

**MULTI-DIMENSIONAL UNCERTAINTY
AND HERD BEHAVIOR IN
FINANCIAL MARKETS**

by

C. AVERY*
and
P. ZEMSKY**

96/79/FIN/SM

- * Assistant Professor at Kennedy School, Harvard University, Boston, USA.
- ** Assistant Professor of Strategy at INSEAD, Boulevard de Constance, Fontainebleau 77305 Cedex, France.

A working paper in the INSEAD Working Paper Series is intended as a means whereby a faculty researcher's thoughts and findings may be communicated to interested readers. The paper should be considered preliminary in nature and may require revision.

Printed at INSEAD, Fontainebleau, France.

Multi-Dimensional Uncertainty and Herd Behavior in Financial Markets

Christopher Avery and Peter Zemsky *

November 1, 1996

Abstract

We study the relationship between rational herd behavior and asset prices. We define herd behavior as occurring when an agent trades against his initial assessment and instead follows the trend in previous trade. When traders have an informational advantage on a single dimension (the new asset value), price adjustments by a competitive market maker prevent any herd behavior. If the market maker is additionally uncertain as to whether the underlying asset value has changed, we show that herd behavior is possible. However, such herd behavior need not affect the asset price because the market correctly discounts the informativeness of trades during periods of herding. When the market is uncertain about both whether the asset value has changed and whether traders are well or poorly informed on average about the new asset value, then herd behavior can lead to significant, short-run price movements that do not reflect the true asset value.

*Kennedy School of Government and INSEAD, respectively. We are grateful for comments and advice from Anat Admati, Darrell Duffie, Glenn Ellison, Drew Fudenberg, Jim Hines, Allan Kleidon, Paul Milgrom, Paul Pfleiderer, Lones Smith, Peter Sorensen, Xavier Vives and Richard Zeckhauser.

1 Introduction

Walrasian models of general equilibrium rely on simultaneous execution of a large number of trades to produce optimal outcomes. In an Arrow-Debreu economy, all trades take place at the instant the market opens, presuming that the Walrasian auctioneer sets prices correctly. As a challenge to these results, an explosion of papers in the last several years argue that imitative or herd-like behavior can impede the flow of information in a closed economy when consumers act sequentially rather than concurrently. (Bikhchandani, Hirschleifer and Welch (1992), Banerjee (1992), Bulow and Klemperer (1994), Caplin and Leahy (1993, 1994), Chamley and Gale (1992)). 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, Hirschleifer and Welch (BHW) describe that situation as an “informational cascade”. In BHW and Banerjee’s models, an informational 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 markets—especially financial 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 influential market participants and financial economists reportedly still believe that imitative behavior is widespread in financial markets (Devenow and Welch, 1996). This has led some researchers to assert that market participants engage in non-rational herd behavior (e.g. Kirman, 1993, Shleifer and Summers, 1990).

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 following 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

¹The only one of these models which allows explicit price changes is that of Bulow and Klemperer, but their model still fixed the prices locally after each purchase for a sufficient period to produce herding.

for excess volatility? We begin our analysis with an example which motivates a final and no less important question.

1.1 A Simple Example

Our model retains the basic features of BHW, with the notable addition of a price mechanism. It is useful to review the simplest version of the BHW model and to consider what happens when prices are allowed to vary over time in response to trading.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 $\pi_1^t = P(V = 1|H_t)$. 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):

$$\begin{aligned} V^t(x = 0) &= E[V|x = 0, H_t] \\ &= \frac{1 - p}{(1 - p)\pi_1^t + p(1 - \pi_1^t)} \pi_1^t \end{aligned}$$

Assuming all prior agents have acted in accord with their signal, π_1^t 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 $V^t(x = 1) > V^t(x = 0) > 1/2$. 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:

$$\hat{c} = E[V|H_t] = P(V = 1|H_t) = \pi_1^t.$$

Retaining the rest of the BHW assumptions leaves agent valuations unchanged. The key

²We use notation consistent with the rest of this paper.

observation from this simple exercise is that

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

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 is important for the existence of 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.

1.2 Overview of the Paper

In Section 2 we describe a general model and define terms. Of particular importance, we define herd behavior as occurring when an informed agent trades with the trend in past trades even though that trend is counter to his initial information about the new asset value.

In Section 3, we show that there are limits to the distortions that can arise in a financial market where 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, we generalize the example in Section 1.1 to derive a general "monotonicity" condition for private signals such that herding is impossible.

In Section 4 we exhibit a plausible information structure in which herding does occur, adding *event uncertainty* to the information structure of Section 3. With event uncertainty, the market is uncertain as to whether the value of the asset has changed from its initial expected value, while informed traders know for sure whether or not the value has changed. We show that any amount of event uncertainty produces the possibility of herd

behavior. As event uncertainty grows large (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 does change. 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 even clearly at odds with optimal social learning in our setting. Herd behavior arises in situations where traders believe that the market price does not accurately reflect the implications of past trade. As a result, herd behavior can actually aid in price convergence.

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 5 we investigate what happens when the market faces both event uncertainty and what we call *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 leads (almost surely) to a price bubble and crash. In these states, market participants have a mistaken, but rational, belief that most traders possess very accurate information even though everyone has actually received poorly informative signals. 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 Grossman (1988) and Jacklin et al. (1992) for the stock market crash of 1987: traders underestimated the prevalence of non-informative 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 6 we consider the converse of herding, "contrarian behavior", where agents ignore their private information about the asset value to trade *against* the trend in past trades. We show that composition uncertainty can give rise to such behavior. In Section 7 we argue that the herd behavior in our model does not introduce the possibility of

market manipulation by a long-lived, but uninformed, trader. Section 8 concludes.

2 The General Model

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. 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 that an informed trader arrives in any given period; $1 - \mu$ is the probability that a noise trader arrives. For convenience, we assume that noise traders buy, sell and don't 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 The probability that a trader of type θ arrives is denoted $\mu_\theta > 0$.

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 non-trivial 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 $E[V|x_\theta, H_t] \neq E[V|H_t]$ for $x_\theta \in R$; and moreover, if $|E[V|H_t] - V| = \delta > 0$ then for some $\epsilon(\delta) > 0$, $|E[V|x_\theta, H_t] - E[V|H_t]| > \epsilon(\delta)$.

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

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

⁴Thus, a trader potentially has *two* pieces of private information, the value of $x_\theta \in [0, 1]$ and his type θ . Type constitutes private information if there is uncertainty about the composition of the market. See Section 6 for details.

is, the trader who arrives in period t faces a bid, B^t , and an ask, A^t , which satisfy:

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

and

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

where h_t is the action taken by the trader who arrives in period t , where $h_t = B$ indicates a buy, $H_t = S$ indicates a sell and $h_t = NT$ indicates no trade. The market maker's expected value for the asset given public information, $E[V|H_t]$, plays an important role in the paper. We define $V_m^t = E[V|H_t]$, which we shall sometimes refer to as the price.⁵ Finally, we define the market maker's assessed distribution for the possible values as $\pi_v^t = P(V = v|H_t)$. By Bayes' theorem, these priors respond to trade as follows:

$$\pi_v^{t+1} = \pi_v^t \frac{P(h_t|V = v)}{P(h_t)}, \quad (1)$$

where $P(h_t) = \sum_v \pi_v^t P(h_t|V = v)$.

