we resort to a simulation at this point in the analysis.

²⁰ The price path in Figure 1 is typical in that most price paths take on extreme values and then return to prices of around $1/2$. It is approximately equally likely that the extreme value is 0 or 1.




FIGURE 1. AN EXAMPLE OF A PRICE BUBBLE

occurred. Then it learns that the market is poorly informed. Only when these first two dimensions are sufficiently resolved, does the market begin to aggregate information about value uncertainty. The price bubble arises because the market mistakenly thinks that it is learning about both event and value uncertainty.

C. Discussion and Connections to Prior Research

It is possible to formalize the above example to show that price bubbles consistently occur under certain identifiable, but unlikely, conditions. There are three key features of the example. First, an information event is very unlikely, $\pi_{1/2}^0 \rightarrow 1$. This assures that the price is fixed for a long period of time, so that a substantial amount of herd behavior occurs. Second, it is very likely that the market is well informed, $\pi_{v,w}^0/\pi_{v,p}^0 \rightarrow 1$. This assures that at first the market maker completely discounts the possibility that the market is poorly informed and that the substantial imbalance in trade is due to herding. Third, all informed traders are of type L in a poorly informed market, $\mu_H^P = 0$. Then a poorly informed market with herd behavior behaves exactly like a well-informed market and nothing is learned about composition uncertainty if there is herding. These effects combine to create highly volatile prices

when the market is poorly informed about an information event. In particular, the price tends arbitrarily close to an extreme value, then returns to $1/2$. The extreme value can be either 0 or 1.²¹

The existence of price bubbles in IS II for extreme parameters is consistent with a general intuition. The combination of event and composition uncertainty leads herd behavior to distort asset prices. So long as the market maker can not completely distinguish between a well-informed market and a poorly informed market during periods of herding, the herd behavior that arises from event uncertainty will distort prices. The more the market maker is surprised that the market is poorly informed, the more prices will respond to the herd behavior.

Our conclusion that rational herding can explain price bubbles and crashes contrasts with several papers which argue implicitly that herding and crashes, specifically the stock market crash of 1987, cannot be explained in models of rational trading (Robert J. Shiller [1989] gives a collection of papers to this effect; Allan W. Kleidon [1992] summarizes and criticizes this line of thought). For example, several papers explain the failure of markets to produce effective prices as the

²¹ We formalize this result in an earlier version of this paper (available from the authors), where we make some additional simplifying assumptions.

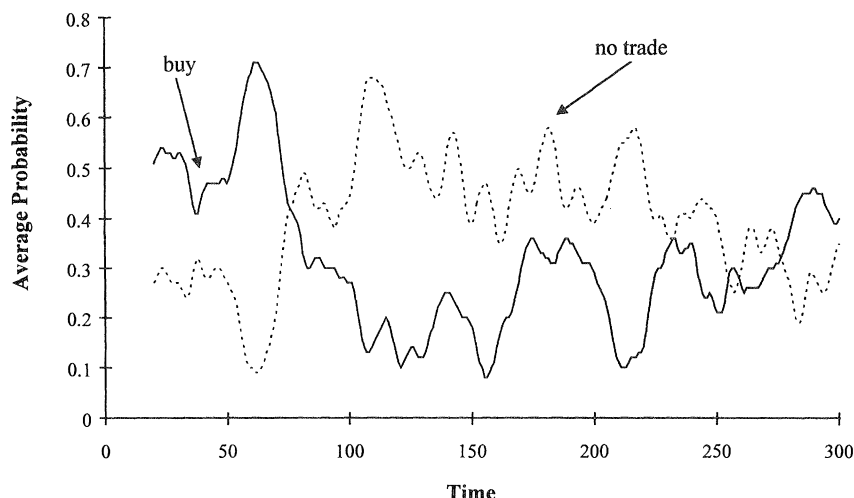


FIGURE 2. TRADING ACTIVITY (MOVING AVERAGE)

result of unsophisticated strategies and sub-optimal behavior by market participants (e.g., Gerard Genotte and Hayne Leland, 1990; Shleifer and Summers, 1990).

Of prior work, Jacklin et al. (1992) and David Romer (1993) come closest to providing a rational actor theory of price bubbles. Both papers have two dimensions of uncertainty: value uncertainty and composition uncertainty. We believe that our use of three dimensions of uncertainty (value, event, and composition uncertainty) provides a more complete theory of price bubbles. Romer studies a noisy rational expectations model with a form of composition uncertainty very much like that in IS II. His theory explains how price corrections can occur without contemporaneous changes in fundamentals: when markets learn about the composition of the market, they reevaluate the information contained in past trades. However, his theory relies on the exogenous mispricing of assets²² and on the exogenous arrival of information about the

composition of the market.²³ In contrast, mispricing in our model arises endogenously through the interaction of herd behavior and composition uncertainty and composition uncertainty is endogenously resolved through the pattern of trade.

Jacklin et al. consider a market with a class of insurance traders who buy stock when the price rises and sell when it declines. They show that such insurance trading creates a positive feedback loop which can produce bubbles and crashes when the market is surprised by the extent of insurance trading. While such insurance trading has some desirable properties when investors hold a diverse portfolio of stock, Jacklin et al. take the use of these strategies as exogenously given. In contrast, the herding strategy that produces our bubble is endogenous.

VII. Contrarian Behavior

In Section V, subsection A, we began to address the puzzle of price charting. We show

²² In Romer's model, mispricing is driven by noise trading and an assumption that traders receive signals which are inaccurate even when perfectly aggregated. Similarly, Lee's (1995) theory of sudden market corrections does not explore the mechanisms by which asset prices become mispriced, beyond noise trading or the arrival of many misinformed traders.

²³ There is a further limitation to Romer's theory. Our results in Section VII below suggest that in a *sequential* trading model, there is a countervailing force that opposes price bubbles when there is composition uncertainty but no event uncertainty. Composition uncertainty creates an incentive for poorly informed traders to trade against the trend in prices. This should limit the formation of price bubbles.

there that an agent’s trading strategy can be strongly based on the history of past trades. In particular, we show that with event uncertainty a trader may ignore his private information about value uncertainty in order to trade *with* the trend in past trades. However, such herd behavior is only one of two possible types of strongly history-dependent behavior. The other possibility is trade which *opposes* the trend in past trades at the expense of private information about value uncertainty. We start with a formal definition of such contrarian behavior.