Our model is a special case of the model developed by 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, B^t , and ask, A^t , prices which satisfy $B^t \leq V_m^t \leq A^t$. V_m^t and π_v^t 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, the ask price (the price for a buy order) is above the current assessment of the asset's value and the bid price is below the current assessment of the asset's value. That is, $A^t > V_m^t$ and $B^t < V_m^t$. V_m^t and π_v^t 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^t and π_v^t would be systematically mistaken in a manner which should be predictable to him.

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

2.1 The Definition of Herd Behavior

We differentiate between an informational cascade and herd behavior. In the example of section 1.1, herd behavior always implies an informational cascade. With the simple information structure used there, no information reaches the market when traders with bad signals ($x = 0$) imitate traders with good signals ($x = 1$). However, in a more general model (e.g. with multiple types of traders), 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.

Definition 2 *A trader with private information x_θ engages in herd behavior at time t if either he buys when $E[V|x_\theta] < E[V] < E[V|H_t]$ or he sells when $E[V|x_\theta] > E[V] > E[V|H_t]$; 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: $E[V|x_\theta] < E[V]$. Second, the history of trading must be positive: $E[V] < E[V|H_t]$.⁶ Finally, the trader must want to buy given this positive history and his signal, which implies that $E[V|x_\theta, H_t] > E[V] \geq E[V|H_t]$. 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 $E[V|H_t]$.

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

⁶We write $E[\cdot]$ for $E[\cdot|H_0]$.

reliance on public information (see Vives, 1995a).⁷ In contrast, we start with a behavioral definition of herding and then study the extent to which such behavior leads to distortions and inefficiencies.

3 Bounds on the Effect of Herding

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

Proposition 2 *An informational cascade never occurs in market equilibrium.*⁸

Proof: Suppose there is an informational cascade in period t . In an informational cascade, the market maker learns nothing from a trade and hence $B^t = A^t = V_m^t = E[V|H_t]$. 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 non-trivial set of traders with useful information, (i.e. traders for whom $E[V|x_\theta, H_t] \neq E[V|H_t]$). With $B^t = A^t = E[V|H_t]$, 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 $E[V|x_\theta \in R_B, H_t] = E[V|H_t]$. This contradicts all types $x_\theta \in R_B$ buying. Similarly, no positive measure of informed traders can be selling. But this contradicts a non-trivial set of traders having useful information. We conclude that an informational cascade is impossible. \square

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 observed actions are not independent of the state. Hence, an informational cascade is impossible. Like

⁷In 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 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 (see the discussion in Section 4.3). 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 (1995a,1995b)).

⁸We are very grateful to an anonymous referee for suggesting this result.

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 depends on there being 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 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 $E[V|x_\theta, H_t]$ is always (weakly) between $v(x_\theta)$ and $E[V|H_t]$ 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 1.1, which are often used in Glosten-Milgrom style models, are monotonic because $E[V|x, H_t] \in [x, E[V|H_t]]$. In addition, noisy rational expectations models (e.g. Grossman and 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 micro-structure 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 $E[V|x_\theta, H_t] > A^t \geq E[V|H_t]$. Since the signal is monotonic, this implies that $v(x_\theta) > E[V|H_t]$. But since $E[V|H_t] > E[V]$ for herd buying, monotonicity also implies that $E[V|x_\theta] > E[V]$ and the trader was not originally pessimistic, which is a contradiction. Similarly, herd selling never occurs. \square

With monotonicity, 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.⁹ If we abstract from

⁹Vives (1995b) develops a dynamic noisy rational expectations model which compliments our analysis. Consistent with Proposition 3 and his use of monotonic signals, traders in Vives' model never engage in herd behavior 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.

the existence 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^t , and sell if it is below. 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^t , can summarize for traders all that 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.¹⁰

In Section 4 we show that there exist plausible non-monotonic 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: 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.

□

¹⁰We 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 multi-dimensional prices as arising when there are derivative securities (such as options) that are traded.

From Glosten and Milgrom we already know that all private information becomes public over time as long as the market does not break down. Convergence is then a direct consequence of the assumption that a non trivial amount of useful private information always exists (so long as the true value is not yet identified by the trading history). Then, the only way for all private information to become public is for the price to converge.

While our 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 $\Delta V_m^t = E[V|H_t] - E[V|H_{t-1}]$.

Corollary 5 *The variance of price paths is bounded as follows.*¹¹

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

Additionally, for a fixed t_1

$$\text{Var}(V - V_m^{t_1}) = \text{Var}(V) - \text{Var}(V_m^{t_1}).$$

Proof: 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 $V_m^{t^*} = V_m^0 + \sum_{t=1}^{t=t^*} \Delta V_m^t$, we can write the variance of $V_m^{t^*}$ as the sum of variances of ΔV_m^t : $\text{Var}(V_m^{t^*}) = \sum_{t=1}^{t=t^*} \text{Var}(\Delta V_m^t)$. 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 $\text{Var}(V_m^{t_2} - V_m^{t_1}) = \sum_{t=t_1+1}^{t=t_2} \text{Var}(\Delta V_m^t) = \text{Var}(V_m^{t_2}) - \text{Var}(V_m^{t_1})$. The result follows by taking the limit as t_2 grows large. \square

The first part of Corollary 5 states that the expected volatility is bounded by the fundamental uncertainty over V . Hence, it is not possible to explain “excess” volatility 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

¹¹We are grateful to Paul Milgrom for suggesting this result.

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 6 Consider some $a < E[V|H_t]$. Then $P(V < a|H_t) \leq (1 - E[V|H_t])/(1 - a)$.

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

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

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 $E[V|V \geq a, H_t] > E[V|H_t]$ 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 $E[V|V < a, H_t] = a$. \square

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 6 implies that $P(V = 0) = 1 - E[V|H_t]$. Hence, as the price ($E[V|H_t]$) 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, we note that 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.

4 Event Uncertainty and Herd Behavior

4.1 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 (Brown and Jennings, 1990). This 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 $v(x_\theta) < B^t$.¹²

While it is certainly not difficult to specify non-monotonic signals, the more interesting question is whether such signals are likely to be common in financial markets. Consider 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* (Easley and O'Hara (1992)). We offer the following formalization of event uncertainty.

Definition 4 *There is event uncertainty when $1 > P(V = E[V]) > 0$, where $E[V]$ is the market maker's initial expected value for the asset.*¹³

We now extend the simple BHW information structure used in the example of Section 1.1 to incorporate event uncertainty. Traders are informed if there is an information event (i.e. $V \neq E[V]$) 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 $\pi_{\frac{1}{2}}^0 > 0$ and $\pi_1^0 = \pi_0^0 > 0$. Then, $E[V] = \frac{1}{2}$ and there is event uncertainty. There is a single type of informed trader with signal x , where

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

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

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

¹²Hence, 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.

¹³The event uncertainty that we study below satisfies the additional property that informed traders know whether or not an information event has occurred. That is, $P(V = E[V]|x_\theta) \in \{0, 1\}$.

where $1 \geq p > \frac{1}{2}$. Note that if there were no event uncertainty (i.e. $\pi_{\frac{1}{2}}^0 = 0$), then signals are monotonic. The addition of event uncertainty in this way makes signals non-monotonic and herd behavior possible.¹⁴

Proposition 7 *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.*¹⁵

Proof: 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 $V \neq \frac{1}{2}$. 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 $\frac{1}{2} + \epsilon$. 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 $\pi_{\frac{1}{2}}^{m+n+1} > 1 - \epsilon$. Under these conditions $A^{m+n} < \frac{1}{2} + \epsilon$ 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. \square

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 value of the asset 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

¹⁴It 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 for sure whether or not an information event has occurred, which we find to be a natural feature of event uncertainty.

¹⁵Easley 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. There is one minor difference with their paper. They assume that informed traders are only active in the market if there is an information event. We assume that informed traders do trade if there is no information event and the bid-ask spread does not contain their expected value of $\frac{1}{2}$.

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. The addition of event uncertainty causes informed traders to differ from the market maker in their interpretation of the trading history and hence herd behavior becomes possible.

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 market maker, who must consider the possibility that there has been no change in the underlying value of the asset and the imbalance is due to noise traders. Thus, event uncertainty dulls price adjustment in the short run. Given the results of the literature on herd behavior with fixed prices, it is not surprising that sufficiently dulled price adjustment leads to herd behavior.

The final step to understanding Proposition 7 is to see why for any amount of event uncertainty (i.e. for any $\pi_{\frac{1}{2}}^0 > 0$), price adjustment can become sufficiently dulled to create herding. If informed agents are only expected to buy or sell when there has been an information event (as happens at the outset of trade), then a long period of no trade reduces the market maker's assessed probability of an information event. A sufficiently long period without trade drives $\pi_{\frac{1}{2}}^t$ towards one. Then, the bid and ask prices tend towards $\frac{1}{2}$ for any given imbalance in trading since the market maker is heavily discounting the possibility of an information event.

Proposition 8 *Consider IS I and some trading history H_t that results in priors $\pi_v^t = P(V = v|H_t)$. For $\pi_1^t = \pi_0^t$, there is no herd behavior. For $\pi_1^t \neq \pi_0^t$, there exists a critical value for the precision of traders signals $\bar{p}(\mu, \pi_0^t, \pi_1^t)$ such that traders engage in herd behavior in period t iff $p < \bar{p}$. This \bar{p} decreases with μ and increases with $\pi_{\frac{1}{2}}^t$ (holding π_1^t/π_0^t constant). If $\pi_1^t > \pi_0^t$, then any herd behavior involves buying and \bar{p} increases with π_1^t/π_0^t (holding $\pi_{\frac{1}{2}}^t$ constant). If $\pi_1^t < \pi_0^t$ then any herd behavior involves selling and \bar{p} increases with π_0^t/π_1^t (holding $\pi_{\frac{1}{2}}^t$ constant).*

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

$$\frac{(1-p)\pi_1^t}{(1-p)\pi_1^t + p\pi_0^t} > \frac{\frac{1}{2}\gamma\pi_{\frac{1}{2}}^t + \pi_1^t(\gamma + p\mu)}{\gamma + \pi_1^t p\mu + \pi_0^t(1-p)\mu}$$

The above condition is equivalent to

$$\begin{aligned}\Delta(p) &= \pi_1^t(1 + \pi_0^t - \pi_1^t)(1 - \mu) + 6\mu\pi_1^t\pi_0^t \\ &\quad - p \left((1 - \mu)(\pi_1^t(1 + \pi_0^t - \pi_1^t) + \pi_0^t(1 + \pi_1^t - \pi_0^t) + 12\mu\pi_1^t\pi_0^t) \right) \\ &> 0.\end{aligned}$$

We have $\Delta(1) < 0$ and $\Delta(\frac{1}{2}) = (\pi_1^t(1 - \pi_1^t) - \pi_0^t(1 - \pi_0^t))(1 - \mu)$, which is positive iff $\pi_1^t > \pi_0^t$. Hence, when $\pi_1^t > \pi_0^t$, there exists a unique $\bar{p} \in (\frac{1}{2}, 1)$ such that $\Delta(\bar{p}) = 0$. Since $\partial\Delta/\partial p < 0$, there is herd buying for $p < \bar{p}(\mu, \pi_0^t, \pi_1^t)$, where \bar{p} is yet to be determined. Solving $\Delta(\bar{p}) = 0$ yields a closed form expression for \bar{p} . It is then straight forward to show that $\partial\bar{p}/\partial\mu < 0$. To show that \bar{p} is increasing in $\pi_{\frac{1}{2}}^t$, take the expression for \bar{p} and set $\pi_0^t = \alpha\pi_0^t$, $\pi_1^t = \alpha\pi_1^t$ and $\pi_{\frac{1}{2}}^t = 1 - \alpha(\pi_0^t + \pi_1^t)$ and then note that $\partial\bar{p}/\partial\alpha < 0$. Finally, to show that \bar{p} is increasing in π_1^t/π_0^t , set $\pi_0^t = k - \pi_1^t$ and note that $\partial\bar{p}/\partial\pi_1^t > 0$.

The results extend to herd selling by symmetry. \square

Proposition 8 identifies the forces that produce herding. In general, 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 $|\pi_1^t - \pi_0^t|$ increases the amount of information contained in the history. Either change makes it easier for trading history to overwhelm the information about value uncertainty in private signals and thus makes it easier for herding to arise. Herding becomes more prevalent when prices become less responsive to trading history. As the probability of an information event decreases (i.e. $\pi_{\frac{1}{2}}^t$ increases), prices respond less to the trading history and thus more of the information in the trading history is private.¹⁶ Finally, the effect of an increase in the proportion μ of informed traders operates through the bid-ask spread. An increase in μ causes the market maker to set a wider bid-ask spread because he is losing more to informed trading. That increase in the bid-ask spread opposes herding since traders are then less inclined to trade.

4.2 Existence of Pronounced Herd Behavior

Proposition 7 shows that herding is *possible* for any $p < 1$ and $\pi_{\frac{1}{2}}^t > 0$. We now show that as event uncertainty becomes extreme, herd behavior becomes pronounced, resembling

¹⁶Note that the information gleaned by informed traders from observing trading history is fixed regardless of the value of $\pi_{\frac{1}{2}}^t$ because they know whether or not an information event has occurred.

the cascades of BHW.

Proposition 9 *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. $\pi_{\frac{1}{2}}^0 \rightarrow 1$), the probability that there is some herd behavior goes to 1. Moreover, the trading history almost surely takes the following form:*

A finite, initial period of trading during which herd behavior does not occur,

An arbitrarily long period of herd behavior of one type (i.e. always buy or always sell).

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. $\pi_{\frac{1}{2}}$ close to 1), the market maker discounts almost completely the informativeness of trading. The price remains fixed at the initial expected asset value of $\frac{1}{2}$ 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.¹⁷

There is an important distinction between the herding of Proposition 9 and informational cascades, which cannot occur in equilibrium. In an informational cascade, no new information reaches the market, as when all traders take the same action regardless of their information. Under the conditions of Proposition 9, behavior resembles a cascade when there has been an information event. For a very long period, all informed traders act as buyers (or sellers). However, new information still reaches the market during the period of herding because the consensus among the informed traders is itself informative. Had there been no information event, the volume of trade and the percentage of buy

¹⁷The 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 non adopters to start an informational cascade. Here the imbalance between 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 with the same signal have the wrong signal $(1-p)^2/(p^2 + (1-p)^2)$. Here λ is between the probability that two traders with the same signal have the wrong signal and the probability that just one trader has the wrong signal.

orders would have been much 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 value uncertainty that prices adjust, which ends herding.