Definition 6: A trader with private information x_θ engages in *contrarian behavior* at time t if either he buys when $V_\theta^0(x_\theta) < V_m^0$ and $\bar{v}(x_\theta) < V_m^1 < V_m^0$ or he sells when $V_\theta^0(x_\theta) > V_m^0$ and $\bar{v}(x_\theta) > V_m^1 > V_m^0$, where $\bar{v}(x_\theta)$ is defined as follows:

$$\bar{v}(x_\theta) = \lim_{n \rightarrow \infty} E[V | n \text{ draws from the distribution } f_\theta(\cdot | V) \text{ all have the value } x_\theta].$$

The definition of contrarian behavior is the analogue of herd behavior with the additional requirement that the trend in past trades does not overshoot the “limit value” $\bar{v}(x_\theta)$ of the signal. To see why such an addition is necessary, consider a trader who knows that $V = 3/4$ and who trades in a market where $V_m^0 = 1/2$. Initially, the trader wants to buy. If the trend in past trades pushes the price above $3/4$, he will sell. This is not history-dependent behavior. The trading strategy depends only on the price and the signal value (e.g., buy if and only if the asking price is less than $3/4$). In defining contrarian behavior, we seek to exclude situations where the trader reverses his behavior simply because the trend in past trades has become more positive (or negative) than the trader’s information about value uncertainty. With IS II, $\bar{v}(x_L) = x_L$, so that there is contrarian selling if a trader with $x_L = 1$ sells when the trend in trade is positive $V_m^1 > V_m^0$.

PROPOSITION 11: *A trader with a monotonic signal never engages in contrarian behavior.*

PROOF:

Suppose a trader with monotonic signal x_θ engages in contrarian buying at time t . Then $V_\theta^1(x_\theta) \geq A^t \geq V_m^1$. Since the signal is monotonic, this implies that $v(x_\theta) \geq V_m^1$. But since $\bar{v}(x_\theta) = v(x_\theta)$, this contradicts contrarian buying. Similarly, contrarian selling never occurs.

Monotonicity is sufficient to rule out contrarian behavior. Thus, Propositions 3 and 11 demonstrate that an assumption of monotonic signals is inconsistent with strongly history-dependent behavior of both the herd and contrarian variety. While event uncertainty can produce herd behavior, we now show that composition uncertainty can produce contrarian behavior.

PROPOSITION 12: *Consider IS II without event uncertainty (i.e., $\pi_{1/2}^0 = 0$). A sufficient condition for L types to engage in contrarian behavior with positive probability is*

$$(2) \quad \frac{\mu_L^P}{\mu_L^W} > \left(\frac{p_L}{1 - p_L} \right) \left(\frac{\gamma + \mu}{\gamma + \mu_H^W} \right).$$

PROOF:

See Appendix.

When there is composition uncertainty (and no event uncertainty), traders of type L place less weight on previous trades than does the market maker. Why? By definition, an informed trader is more likely to get a low-precision signal in a poorly informed market than in a well-informed market (i.e., $\mu_L^P > \mu_L^W$). Hence, an L type trader assigns a higher probability to $I = P$ than does the market maker: “If this is such a well-informed market, why did I receive such poor information?” In a poorly informed market, a given imbalance between buys and sells is less informative than in a well-informed market. Hence, with composition uncertainty, the market maker adjusts his expected value more in response to past trading than does a trader of type L .²⁴

²⁴ We suspect that composition uncertainty can also create herd behavior. Note that H types believe the market

The sufficient condition in Proposition 12 has three parts. The left-hand side is a measure of the amount of information L types have about the composition of the market. The term $p_L/(1 - p_L)$ on the right-hand side is a measure of the amount of information that L types have about the new asset value. The second right-hand-side term results from the existence of the bid-ask spread: as $\mu \rightarrow 0$, the bid-ask spread goes to zero and $(\gamma + \mu)/(\gamma + \mu_H^W) \rightarrow 1$. Hence, contrarian behavior due to composition uncertainty is shown to be possible when the information of L 's about the composition of the market is large relative to their information about value uncertainty and relative to the bid-ask spread.

With event uncertainty, herd behavior is possible for any imperfect signal (see Proposition 5), but here contrarian behavior only occurs if signals are sufficiently imprecise. The difference arises because with event uncertainty, informed traders know that some states are impossible (i.e., $V \neq 1/2$), while with composition uncertainty, L 's only believe that some states are less likely than the market maker. Note that with $\mu_L^W = 0$, L types know for sure that the market is poorly informed and the sufficient condition is satisfied for $p_L < 1$, which parallels the result for event uncertainty and herd behavior. We draw the following general conclusion. The existence of history-dependent behavior (in either its herd or contrarian form) requires (i) that there exist multiple dimensions of uncertainty, and (ii) that traders' asymmetric information about value uncertainty be sufficiently poor relative to their information about one of the other dimensions of uncertainty.

Multidimensional uncertainty provides a justification for the phenomena of price charting (Easley and O'Hara [1992] and Lawrence Blume et al. [1994] reach similar conclusions). A trader who wants to make optimal

use of all dimensions of his information needs to know more about the trading history than just the price.²⁵

VIII. Conclusion

We reexamined the role of the price mechanism in the aggregation of dispersed private information in an economy when trade is sequential rather than simultaneous. In our general model, the price mechanism assures that long-run choices are efficient and with simple information structures, it assures that herd behavior is impossible. However, we show that more complex information structures can lead to herd behavior and that a sufficiently complex information structure makes price bubbles possible. Price is a single-dimensional instrument and it only assures that the economy learns about a single dimension of uncertainty at one time. As a result, multiple dimensions of uncertainty can "overwhelm" the price mechanism during some stretches of trading. Then, interesting short-run behavior—such as herding, price bubbles, and contrarian behavior—become possible.