Adding a second dimension of uncertainty makes herd behavior possible and even extreme. However, this herd behavior does not distort the asset price.

Corollary 10 *During the entire period of herding identified in Proposition 9, the movement in the asset price from $\frac{1}{2}$ is less than*

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

The maximum possible price rise is thus $p - \frac{1}{2}$ and in the limit as either $\mu \rightarrow 0$ or $p \rightarrow \frac{1}{2}$, $\Delta = 0$.

Proof: See Appendix.

No one is fooled by herd behavior in IS I. Herding keeps information about the new asset value from entering the market and rational actors (including the market maker) account for this. With extreme event uncertainty, herding starts when the value of an informed trader crosses $\frac{1}{2}$, which occurs as soon as the history of trades conveys information about value uncertainty equivalent to just one signal. Then, the valuations of traders are fixed for the entire period of herding. Therefore, herding cannot move the price away from $\frac{1}{2}$ by more than the information contained in one trade. In some circumstances, the change is small. In particular, for a small probability of an informed trader arriving or for low precision signals, each trade conveys little information and has a small effect on valuations.

In conclusion, any price rise during periods with herding in IS I 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.

4.3 Efficiency and Herd Behavior

Vives (1995a) studies the relationship between efficiency and imitative behavior in a setting with fixed prices. 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. If it is very important to learn whether or not an information event has occurred, then herding is actually the socially preferred trading strategy.¹⁸

Proposition 11 *Consider IS I and suppose that the social objective is to maximize, with each trade, awareness of information events. That is, the objective is to make $E[\pi_{\frac{1}{2}}^{t+1}|V \neq \frac{1}{2}, H_t]$ as small as possible.¹⁹ Then, herding is always more socially efficient than having agents trade according to their information about value uncertainty.*

Proof: We work with the dual objective of maximizing $E[\pi_{\frac{1}{2}}^{t+1}|V = \frac{1}{2}, H_t]$. With herd buying, all informed traders buy when $x \neq \frac{1}{2}$ and either sell or refrain from trading when $x = \frac{1}{2}$. Then, using Equation (1) we have

$$\phi_h = E[\pi_{\frac{1}{2}}^{t+1}|V = \frac{1}{2}, H_t] = \pi_{\frac{1}{2}}^t \left(\frac{\gamma^2}{\gamma + \mu(1 - \pi_{\frac{1}{2}}^t)} + \gamma + \frac{(\gamma + \mu)^2}{\gamma + \mu\pi_{\frac{1}{2}}^t} \right).$$

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 $x = \frac{1}{2}$. Then,

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

¹⁸The resolution of event uncertainty is of primary importance when decision makers in the real economy have access to a technology which enables them to learn about what is going on inside a firm once they are alerted to the existence of an information event.

¹⁹Since $\pi_{\frac{1}{2}}^t$ is a martingale, this is equivalent to making $E[\pi_{\frac{1}{2}}^{t+1}|V = \frac{1}{2}, H_t]$ as large as possible.

The difference between these two quantities takes the following form

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

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

While herding is costly in that it obscures information about value uncertainty, it has a benefit. It is more effective (than trading based on information about value uncertainty) 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.

The social objective in Proposition 11 is extreme in two ways. First, it puts no weight on learning about whether an information event is positive or negative. Second, it assumes that the decision maker in the real economy has access to the whole trading history in addition to the current asset price. As Corollary 10 shows, herding can have little effect on the price—even as it reveals that an information event has occurred. An objective that might better capture the interests of more decision makers is one which leads prices to reflect the nature of any information event. Such an objective is relevant to an individual who is going to relocate his family to join one of several firms. What such an individual wants to know is whether there has been a good or bad shock to the fortunes of one of his perspective employers. Surprisingly, such an individual may be quite happy with the functioning of a decentralized market.

Proposition 12 *Consider IS I and suppose that the social objective is to maximize, with each trade, the movement of the price towards any new value of the asset, where the importance of moving to some new value V is weighted by π_V^t . That is, the social planner seeks to make $\pi_1^t E[1 - V_m^{t+1} | V = 1] + \pi_0^t E[V_m^{t+1} - 0 | V = 0]$ as small as possible. Then, as $\mu \rightarrow 0$ herding is preferred by the social planner to having agents trade with their signals about value uncertainty precisely when they do so in the decentralized economy.²⁰*

Proof: Let ϕ_h be the value of the social objective, $\pi_1^t E[1 - V_m^{t+1} | V = 1] + \pi_0^t E[V_m^{t+1} - 0 | V = 0]$, when traders herd and ϕ_v be the value when traders trade based on their

²⁰The result does not hold for general μ . Numerical analysis suggests that in general there is too little herding for $\mu > 0$ in the decentralized economy. Note that the objective is myopic. An investigation of a dynamic notion of efficiency, while desirable, is beyond the scope of this paper. Nor does the objective negatively weight movements in the asset price from 1/2 when there is no information event.

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

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

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

When traders engage in herd buying, $P(h_t = S|V) = P(h_t = NT|V \neq \frac{1}{2}) = P(h_t = B|V = \frac{1}{2}) = \gamma$, $P(h_t = B|V \neq \frac{1}{2}) = P(h_t = NT|V = \frac{1}{2}) = \gamma + \mu$. When traders trade with their information about value uncertainty, $P(h_t = S|V = \frac{1}{2}) = P(h_t = B|V = \frac{1}{2}) = P(h_t = NT|v \neq \frac{1}{2}) = \gamma$, $P(h_t = NT|V = \frac{1}{2}) = \gamma + \mu$, $P(h_t = B|V = 1) = P(h_t = S|V = 0) = \gamma + \mu p$, 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

$$\alpha_2 = \pi_1^t(1 + \pi_0^t - \pi_1^t)(1 - p) + \pi_0^t(1 + \pi_1^t - \pi_0^t),$$

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 8. Hence, for μ sufficiently small, traders in the decentralized economy herd precisely when a social planner prefers herding. \square

The decentralized economy can come arbitrarily close to maximizing the movement of the asset price towards its new value. 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 8) to see why they might be socially desirable. *Ceteris paribus* traders herd when they have low precision signals, but that is when the social cost to herd behavior (lost information about the new asset value) is small. They do not herd when $\pi_1^t = \pi_0^t$, but this is when information that an event has occurred has no impact on prices. Conversely, traders herd when there is already good information about the new asset value (i.e. $|\pi_1^t - \pi_0^t|$ large), but this is when the social benefit to informed trading (more information about the new asset value) is small. They herd when there is little awareness of an information event ($\pi_{\frac{1}{2}}^t$ large), but then trading on information about value uncertainty is not desirable because price responds only sluggishly to information

about the new asset value.

Proposition 12 is a fitting end to Sections 3 and 4, which sing the praises of informationally efficient financial markets populated by rational traders.

5 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 martingale that converges to the true value, it is not possible to have excess volatility in our general model (Corollary 5), nor can price bubbles be both likely to occur and extreme (Corollary 6). 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.

5.1 Uncertainty about the Precision of Aggregate Information

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. For example, without detailed information about a merger partner it may be difficult to estimate the synergies. 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 with 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 $V \in \{0, \frac{1}{2}, 1\}$. 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 $1 > p_L > \frac{1}{2}$. 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^I 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 then $\pi_{V,I}^t$.

5.2 An Example of a Price Bubble

We consider the effect of adding composition uncertainty to the extreme event uncertainty studied in Section 4.2. As with IS I, prices remain close to $\frac{1}{2}$ initially, and traders with

imprecise signals (the L types) may engage in herd behavior. We now develop an example to demonstrate how composition uncertainty serves to link herd behavior to potentially extreme price movements.

We simulate a price bubble in IS II using the following parameters: $p_L = .51$, $\gamma = .25$, $\mu_H^W = .125$, $\mu_L^W = .125$, $\mu_H^P = 0$, and $\mu_L^P = .25$.²¹ The initial priors are $\pi_{\frac{1}{2}}^0 = .9999$, $\pi_{v,W}^0/\pi_{v,P}^0 = 99$. 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. With our choice of 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 $p_L = .51$). Figure 1 shows a typical simulated price path.²² 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).




Figure 1: Example of a Price Bubble

²¹The analysis of price paths is a non trivial exercise. The stochastic process that generates prices is especially complex with herding. In any period the history takes one of three possible value 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 simulations and in Proposition 13 to the study of extreme parameters.

²²It is typical in that most price paths take on extreme values and then return to prices of around $\frac{1}{2}$. It is approximately equally likely that the extreme value is 0 or 1.



Figure 2: Trading Activity (moving average)

With extreme event uncertainty, the price is quite stable for the first 25-30 periods. During this initial interval, however, Figure 2 shows a large build-up 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 assessment of the likelihood of an informational event as buy orders continue to arrive at a high rate. However, unlike the case of event uncertainty alone, the price moves dramatically as the market maker concludes that there has been an information event.

Since it is impossible to distinguish between well and poorly informed economies during the period of 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 as if the market were well informed.

Eventually, the market maker ends herding and breaks the partial informational cascade by increasing the price beyond the valuation of L types (i.e. to 0.94 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 $\frac{1}{2}$ 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. It is only after the market learns that there has been an information event and that the market is poorly informed, that information about the new asset value begins to arrive.

5.3 A Formal Result

To formalize the above example, we make several strong assumptions to show that a bubble can occur with probability near 1 when it is very likely that the market is well informed and very unlikely that an information event has occurred, but an event occurs about which the market is poorly informed.

- (A1) Type L traders receive a signal of precision $\frac{1}{2} + \epsilon$ where ϵ is assumed to be small.
- (A2) An information event is very unlikely: $\pi_{\frac{1}{2}}^0 \rightarrow 1$.
- (A3) It is very likely that the market is well informed: $\pi_{v,W}^0 / \pi_{v,P}^0 \rightarrow 1$ for $V \neq \frac{1}{2}$.
- (A4) All informed traders are of type L in a poorly informed market: $\mu_H^P = 0$; $\mu_L^P = \mu$. Almost all informed traders are of type H in a well informed market: $\mu_H^W \rightarrow \mu$; $\mu_L^W \rightarrow 0$.
- (A5) A trader of type L continues to place great weight on the likelihood of a well informed market:

$$\frac{\mu_L^W \pi_{v,W}^0}{\mu_L^W \pi_{v,W}^0 + \mu_L^P \pi_{v,P}^0} \rightarrow 1.$$

The main effect of the above assumptions is as follows. (A1) assures that L types have such a weak signal that they will herd based on a single imbalance in the trading. (A2) assures that the price is fixed for a long period of time, so that a substantial amount of

herd behavior occurs. (A3) 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. (A4) assures that the market maker learns nothing about the composition of the market during herding. Under (A4) a poorly informed market with herd behavior behaves exactly like a well informed market. (A5) assures that L types also believe initially that the market is well informed. (As discussed in the Section 6, L types update their priors on I based on receiving a low precision signal.) Hence, L types continue to engage in herd behavior even as the price deviates significantly from $\frac{1}{2}$. These effects combine to create highly volatile prices when the market is poorly informed about an information event.

Proposition 13 *Consider IS II when (A1) - (A5) hold and there is an information event ($V \neq \frac{1}{2}$) about which the market is poorly informed ($I = P$). Herd behavior by L types drives the asset price arbitrarily close to 0 or 1, where the probability that the price moves towards $1 - V$ is $\frac{1}{2} - \epsilon$.*

Proof: See Appendix.

Except for one initial informative trade which triggers herding, there is no information revealed about value uncertainty during the dramatic price movement as the L types are herding. Hence, the price movement is not sustainable.

Corollary 14 *Consider IS II when (A1) - (A5) hold and there is an information event about which the market is poorly informed. The price returns to $(\frac{1}{2} - \epsilon), (\frac{1}{2} + \epsilon)$ with probability 1 after it takes on an extreme value arbitrarily close to 0 or 1.*

Hence, assumptions (A1) to (A5) lead to a true price bubble. The price tends arbitrarily close to an extreme value, then returns to $\frac{1}{2}$. It does not matter whether herding drives the price in the right or the wrong direction, for the price must return to $\frac{1}{2}$ in either case. In effect, the market is only learning about one dimension of uncertainty at a time. At first, the market learns that an information event has occurred. Then it learns that the market is poorly informed. Only when these first two dimension 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.

Others have shown that markets do not function well when surprising events occur. For example, Ait-Sahalia (1994) shows that the current pricing of options for very low

probability events is inaccurate. Our results show that this property of financial markets can link herd behavior and price bubbles.

We relied on extreme assumptions to produce a price bubble with probability one in Proposition 13. The result is consistent with a more 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.

5.4 Connection to Prior Work

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 (Shiller (1989) gives a collection of papers to this effect; Kleidon (1992) summarizes and criticizes this line of thought.) For example, several papers explain the failure of markets to produce effective prices as the result of unsophisticated strategies and suboptimal behavior by market participants (e.g. Genotte and Leland (1990), Shleifer and Summers (1990)).

Of prior work, Romer (1993) and Jacklin et al. (1992) 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 explanation of price bubbles. Romer studies a noisy rational expectations model with a form of composition uncertainty very much like that in IS II. Traders vary in the precision of their signals and the proportion of traders with different signals is uncertain. His theory has two shortcomings. First, because he uses a rational expectations model and non-monotonic signals, he is unable to derive analytic results. More fundamentally, his theory only explains how price corrections can occur without contemporaneous changes in fundamentals. The sources of the mispricing are essentially exogenous.²³ In contrast,

²³In Romer's model, mispricing is driven by noise trading and an assumption that traders receive signals which are inaccurate even when perfectly aggregated. In Ho Lee (1995) develops an explanation of sudden market movements based on his observation that transaction costs produce informational cascades. If new information arrives that breaks the cascade, then the previously blocked information can suddenly flow into the market. Then, a mispriced asset may suddenly be corrected. Like Romer, his theory does not explore the mechanisms by which asset prices become mispriced, beyond noise trading or the arrival of many misinformed traders.

mispricing in our model arises endogenously through the interaction of herd behavior and composition uncertainty.²⁴

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.