Our results are consistent with the literature on trading and common knowledge. Repeated communication leads all agents to agree in their assessments of the true value: they cannot "agree to disagree" (John D. Geanakoplos and Heraklis M. Polemarchakis, 1982). In the simplest examples discussed by Geanakoplos (1992), a single round of communication causes agents to unite in their beliefs; a richer set of possible outcomes necessitates further rounds of communication before the agents agree in their assessments. Adding a new dimension of uncertainty in our model is analogous to enriching the set of outcomes in a common knowledge game. Our results show that communication need not

is more likely well informed than does the market maker. Hence, they put higher weight on the history of trade than does the market maker. Because we assume H types have a perfect signal about value uncertainty ($p_H = 1$), their signal always drives their trading. However, if their signal were imperfect then their private information about market composition might lead them to engage in herd behavior.

²⁵ Adding simple aggregate statistics such as the trading volume or the imbalance of the market maker's sales need not produce a sufficient statistic for the complete history. The meaning of a buy or sell at time t depends on the extent of herd and contrarian behavior at that time.

happen uniformly in a financial market. The market may only be learning about one dimension of uncertainty at a time and with a sufficient number of dimensions, this can lead to highly volatile price paths.

We have investigated whether herding might allow an arbitrageur to profitably manipulate the market's learning process. Consider an arbitrageur who buys in hopes of generating herd buying so as to sell at a later period for a profit. We show elsewhere (Avery and Zemsky, 1998) that such simple trading strategies can not be profitable for a trader with the same information as the market maker.

We close with some ideas for future work. First, we hypothesize that herding and bubbles are less pronounced when prices have multiple dimensions. A natural source of multidimensional prices is derivative securities such as options. Second, while excess volatility can not be explained in our general model, our results demonstrate that volatility concentrates in certain identifiable situations. Our identification of conditions under which price bubbles arise is but a first step in investigating the pooling of variance. Third, we have not fully explored the topic of multidimensional uncertainty. For example, we look at a market where there is either one or no information events. In an economy in which information events arrive stochastically, the market might be uncertain as to how many information events are unfolding at any given time. Empirically, it would be useful to know more about how traders use price history in their trading strategies.

APPENDIX

PROOF OF PROPOSITION 1:

We prove existence and uniqueness of an equilibrium ask price. The proof is similar for bid prices. Let $h_t = B$, the event that there is a buy order at time t . An equilibrium ask price satisfies $E[V|H_t, A^t, h_t] - A^t = 0$. The conditional expected value given a buy order is the weighted average of two terms: the weighted expected value of informed buyers whose assessments sat-

isfy $V_\theta(x_\theta) \geq A^t$, and the assessment of a noise trader V_m^t .

At $A^t = V_m^t$, $E[V|H_t, A^t, h_t] \geq A^t$. As A^t increases from V_m^t , $E[V|H_t, A^t, h_t]$ changes only when A^t outstrips the assessment of some informed traders, at which point those traders stop buying. Consequently, $E[V|H_t, A^t, h_t]$ is weakly increasing in A^t whenever $A^t < E[V|H_t, A^t, h_t]$, and weakly decreasing in A^t whenever $A^t > E[V|H_t, A^t, h_t]$.

The implication is that $E[V|H_t, A^t, h_t] - A^t$ is strictly decreasing once it reaches zero. If there is an equilibrium price, it is unique. $E[V|H_t, A^t, h_t] - A^t$ is continuous in A^t except at finitely many points, where these discontinuities never change the sign of $E[V|H_t, A^t, h_t] - A^t$. In addition, $E[V|H_t, A^t, h_t] - A^t$ is nonpositive at $A^t = V_m^t$, where the market maker gains nothing from noise traders, and nonnegative at $A^t = 1$, where the market maker loses nothing to informed traders. Therefore, a zero profit price must exist, and we know from above that it is unique and that $A^t \geq V_m^t$.

The market maker's expected value for the asset and priors are martingales with respect to H_t since H_t contains all of the market maker's information.

PROOF OF PROPOSITION 2:

Suppose there is an informational cascade in period t . In an informational cascade, the market maker learns nothing from a trade and hence $B^t = A^t = V_m^t$. With noise trading, all histories occur with strictly positive probability in all states of the world. Hence, there does not exist a v such that $P(V = v|H_t) = 1$ and there is a nontrivial set of traders with useful information, [i.e., traders for whom $V_\theta(x_\theta) \neq V_m^t$]. With $B^t = A^t = V_m^t$, these traders must be buying or selling.

Suppose traders with signals $x_\theta \in R_B$ are buying and $P(x_\theta \in R_B|H_t) > 0$. Since this is an informational cascade, $P(x_\theta \in R_B|V, H_t) = P(x_\theta \in R_B|H_t) \forall V$, which implies that $E[V|x_\theta \in R_B, H_t] = V_m^t$. This contradicts all types $x_\theta \in R_B$ buying. Similarly, no positive measure of informed traders can be selling. But this contradicts a nontrivial set of traders having useful information. We conclude that an informational cascade is impossible.

PROOF OF PROPOSITION 4:

This result is a direct consequence of Proposition 4 of Glosten and Milgrom (1985), which states that the beliefs of informed traders and the market maker converge over time so long as trade on both sides of the market is bounded away from zero. Since we assume a stationary probability that noise traders buy and sell in each period, there is a positive probability for a buy order and a positive probability for a sell order in each period. Therefore, the Glosten and Milgrom result applies and the expectations of the market maker and of all the informed traders converge over time. If expectations converge to the true value V , then the prices must do so as well.

Suppose that the market maker's expectation does not converge to V . Then for some $\delta > 0$, there is strictly positive probability in each period that the market maker's expectation differs from V by at least δ . But then, there is a strictly positive probability (in each period) that an informed trader's assessment differs from the market maker's assessment by at least $\epsilon(\delta) > 0$. This contradicts the convergence of these assessments.

PROOF OF COROLLARY 1:

Since V_m^t is a martingale, $E[\Delta V_m^{t_1} \cdot \Delta V_m^{t_2}] = 0$ for each $t_1 \neq t_2$. That is, $\text{Cov}(\Delta V_m^{t_1}, \Delta V_m^{t_2}) = 0$. Since $V_m^{t^*} = V_m^0 + \sum_{t=1}^{t^*} \Delta V_m^t$, we can write the variance of $V_m^{t^*}$ as the sum of variances of ΔV_m^t : $\text{Var}(V_m^{t^*}) = \sum_{t=1}^{t^*} \text{Var}(\Delta V_m^t)$. As $t^* \rightarrow \infty$, $V_m^{t^*}$ converges almost surely to V , so $\text{Var}(V_m^{t^*})$ converges to $\text{Var}(V)$ and the first part of the proposition follows.

For the second part, note that $\text{Var}(V_m^{t_2} - V_m^{t_1}) = \sum_{t=t_1+1}^{t_2} \text{Var}(\Delta V_m^t) = \text{Var}(V_m^{t_2}) - \text{Var}(V_m^{t_1})$. The result follows by taking the limit as t_2 grows large.

PROOF OF COROLLARY 2:

Let $m = P(V < a | H_t)$. Because V_m^t is a martingale converging to V ,

$$V_m^t = mE[V | V < a, H_t] + (1 - m)E[V | V \geq a, H_t].$$

Hence,

$$m = \frac{E[V | V \geq a, H_t] - E[V | H_t]}{E[V | V \geq a, H_t] - E[V | V < a, H_t]}.$$

Since $E[V | V \geq a, H_t] > V_m^t$ and $E[V | V < a, H_t] < a$, an upper bound on m is given by setting $E[V | V \geq a, H_t] = 1$ and $E[V | V < a, H_t] = a$.

PROOF OF PROPOSITION 5:

Suppose $p = 1$. Then $E[V | x, H_t] = x$ and signals are monotonic. Hence there is no herd behavior.

Suppose that $p < 1$ and $V \neq 1/2$. Because of noise trading, any finite history occurs with positive probability. Suppose that there is probability 0 of herding in the first N trades for each finite N . Fix $\epsilon > 0$. Without herding in the first N trades, each buy order increases the expected value of the asset (with an upper limit of 1). Choose n such that an informed trader who observes $n - 1$ buy orders and a signal $x = 0$ has expected value for the asset greater than $1/2 + \epsilon$. Note that each no trade increases $\pi_{1/2}^t$ with an upper limit of 1. Consider a history of length $t = m + n < N$ which consists of m no trades followed by n buy orders, where m is sufficiently large that $\pi_{1/2}^{m+n+1} > 1 - \epsilon$. Under these conditions $A^{m+n} < 1/2 + \epsilon$ and an informed trader with $x = 0$ will buy at time $m + n$. Further, this history occurs with positive probability, contradicting the assumption that there was no herding at time $m + n < N$. A similar argument establishes that herd selling also occurs with positive probability. Therefore, herding in the wrong direction occurs with positive probability.

PROOF OF PROPOSITION 6:

There is herd buying if $E[V | x = 0, H_t] > A^t$, which is equivalent to

$$\begin{aligned} & \frac{(1 - p)\pi_1^t}{(1 - p)\pi_1^t + p\pi_0^t} \\ & > \frac{\frac{1}{2}\gamma\pi_{1/2}^t + \pi_1^t(\gamma + p\mu)}{\gamma + \pi_1^t p\mu + \pi_0^t(1 - p)\mu}. \end{aligned}$$

Setting $\gamma = (1 - \mu)/3$ and $\pi'_{1/2} = 1 - \pi'_0 - \pi'_1$, the above condition is equivalent to

$$\begin{aligned} \Delta(p) &= \pi'_1(1 + \pi'_0 - \pi'_1)(1 - \mu) + 6\mu\pi'_1\pi'_0 \\ &\quad - p((1 - \mu)(\pi'_1(1 + \pi'_0 - \pi'_1) \\ &\quad + \pi'_0(1 + \pi'_1 - \pi'_0)) + 12\mu\pi'_1\pi'_0) \\ &> 0. \end{aligned}$$

We have $\Delta(1) < 0$ and $\Delta(1/2) = (\pi'_1(1 - \pi'_1) - \pi'_0(1 - \pi'_0))(1 - \mu)$, which is positive if and only if $\pi'_1 > \pi'_0$. Hence, when $\pi'_1 > \pi'_0$, there exists a unique $\bar{p} \in (1/2, 1)$ such that $\Delta(\bar{p}) = 0$. Since $\partial\Delta/\partial p < 0$, there is herd buying for $p < \bar{p}(\mu, \pi'_0, \pi'_1)$, where the closed-form expression for \bar{p} comes from solving $\Delta(\bar{p}) = 0$. It is then straightforward to show that $\partial\bar{p}/\partial\mu < 0$. To show that \bar{p} is increasing in $\pi'_{1/2}$, take the expression for \bar{p} and set $\pi'_0 = \alpha\pi'_0$, $\pi'_1 = \alpha\pi'_1$ and $\pi'_{1/2} = 1 - \alpha(\pi'_0 + \pi'_1)$ and then note that $\partial\bar{p}/\partial\alpha < 0$. Finally, to show that \bar{p} is increasing in π'_1/π'_0 , set $\pi'_0 = k - \pi'_1$ and note that $\partial\bar{p}/\partial\pi'_1 > 0$. The results extend to herd selling by symmetry.

PROOF OF PROPOSITION 7:

Suppose that in the first t periods there is no herding and there are b buy orders and s sell orders, where $wlog b \geq s$. In these t periods, informed traders with $x = 1$ submit buy orders while those with $x = 0$ submit sell orders. At time $t + 1$, the assessment of an informed trader with signal $x = 0$ is

$$V^{t+1}(0) = \frac{(1 - p)(\mu p + \gamma)^{b-s}}{(1 - p)(\mu p + \gamma)^{b-s} + p(\mu(1 - p) + \gamma)^{b-s}},$$

and $V(0)$ is simply a function $G(b - s)$. $G(b - s) > 1/2$ whenever $(b - s)$ is equal to or greater than a critical value \hat{n} . Define $\delta_1 = G(\hat{n}) - 1/2$, and $\delta_2 = 1/2 - G(\hat{n} - 1)$. Generically, $\delta_2 > 0$ as well.

Let $\delta^* = \min(\delta_1, \delta_2)$. Assessments of informed traders always differ from $1/2$ by at least δ^* when there is an imbalance in prior trades.