6 Contrarian Behavior

In Section 4.1 we began to address the puzzle of price charting. We show 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 $E[V|x_\theta] < E[V]$ and $\bar{v}(x_\theta) > E[V|H_t] > E[V]$ or he sells when $E[V|x_\theta] > E[V]$ and $\bar{v}(x_\theta) < E[V|H_t] < E[V]$, where $\bar{v}(x_\theta)$ is defined as follows:*

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

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 to assure that the trade is not based on information about value uncertainty alone, consider a trader who knows for

²⁴There is a further limitation to Romer's theory. Our results in Section 6 below suggest that in a *sequential* trading model, there is a countervailing force that opposes price bubbles due to composition uncertainty alone. We show that with composition uncertainty, poorly informed traders have an incentive to engage in contrarian behavior, where they trade against the trend in prices, which should work against the formation of price bubbles in poorly informed markets.

sure that $V = 3/4$ and who trades in a market where $E[V] = 1/2$. 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 iff ask 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.

Proposition 15 *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 $E[V|x_\theta, H_t] \geq A^t \geq E[V|H_t]$. Since the signal is monotonic, this implies that $v(x_\theta) > E[V|H_t]$. But since $\bar{v}(x_\theta) = v(x_\theta)$, this contradicts contrarian buying. Similarly, contrarian selling never occurs. \square

Monotonicity is sufficient to rule out contrarian behavior. Thus, Propositions 3 and 15 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 16 *Consider IS II without event uncertainty (i.e. $\pi_{\frac{1}{2}}^0 = 0$). A sufficient condition for L types to engage in contrarian behavior with positive probability is*

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

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 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 .²⁵

The sufficient condition in Proposition 16 has three parts. The LHS term is a measure of the amount of information L types have about the composition of the market. The term $p_L/(1-p_L)$ is a measure of the amount of information they have about the new asset value. The second RHS term results from the existence of the bid-ask spread.²⁶ Hence, herding 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. This mirrors the results for herd behavior in Proposition 8.

With event uncertainty, herd behavior is possible for any imperfect signal (see Proposition 7), but here contrarian behavior only arises if signals are sufficiently imprecise. The difference arises because with event uncertainty, informed traders know that some states are impossible (i.e. $V \neq \frac{1}{2}$) 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 again 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 1) that there exist multiple dimensions of uncertainty and 2) that traders' asymmetric information about value uncertainty be sufficiently poor relative to their information about one of the other dimensions of uncertainty.

Thus, multi-dimensional uncertainty (and the resulting non-monotonic signals) provides a justification for the phenomena of price charting (Easley and O'Hara (1992), Blume, Easley and O'Hara (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. He needs to compute his current expected value of the asset, which requires closely tracking the history of trade in order to update his priors on all states of the world.²⁷

²⁵We suspect that composition uncertainty can also create herd behavior. Note that H types believe the market 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, the 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.

²⁶Note that for $\mu \rightarrow 0$, the bid-ask spread goes to zero and $(\gamma + \mu)/(\gamma + \mu_H^W) \rightarrow 1$.

²⁷Adding simple aggregates such as trading volume and 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.

7 Manipulation and Herding

The existence of substantial price effects due to herding raises the question of whether the market is robust to efforts to manipulate price changes. Our results suggest that in certain instances there may be a huge trend in short-run prices which is unanticipated by the market maker. A far-seeing investor with the opportunity to trade on multiple occasions could then have an incentive to make trades which will trigger such a trend.

Allen and Gale (1992) and Benabou and Laroque (1992) demonstrate profitable manipulation in settings where the manipulator takes a conspicuous action in period 0 and then may attempt to profit on any mispricing which results in period 1. Allen and Gale's manipulator threatens a takeover, while Benabou and Laroque's manipulator makes a prediction about the future course of the market. In each case, the conspicuous action may reveal a large amount of private information, or it may be deceptive. A manipulation strategy has a similar motivation in our model, but it is much less likely to be profitable because the manipulator is unable to take a particularly conspicuous action.

The success of a manipulative trade in inducing a herd depends critically on its timing. Our earlier results demonstrate that a single trade can trigger a period of herding, but only under very specialized circumstances. In this section we assume that a manipulator arrives at exactly the right moment, t_1 , to induce herding. We assume that the manipulator has no more information than the market maker, and is simply trying to profit from his ability to deceive the market maker (and the other traders) with a calculated action which does not actually reflect information. Since the manipulator has only the same information as the market maker, his trade only induces herding in some states of the world. We assume that the market maker and other traders are completely unaware of the possibility of manipulation, leaving the pricing and trading rules unchanged from our previous analysis.

We require the manipulator to liquidate his trade, but allow him to choose to do so at any time during the period of herding. In other words, we allow the manipulator to "jump the line" to become the agent who trades with the market maker at any time, t_2 , prior to the end of herding. Once the period of herding ends, we require the manipulator to liquidate immediately by submitting a sell order.²⁸

Proposition 17 *Suppose that there is a manipulative buy order in period t_1 and then*

²⁸We define the end of the period of herding as the first period, t_2^* when there is no herding for any state of the world.

herding proceeds in each period between t_1 and t_2^* . In IS I, V_m^t is a supermartingale in each period in the interval (t_1, t_2^*) . In IS II, V_m^t is a supermartingale in each period in the interval (t_1, t_2^*) such that $B_t > 1/2$.

Proof: See Appendix.

The apparent advantage of the buy order from the manipulator's standpoint is that it may help to trigger herd buying. However, it also distorts the market maker's assessed probabilities for the actions of future traders, including the market maker's assessment of the probability of herd buying. We distinguish the manipulator's information at time t_1 and beyond, Z_t from the market maker's information, which is simply H_t . In contrast to the original model, the market maker's assessed value, V_m^t , is then not a martingale with respect to Z_t after the manipulative buy order.

The manipulator's action influences the market maker's assessment in two ways: first, the market maker overweights the probability that an information event has occurred, and second, the market maker overweights the probability of a positive information event relative to the probability of a negative one. In most instances, this divergence will work against the manipulator rather than in her favor.

Intuitively, the manipulative buy order causes the market maker to be overly optimistic in expecting more buy orders in the future. If the market maker's assessments were correct, then V_m^t would be a martingale. Instead, the "overly optimistic" beliefs induced by the manipulative order cause V_m^t to become a supermartingale with respect to Z_t throughout the period of herding in IS I. In IS II, there are a wider variety of traders and it is only possible to assure that V_m^t is a supermartingale with respect to Z_t after a manipulative buy order when informed traders are sellers if there is no information event (i.e. when $B^t > 1/2$).

Corollary 18 *Manipulation is not profitable for any trading rule for the manipulator in IS I nor is it profitable in IS II under the conditions of Proposition 13.*

Proof: See Appendix.