So long as prices remain in the range $(1/2 - \delta^*, 1/2 + \delta^*)$, the market behaves as if the price were fixed at $1/2$. We now observe that as $\pi^0_{1/2} \rightarrow 1$, both the bid and ask prices remain in this range $(1/2 - \delta^*, 1/2 + \delta^*)$ for an arbitrarily long time.

If there is no herding through period t , then the market maker's assessed probability of an information event is bounded below by his assessment after a history with t consecutive trades:

$$\pi^{t+1}_{1/2} \geq \frac{\pi^0_{1/2}(\gamma)^t}{\pi^0_{1/2}(\gamma)^t + (1 - \pi^0_{1/2})(\mu + \gamma)^t}.$$

For $\pi^0_{1/2}$ sufficiently close to 1, this implies that $\pi^{t+1}_{1/2} > 1 - 2\delta^*$. As a result, the bid and ask prices must be in the range $(1/2 - \delta^*, 1/2 + \delta^*)$ in each period prior to t . Define t^* as the last period such that any trading history produces bid and ask prices in the range $(1/2 - \delta^*, 1/2 + \delta^*)$ for each period through t^* . As $\pi^0_{1/2} \rightarrow 1$, $t^* \rightarrow \infty$.

Prior to time t^* , informed traders with $x = 0$ sell while traders with $x = 1$ buy so long as the absolute difference in buy and sell orders is less than \hat{n} . If the imbalance reaches \hat{n} in period $q < t^*$, herding begins and no new information about value uncertainty reaches the market. Hence, $V^t(x) = V^q(x)$ for $t^* \geq t \geq q$ and herding continues until at least t^* , when it becomes possible for the price to adjust enough to break the herd.

The imbalance between buy and sell orders is a random walk with drift $\pm 2\mu p$ prior to time t^* . As t^* grows large (i.e., $\pi^0_{1/2}$ approaches 1), the law of large numbers implies that the probability that the absolute imbalance reaches \hat{n} prior to time t^* increases to 1. Therefore, herding arises with probability 1 as $\pi^0_{1/2} \rightarrow 1$.

The probability of herding in the wrong direction is

$$\begin{aligned} \lambda &= \frac{(\gamma + \mu(1 - p))^{\hat{n}}}{(\gamma + \mu(1 - p))^{\hat{n}} + (\gamma + \mu p)^{\hat{n}}} \\ &= \left(1 + \left[\frac{\gamma + \mu p}{\gamma + \mu(1 - p)} \right]^{\hat{n}} \right)^{-1}. \end{aligned}$$

It follows from $G(\hat{n} - 1) \leq 1/2 \leq G(\hat{n})$, that

$$\left[\frac{p}{1-p} \right]^2 \leq \left[\frac{\gamma + \mu p}{\gamma + \mu(1-p)} \right]^{\hat{n}} \leq \frac{p}{1-p}.$$

Hence,

$$\lambda \in \left[\frac{(1-p)^2}{p^2 + (1-p)^2}, (1-p) \right].$$

As $\mu \rightarrow 0$, \hat{n} gets arbitrarily large, $G(\hat{n}) \rightarrow 1/2$ and $\lambda \rightarrow (1-p)$.

PROOF OF PROPOSITION 8:

Suppose there is no herding in period t , $h_t = B$, and herd buying begins in period $t+1$ and continues through period T . Then, $V^{t+1}(x=0) > B^{t+1} > B^t > V^t(x=0)$ and since $\pi_1^{t+1} > \pi_0^{t+1}$ for herd buying, we also have $B^t > 1/2$. Since $V^s(x=0) = V^{t+1}(x=0)$ for $t+1 < s \leq T$, an upper bound on the change in price from t to T , is given by

$$(A1) \quad V^{t+1}(x=0) - \max\{V^t(x=0), \frac{1}{2}\}.$$

Defining $v = V^t(x=0)$, we can write

$$V^{t+1}(x=0) = \frac{v(\mu p + \gamma)}{\gamma + \mu p v + \mu(1-p)(1-v)}.$$

Since $0 < \partial V^{t+1}(x=0)/\partial v < 1$ for $1/2 \leq v$ and $0 < \partial V^{t+1}(x=0)/\partial v$ for $v < 1/2$, an upper bound on price changes is given by expression (A1) evaluated at $V^t(x=0) = 1/2$, which simplifies to give $\Delta = 3\mu(p - 1/2)/(2 + \mu)$.

PROOF OF PROPOSITION 9:

Let $\phi_h = E[\pi_{1/2}^{t+1} | V = 1/2, H_t]$ when informed traders herd and let ϕ_v be the same quantity when traders trade based on their information about value uncertainty. With herd buying, all informed traders buy when $x \neq 1/2$ and either sell or refrain from trading when $x = 1/2$. Then, using equation (1) we have

$$\phi_h = \pi_{1/2}^t \left(\frac{\gamma^2}{\gamma + \mu(1 - \pi_{1/2}^t)} + \gamma + \frac{(\gamma + \mu)^2}{\gamma + \mu\pi_{1/2}^t} \right).$$

The expression for herd selling is identical. When traders trade with their information about value uncertainty, they buy when $x = 1$, sell when $x = 0$ and refrain from trading when $x = 1/2$. Then,

$$\phi_v = \pi_{1/2}^t \left(\frac{\gamma^2}{\gamma + \mu(p\pi_1^t + (1-p)\pi_0^t)} + \frac{\gamma^2}{\gamma + \mu(p\pi_0^t + (1-p)\pi_1^t)} + \frac{(\gamma + \mu)^2}{\gamma + \mu\pi_{1/2}^t} \right).$$

The difference between these two quantities takes the following form:

$$\phi_h - \phi_v = [f(a) + f(0)] - [f(b) + f(c)],$$

where $f(x) = (\pi_{1/2}^t \gamma^2)/(\gamma + x)$, $a = \mu(1 - \pi_{1/2}^t)$ and $b + c = a$. The result that herding is more effective at revealing information (i.e., $\phi_h > \phi_v$) follows from the convexity of f .

PROOF OF PROPOSITION 10:

Let $\phi_h = E[(|V_m^{t+1} - V|) | V \neq 1/2, H_t] = \pi_1^t E[1 - V_m^{t+1} | V = 1] + \pi_0^t E[V_m^{t+1} - 0 | V = 0]$ when informed traders herd and let ϕ_v be the same quantity when traders trade based on their information about value uncertainty. *Wlog*, suppose that herding involves buying. We can reduce ϕ_h and ϕ_v to expressions in $P(h_t | V = v)$ and π_v^t using equation (1) for $\pi_v^{t+1}(h_t)$ and the following equations:

$$\begin{aligned} E[V^{t+1} | V = v, H_t] \\ = E[\pi_1^{t+1} + \frac{1}{2}\pi_{1/2}^{t+1} | V = v, H_t], \end{aligned}$$

$$\begin{aligned} E[\pi_v^{t+1} | V = v, H_t] \\ = \sum_{h_t} P(h_t | V = v, H_t) \pi_v^{t+1}(h_t). \end{aligned}$$

When traders engage in herd buying, $P(h_t = S|V) = P(h_t = NT|V \neq 1/2) = P(h_t = B|V = 1/2) = \gamma$, and $P(h_t = B|V \neq 1/2) = P(h_t = NT|V = 1/2) = \gamma + \mu$. When traders trade with their information about value uncertainty, $P(h_t = S|V = 1/2) = P(h_t = B|V = 1/2) = P(h_t = NT|v \neq 1/2) = \gamma$, $P(h_t = NT|V = 1/2) = \gamma + \mu$, $P(h_t = B|V = 1) = P(h_t = S|V = 0) = \gamma + \mu p$, and $P(h_t = S|V = 1) = P(h_t = B|V = 0) = \gamma + \mu(1 - p)$. Hence, $\phi_h - \phi_v$ is a function of the exogenous parameters, μ and p , and the current priors. In particular, one can show that $\phi_v - \phi_h$ takes the form $\alpha_2\mu^2 + \alpha_3\mu^3 + \alpha_4\mu^4$, where

$$\alpha_2 = \pi_1'(1 + \pi_0' - \pi_1')(1 - p)$$

$$+ \pi_0'(1 + \pi_1' - \pi_0'),$$

and α_3 and α_4 are also independent of μ . Hence, as μ goes to zero, the sign of $\phi_v - \phi_h$ is given by α_2 . As $\mu \rightarrow 0$, $\Delta(p) \rightarrow \alpha_2$, where $\Delta(p)$ is as defined in the proof of Proposition 6. Hence, as μ becomes small, herding occurs precisely when it serves to minimize $\phi_v - \phi_h$.

PROOF OF PROPOSITION 12:

The proof shows the possibility of contrarian selling, the result for contrarian buying follows from symmetry. Note that $\bar{v}(x_L) = x_L$, so that if a trader with $x_L = 1$ sells when $V_m^i > V_m^o$, then there is contrarian selling. The proof proceeds in three steps. The first shows that given (2), a sufficiently long period of buys implies that all L types stop buying. Step 2 shows that a sufficiently long period of buying and (2) implies that L types all sell given a condition on the priors when the buying started. Step 3 shows that an initial history of trading exists such that this condition on the priors is satisfied. We say that L types trade with their signal in some period t if those with $x_L = 0$ sell and those with $x_L = 1$ buy.

Step 1: A sufficiently long sequence of buys leads L types with $x_L = 1$ to stop buying given condition (2).

If L 's are trading with their signal, then the ask at time t is

$$\begin{aligned} \hat{A}^t = & [(\gamma + \mu_H^W + \mu_L^W p_L)\pi_{1,W}^t \\ & + (\gamma + \mu_H^P + \mu_L^P p_L)\pi_{1,P}^t] \\ & \div [(\gamma + \mu_H^W + \mu_L^W p_L)\pi_{1,W}^t \\ & + (\gamma + \mu_H^P + \mu_L^P p_L)\pi_{1,P}^t \\ & + (\gamma + \mu_L^W(1 - p_L))\pi_{0,W}^t \\ & + (\gamma + \mu_L^P(1 - p_L))\pi_{0,P}^t]. \end{aligned}$$

L types with a signal $x_L = 1$ do not buy at time t if $\hat{A}^t > V_L^t(1)$, which is equivalent to

$$(A2)$$

$$\begin{aligned} & \frac{\pi_{0,P}^t}{\pi_{1,P}^t} (\mu_L^P(1 - p_L)(\gamma + \mu_H^W) - \gamma p_L \mu_L^W) \\ & + \frac{\pi_{0,W}^t}{\pi_{1,P}^t} (\mu_L^W(1 - p_L)(\gamma + \mu_H^W) - \gamma p_L \mu_L^W) \\ & - \frac{\pi_{0,P}^t}{\pi_{1,W}^t} (\mu_L^P(1 - p_L)(\gamma + \mu_H^P) - \gamma p_L \mu_L^P) \\ & - \frac{\pi_{0,W}^t}{\pi_{1,W}^t} (\mu_L^W(1 - p_L)(\gamma + \mu_H^P) - \gamma p_L \mu_L^P) \\ & > 0. \end{aligned}$$

In each period where $h_t = B$ and L 's are trading with their signal, the market maker's priors are updated as follows

$$\frac{\pi_{0,P}^t}{\pi_{1,P}^t} = \frac{(\gamma + \mu_L^P(1 - p_L)) \pi_{0,P}^{t-1}}{(\gamma + \mu_H^P + \mu_L^P p_L) \pi_{1,P}^{t-1}} = a_1 \frac{\pi_{0,P}^{t-1}}{\pi_{1,P}^{t-1}},$$

$$\frac{\pi_{0,W}^t}{\pi_{1,P}^t} = \frac{(\gamma + \mu_L^W(1 - p_L)) \pi_{0,W}^{t-1}}{(\gamma + \mu_H^P + \mu_L^P p_L) \pi_{1,P}^{t-1}} = a_2 \frac{\pi_{0,W}^{t-1}}{\pi_{1,P}^{t-1}},$$

$$\frac{\pi_{0,P}^t}{\pi_{1,W}^t} = \frac{(\gamma + \mu_L^P(1 - p_L)) \pi_{0,P}^{t-1}}{(\gamma + \mu_H^W + \mu_L^W p_L) \pi_{1,W}^{t-1}} = a_3 \frac{\pi_{0,P}^{t-1}}{\pi_{1,W}^{t-1}},$$

$$\frac{\pi_{0,W}^t}{\pi_{1,W}^t} = \frac{(\gamma + \mu_L^W(1 - p_L)) \pi_{0,W}^{t-1}}{(\gamma + \mu_H^W + \mu_L^W p_L) \pi_{1,W}^{t-1}} = a_4 \frac{\pi_{0,W}^{t-1}}{\pi_{1,W}^{t-1}}.$$

Note that $a_i \in [0, 1]$ and $a_1 > \max\{a_2, a_3, a_4\}$.