Even when we allow the manipulator to observe the series of orders and then to choose to transact at any time between t_1 and t_2^* in these cases, the optimal stopping rule for supermartingales implies that he cannot make money. In essence, the argument observes that the market maker's assessments start at the manipulator's trading price in period $t_1 + 1$ and then drift (on average) back towards the original expected value, prior to the manipulative order.

8 Conclusion

We re-examined 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’ (Geanakoplos and Polemarchikis, 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 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 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 multi-dimensional 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 only a first step in investigating the pooling of variance. Third, we have not fully explored the topic of multi-dimensional uncertainty. For example, we look at a market where there was either one or no information events. In an economy in which information events arrive stochastically, there might be more than one information event unfolding at the same time. Empirically, it would be useful to know more about how traders use price history in their trading strategies.

9 Appendix

Proof of Proposition 1: We prove existence and uniqueness of an equilibrium ask price. The proof is similar for bid prices.

Let the event C denote a buy order at time t . An equilibrium ask price satisfies $E(V|H_t, A^t, C) - A^t = 0$. The conditional expected value given a buy order is the weighted average of two terms: the expected value for an informed buyer whose assessment satisfies $V_\theta(x_\theta) \geq A^t$, and the assessment of a noise trader V_m^t .

At $A^t = V_m^t$, $E(V|H_t, A^t, C) \geq A^t$. As A^t increases from V_m^t , $E(V|H_t, A^t, C)$ changes only when A^t outstrips the assessment of some informed traders (e.g. $V_L(1, H_t)$), at which point those traders drop out of the buying market. Consequently, $E(V|H_t, A^t, C)$ is weakly increasing in A^t whenever $A^t < E(V|H_t, A^t, C)$, and weakly decreasing in A^t whenever $A^t > E(V|H_t, A^t, C)$.

The implication is that $E(V|(H_t, A^t, C) - A^t)$ is strictly decreasing once it reaches zero. If there is an equilibrium price, it is unique. $E(V|(H_t, A^t, C) - A^t)$ is continuous in A^t except at finitely many points, where these discontinuities never change the sign of $E(V|(H_t, A^t, C) - A^t)$. In addition, $E(V|(H_t, A^t, C) - A^t)$ 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. \square

Proof of Proposition 9:

Suppose that there have been t periods of informative trading (no herding) with b buy orders and s sell orders, where (WLOG) $b \geq s$. In these t periods, informed traders with positive signals submitted buy orders and those with negative signals submitted sell orders. At time $t + 1$, the assessment of an informed trader with signal 0 is

$$V_L^{t+1}(0) = \frac{(1 - p_L)(\mu p_L + \gamma)^{b-s}}{(1 - p_L)(\mu p_L + \gamma)^{b-s} + p_L(\mu(1 - p_L) + \gamma)^{b-s}}.$$

For given values of p_L and μ (and thus for $\gamma = (1 - \mu)/3$ as well), $V_L^t(0)$ is simply a function $G(b - s)$, and further, $G(b - s) > \frac{1}{2}$ whenever $(b - s)$ is equal to or greater than a critical value \hat{n} . Define $\theta_1 = G(\hat{n}) - \frac{1}{2}$, and $\theta_2 = \frac{1}{2} - G(\hat{n} - 1)$. Generically, $\theta_2 > 0$ as well.

Let $\theta^* = \min(\theta_1, \theta_2)$. Assessments of informed traders always differ from $\frac{1}{2}$ by at least θ^* when there is an imbalance in prior (informative) trades. So long as prices remain in the range $(\frac{1}{2} - \theta^*, \frac{1}{2} + \theta^*)$, they have no effect on the trading decisions of an informed trader.

We now observe that as $\pi_{\frac{1}{2}}^0 \rightarrow 1$, both the bid and ask prices remain in this range $(\frac{1}{2} - \theta^*, \frac{1}{2} + \theta^*)$ for an arbitrarily long time. In effect, the prices will be fixed for a long period after the outset of trade.

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

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

For $\pi_{\frac{1}{2}}^0$ sufficiently close to 1, this implies that $\pi_{\frac{1}{2}}^{t+1} > 1 - 2\theta^*$. As a result, the bid and ask prices must be in the range $(\frac{1}{2} - \theta^*, \frac{1}{2} + \theta^*)$ 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 $(\frac{1}{2} - \theta^*, \frac{1}{2} + \theta^*)$ for each period through t^* . As $\pi_{\frac{1}{2}}^0 \rightarrow 1, t^* \rightarrow \infty$.

Prior to time t^* , informed traders follow their own signals so long as the absolute difference in buy and sell orders is less than \hat{n} . If the imbalance reaches \hat{n} in period $t < t^*$, any informed trader will discard his signal (again supposing WLOG that there is herd buying rather than herd selling). But then there is an informational cascade in this period. Any sell order must be from a noise trader, while a buy order provides no further information. Once herding starts, it continues until the market maker adjusts the price to break the herd. But, by definition, pricing cannot break the herd until after period t^* .

The imbalance between buy and sell orders is a random walk with drift $\pm 2\mu p_L$ prior to time t^* . As t^* grows large (i.e. $\pi_{\frac{1}{2}}^0$ 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 $\pi_{\frac{1}{2}}^0 \rightarrow 1$.

The probability of herding in the wrong direction is

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

Because $G(\hat{n} - 1) \leq \frac{1}{2} \leq G(\hat{n})$, it must be that

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

Hence,

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

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

Proof of Corollary 10: WLOG consider herd buying that starts with an imbalance between buys and sells of \hat{n} . Because the valuations of informed traders do not change during periods of herding, an upper bound on the movement of prices is given by $G(\hat{n}) - \frac{1}{2}$. We can write

$$G(n) = \frac{1 - p_L}{1 - p_L + p_L a^n},$$

where $a = (\mu(1 - p_L) + \gamma)/(\mu p_L + \gamma)$.

Because there is no herding with an imbalance of $\hat{n} - 1$,

$$G(\hat{n} - 1) \leq \frac{1}{2} \Leftrightarrow \frac{1 - p_L}{p_L} \leq a^{\hat{n}-1}.$$

Setting $a^{\hat{n}-1}$ to its lowest possible value yields an upper bound on the movement in prices:

$$\begin{aligned}\Delta &= G(\hat{n}) \Big|_{a^{\hat{n}-1} = \frac{1-p_L}{p_L}} - \frac{1}{2} \\ &= \frac{1-a}{2(1+a)} \\ &= \frac{3\mu(p_L - \frac{1}{2})}{2+\mu}\end{aligned}$$

□

Proof of Proposition 13: Similar arguments to those in the proof of Proposition 9 imply that the bid and ask prices have no effect on trading up to any time t^* so long as $\pi_{\frac{1}{2}}^0$ is sufficiently close to 1.

Consider then the path of trading prior to t^* for a fixed price of $\frac{1}{2}$ in a poorly informed market. In a well informed market, all informed traders know the true value of V , buying if $V = 1$ and selling if $V = 0$. In a poorly informed market, type L traders follow their signal if the number of previous buy and sell orders are equal. If there is any imbalance, type L traders then discard their signal and follow that imbalance.