Suppose $h_t = B$ and L 's trade with their signal for $S \leq t < S + T$. Condition (A2) for period $S + T$ can be written as

$$\alpha_1 a_1^T \frac{\pi_{0,P}^S}{\pi_{1,P}^S} + \alpha_2 a_2^T \frac{\pi_{0,W}^S}{\pi_{1,P}^S} + \alpha_3 a_3^T \frac{\pi_{0,P}^S}{\pi_{1,W}^S} + \alpha_4 a_4^T \frac{\pi_{0,W}^S}{\pi_{1,W}^S} > 0,$$

where $\alpha_1 = (\mu_L^P(1 - p_L)(\gamma + \mu_H^W) - \gamma p_L \mu_L^W)$, $\alpha_2 = (\mu_L^W(1 - p_L)(\gamma + \mu_H^W) - \gamma p_L \mu_L^W)$, etc. Hence condition (A2) is satisfied for T sufficiently large if $\alpha_1 > 0$. Solving $\alpha_1 > 0$ for $(1 - p_L)/p_L$ gives

$$\frac{1 - p_L}{p_L} > \left(\frac{\mu_L^W}{\mu_L^P} \right) \left(\frac{\gamma}{\gamma + \mu_H^W} \right),$$

which is implied by condition (2).

Step 2: Show that a sufficiently long sequence of buys starting in period S leads to herd selling if $\pi_{0,W}^S$ is sufficiently smaller than $\pi_{0,P}^S$. Suppose L types sell regardless of their signal at some time t . Then the bid at time t is

$$(A3) \quad \hat{B}^t = \frac{(\gamma + \mu_L^W)\pi_{1,W}^t + (\gamma + \mu_L^P)\pi_{1,P}^t}{(\gamma + \mu_L^W)\pi_{1,W}^t + (\gamma + \mu_L^P)\pi_{1,P}^t + (\gamma + \mu)(\pi_{0,W}^t + \pi_{0,P}^t)}$$

All L types sell at time t if $\hat{B}^t \geq V_L^t(1)$, which is equivalent to

$$(A4) \quad \frac{\pi_{0,P}^t}{\pi_{1,P}^t} (\mu_L^P(1 - p_L)(\gamma + \mu_L^W) - p_L \mu_L^W(\gamma + \mu)) + \frac{\pi_{0,W}^t}{\pi_{1,P}^t} (\mu_L^W(1 - p_L)(\gamma + \mu_L^W) - p_L \mu_L^W(\gamma + \mu)) - \frac{\pi_{0,P}^t}{\pi_{1,W}^t} (\mu_L^P(1 - p_L)(\gamma + \mu_L^P) - p_L \mu_L^P(\gamma + \mu)) - \frac{\pi_{0,W}^t}{\pi_{1,W}^t} (\mu_L^W(1 - p_L)(\gamma + \mu_L^P) - p_L \mu_L^P(\gamma + \mu)) > 0.$$

In each period where $h_t = B$ and no L types are buying, the market maker's priors are updated as follows

$$\frac{\pi_{0,P}^t}{\pi_{1,P}^t} = \frac{\gamma}{(\gamma + \mu_H^P + \mu_L^P p_L)} \frac{\pi_{0,P}^{t-1}}{\pi_{1,P}^{t-1}} = b_1 \frac{\pi_{0,P}^{t-1}}{\pi_{1,P}^{t-1}},$$

$$\frac{\pi_{0,W}^t}{\pi_{1,P}^t} = \frac{\gamma}{(\gamma + \mu_H^P + \mu_L^P p_L)} \frac{\pi_{0,W}^{t-1}}{\pi_{1,P}^{t-1}} = b_1 \frac{\pi_{0,W}^{t-1}}{\pi_{1,P}^{t-1}},$$

$$\frac{\pi_{0,P}^t}{\pi_{1,W}^t} = \frac{\gamma}{(\gamma + \mu_H^W + \mu_L^W p_L)} \frac{\pi_{0,P}^{t-1}}{\pi_{1,W}^{t-1}} = b_2 \frac{\pi_{0,P}^{t-1}}{\pi_{1,W}^{t-1}},$$

$$\frac{\pi_{0,W}^t}{\pi_{1,W}^t} = \frac{\gamma}{(\gamma + \mu_H^W + \mu_L^W p_L)} \frac{\pi_{0,W}^{t-1}}{\pi_{1,W}^{t-1}} = b_2 \frac{\pi_{0,W}^{t-1}}{\pi_{1,W}^{t-1}}.$$

Note that $b_i \in [0, 1]$ and $b_1 > b_2$.

Suppose $h_t = B$, and L 's are trading with their signal for all periods t such that $S \leq t < S + T + U$, where T takes on the smallest value such that condition (A2) is satisfied. Then condition (A4) for period $S + T + U$ can be written as

$$\beta_1 a_1^T b_1^U \frac{\pi_{0,P}^S}{\pi_{1,P}^S} + \beta_2 a_2^T b_1^U \frac{\pi_{0,W}^S}{\pi_{1,P}^S} + \beta_3 a_3^T b_2^U \frac{\pi_{0,P}^S}{\pi_{1,W}^S} + \beta_4 a_4^T b_2^U \frac{\pi_{0,W}^S}{\pi_{1,W}^S} > 0,$$

where $\beta_1 = \mu_L^P(1 - p_L)(\gamma + \mu_L^W) - p_L \mu_L^W(\gamma + \mu)$, $\beta_2 = \mu_L^W(1 - p_L)(\gamma + \mu_L^W) - p_L \mu_L^W(\gamma + \mu)$, etc. The condition is therefore satisfied for U sufficiently large if $\beta_1 > 0$ and $\pi_{0,W}^S$ sufficiently small relative to $\pi_{0,P}^S$. Solving $\beta_1 > 0$ for $(1 - p_L)/p_L$ gives condition (2).

Step 3: A trading history H_S exists such that $\pi_{0,W}^S/\pi_{0,P}^S$ is arbitrarily close to zero.

Consider some $t < S$. If L types are not trading regardless of their signal and $h_t = NT$, then the market maker's updated priors satisfy

$$\frac{\pi_{0,W}^t}{\pi_{0,P}^t} = \frac{(\gamma + \mu_L^W)}{(\gamma + \mu_L^P)} \frac{\pi_{0,P}^{t-1}}{\pi_{1,P}^{t-1}} < 1.$$

Consider some $t < S$. If L 's are trading with their signal in periods t and $t + 1$, $h_t = B$ and $h_{t-1} = S$, then the market maker's updated priors satisfy

$$\frac{\pi_{0,W}^t}{\pi_{0,P}^t} = \frac{(\gamma + \mu_H^W + \mu_L^W p_L)(\gamma + \mu_L^W(1 - p_L))}{(\gamma + \mu_H^P + \mu_L^P p_L)(\gamma + \mu_L^P(1 - p_L))} \\ \times \frac{\pi_{0,P}^{t-2}}{\pi_{1,P}^{t-2}} < 1.$$

Hence, one can construct an initial series of trades which make $\pi_{0,W}^S/\pi_{0,P}^S$ arbitrarily close to zero.

REFERENCES

- Avery, Christopher and Zemsky, Peter. "Herding and Manipulation in a Financial Market." Mimeo, Harvard University, 1998.
- Banerjee, Abhijit. "A Simple Model of Herd Behavior." *Quarterly Journal of Economics*, August 1992, 107 (3), pp. 797–818.
- Bikhchandani, Sushil; Hirshleifer, David and Welch, Ivo. "A Theory of Fads, Fashion, Custom and Cultural Change as Informational Cascades." *Journal of Political Economy*, October 1992, 100(5), pp. 992–1027.
- Blume, Lawrence; Easley, David and O'Hara, Maureen. "Market Statistics and Technical Analysis: The Role of Volume." *Journal of Finance*, March 1994, 49(1), pp. 153–81.
- Brown, David and Jennings, Robert. "On Technical Analysis." *Review of Financial Studies*, Winter 1989, 2(4), pp. 527–51.
- Bulow, Jeremy and Klemperer, Paul. "Rational Frenzies and Crashes." *Journal of Political Economy*, February 1994, 102(1), pp. 1–24.
- Caplin, Andrew and Leahy, John. "Sectoral Shocks, Learning and Aggregate Fluctuations." *Review of Economic Studies*, October 1993, 60(4), pp. 777–94.
- . "Business as Usual, Market Crashes and Wisdom After the Fact." *American Economic Review*, June 1994, 84(3), pp. 548–66.
- Chamley, Christophe and Gale, Douglas. "Information Revelation and Strategic Delay in a Model of Investment." *Econometrica*, September 1992, 62(5), pp. 1065–86.
- Devenow, Andrea and Welch, Ivo. "Rational Herding in Financial Economics." *European Economic Review*, April 1996, 40(3–5), pp. 603–15.
- Easley, David and O'Hara, Maureen. "Time and the Process of Security Price Adjustment." *Journal of Finance*, June 1992, 47(2), pp. 576–605.
- Geanakoplos, John. "Common Knowledge." *Journal of Economic Perspectives*, Fall 1992, 6(4), pp. 53–82.
- Geanakoplos, John D. and Polemarchakis, Heraklis M. "We Can't Disagree Forever." *Journal of Economic Theory*, October 1982, 28(1), pp. 192–200.
- Genotte, Gerard and Leland, Hayne. "Market Liquidity, Hedging and Crashes." *American Economic Review*, December 1990, 80(5), pp. 999–1021.
- Glosten, Lawrence and Milgrom, Paul. "Bid, Ask and Transaction Prices in a Specialist Market with Heterogeneously Informed Traders." *Journal of Financial Economics*, March 1985, 14(1), pp. 71–100.
- Grossman, Sanford J. "An Analysis of the Implications for Stock and Futures Price Volatility of Program Trading and Dynamic Hedging Strategies." *Journal of Business*, July 1988, 61(3), pp. 275–98.
- Grossman, Sanford J. and Stiglitz, Joseph E. "On the Impossibility of Informationally Efficient Markets." *American Economic Review*, June 1980, 70(3), pp. 393–408.
- Jacklin, Charles J.; Kleidon, Allen W. and Pfeleiderer, Paul. "Underestimation of Portfolio Insurance and the Crash of October 1987." *Review of Financial Studies*, Spring 1992, 5(1), pp. 35–63.
- Kirman, Alan. "Ants, Rationality and Recruitment." *Quarterly Journal of Economics*, February 1993, 108(1), pp. 137–56.
- Kleidon, Allan W. "Market and Environmental Uncertainty," in Peter Newman, Murray Milgate, and John Eatwell, eds., *The new Palgrave dictionary of money and finance*. London: MacMillan, 1992, pp. 651–54.

- Lee, In Ho.** "Market Crashes and Informational Avalanches." Mimeo, University of Southampton, U.K., 1995.
- Milgrom, Paul and Stokey, Nancy.** "Information, Trade and Common Knowledge." *Journal of Economic Theory*, February 1982, 26(1), pp. 17–27.
- Romer, David.** "Rational Asset-Price Movements Without News." *American Economic Review*, December 1993, 83(5), pp. 1112–30.
- Shiller, Robert J.** "Stock Prices and Social Dynamics." *Brookings Papers on Economic Activity*, 1984, (2), pp. 457–98.
- . *Market volatility*. Cambridge, MA: MIT Press, 1989.
- Shleifer, Andrei and Summers, Lawrence H.** "The Noise Trader Approach to Finance." *Journal of Economic Perspectives*, Spring 1990, 4(2), pp. 19–34.
- Vives, Xavier.** "The Speed of Information Revelation in a Financial Market Mechanism." *Journal of Economic Theory*, October 1995, 67(1), pp. 178–204.
- . "Learning from Others: A Welfare Analysis." *Games and Economic Behavior*, August 1997, 20(2), pp. 177–200.