Trading then follows a Markov chain in either type of market. In a well informed market, the probabilities are fixed: $\mu + \gamma$ of a trade in the correct direction (e.g. a buy when $V = 1$), γ of a trade in the wrong direction, and γ of no trade.²⁹ In a poorly informed market, the probabilities adjust to favor any prior imbalance: $\mu + \gamma$ in the direction of any prior imbalance, γ of opposing that imbalance, and γ of no trade.³⁰

Either of these Markov chains drifts away from zero with probability one in finite time, meaning that a permanent imbalance of trade orders arises with probability one. The law of large numbers implies that the long run imbalance tends to a total of μN orders over a period of N periods where N is very large.

Suppose WLOG that the imbalance is in favor of buy orders and that the imbalance arose at time z . Then trading subsequent to time z produced identical trading probabilities in each period for each of the three cases $(V = 1, W), (V = 0, L), (V = 1, W)$: $\mu + \gamma$ of a buy order and γ each for no trade and a sell order. That large imbalance does enable traders to place arbitrarily small weight on the case $(V = 0, W)$, since that would imply a trend of trading strongly in favor of sell rather than buy orders.

Therefore, the market is only able to learn about the relative probabilities of the three remaining cases in the early period of trading before the permanent imbalance of orders. Suppose that there were x buy and x sell orders prior to that permanent imbalance, and that trading balanced (previous buy orders = previous sell orders) in $y \leq x$ periods. In a poorly informed market, then trading was in favor of an imbalance in $x - y$ periods and against it in x periods. Neglecting periods without trade,

$$\mathcal{L}(V = 0, L|H_z) = \mathcal{L}(V = 1, L|H_z) = ((\frac{1}{2} \pm \epsilon)\mu)^y (\mu + \gamma)^{x-y} (\gamma)^x.$$

²⁹We neglect the infinitesimal probability that a type L trader arrives in a well informed market. Though this is crucial to the beliefs of type L traders, it has no effect on trading probabilities in a well informed market.

³⁰If there is no imbalance, then the probabilities are μp_L of a trade in the correct direction, $\mu(1 - p_L)$ of a trade in the wrong direction, γ of no trade at all.

In contrast,

$$\mathcal{L}(V = 1, W|H_z) = (\mu + \gamma)^x (\gamma)^x > \mathcal{L}(V = 1, L|H_z).$$

In summary, the market learns nothing about the relative likelihoods of $(V = 1, W)$, $(V = 1, L)$, $(V = 0, L)$ during the period where the imbalance of buy orders grows permanently from 0. Any trading prior to that period only increases the presumption in favor of $(V = 1, W)$. We conclude that there is an informational cascade until the ask price outstrips the assessment of type L traders with a poor signal.

By (A5), type L traders weight the probability of a well informed market as large relative to a poorly informed market (even after they observe the precision of their signal. In a well informed market, the imbalance of trades suggests a value of 1. Therefore, type L traders assess the value to be at least $\pi_{1,W}/(\pi_{1,W} + \pi_{1,L} + \pi_{0,L})$, which is arbitrarily close to 1.

To break the informational cascade, the market maker must adjust the ask price to be arbitrarily close to 1. \square

Proof of Proposition 16: 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 $E[V|H_t] > E[V]$, then there is contrarian selling. The proof proceeds in three steps. The first shows that given condition (1), a sufficiently long period of buys implies that all L types stop buying. Step 2 shows that a sufficiently long period of buying and condition (1) 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. The result for contrarian buying follows from symmetry.

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

If $s_L^t(B, 1) = 1$ and $s_L^t(B, 0) = 0$, then the ask at time t is

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

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

$$\begin{aligned} & \frac{\pi_{0,P}^t}{\pi_{1,P}^t}(\mu_L^P(1 - p_L)(\gamma + \mu_H^W) - \gamma p_L \mu_L^W) + \frac{\pi_{0,W}^t}{\pi_{1,P}^t}(\mu_L^W(1 - p_L)(\gamma + \mu_H^P) - \gamma p_L \mu_L^W) - \\ & \frac{\pi_{0,P}^t}{\pi_{1,W}^t}(\mu_L^P(1 - p_L)(\gamma + \mu_H^P) - \gamma p_L \mu_L^P) - \frac{\pi_{0,W}^t}{\pi_{1,W}^t}(\mu_L^W(1 - p_L)(\gamma + \mu_H^P) - \gamma p_L \mu_L^P) > 0. \end{aligned} \quad (2)$$

In each period where $h_t = B$, $s_L^t(B, 1) = 1$ and $s_L^t(B, 0) = 0$, the market maker's priors are updated as follows

$$\begin{aligned} \frac{\pi_{0,P}^t}{\pi_{1,P}^t} &= \frac{(\gamma + \mu_L^P(1 - p_L)) \pi_{0,P}^{t-1}}{(\gamma + \mu_H^P + \mu_L^P p_L) \pi_{1,P}^{t-1}} = a_1 \frac{\pi_{0,P}^{t-1}}{\pi_{1,P}^{t-1}} \\ \frac{\pi_{0,W}^t}{\pi_{1,P}^t} &= \frac{(\gamma + \mu_L^W(1 - p_L)) \pi_{0,W}^{t-1}}{(\gamma + \mu_H^P + \mu_L^P p_L) \pi_{1,P}^{t-1}} = a_2 \frac{\pi_{0,W}^{t-1}}{\pi_{1,P}^{t-1}} \end{aligned}$$

$$\begin{aligned}\frac{\pi_{0,P}^t}{\pi_{1,W}^t} &= \frac{(\gamma + \mu_L^P(1-p_L)) \pi_{0,P}^{t-1}}{(\gamma + \mu_H^W + \mu_L^W p_L) \pi_{1,W}^{t-1}} = a_3 \frac{\pi_{0,P}^{t-1}}{\pi_{1,W}^{t-1}} \\ \frac{\pi_{0,W}^t}{\pi_{1,W}^t} &= \frac{(\gamma + \mu_L^W(1-p_L)) \pi_{0,W}^{t-1}}{(\gamma + \mu_H^W + \mu_L^W p_L) \pi_{1,W}^{t-1}} = a_4 \frac{\pi_{0,W}^{t-1}}{\pi_{1,W}^{t-1}}.\end{aligned}$$

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

Suppose $h_t = B$, $s_L^t(B, 1) = 1$ and $s_L^t(B, 0) = 0$ for $S \leq t < S+T$. Condition (1) 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 (1) 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 (1).

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

$$\hat{B}^t = \frac{(\gamma + \mu_L^W) \pi_{1,W}^t + (\gamma + \mu_L^P) \pi_{1,P}^t}{(\gamma + \mu_L^W) \pi_{1,W}^t + (\gamma + \mu_L^P) \pi_{1,P}^t + (\gamma + \mu) (\pi_{0,W}^t + \pi_{0,P}^t)}. \quad (3)$$

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

$$\begin{aligned}\frac{\pi_{0,P}^t}{\pi_{1,P}^t} (\mu_L^P(1-p_L)(\gamma + \mu_L^W) - p_L \mu_L^W(\gamma + \mu)) + \frac{\pi_{0,W}^t}{\pi_{1,P}^t} (\mu_L^W(1-p_L)(\gamma + \mu_L^W) - p_L \mu_L^W(\gamma + \mu)) - \\ \frac{\pi_{0,P}^t}{\pi_{1,W}^t} (\mu_L^P(1-p_L)(\gamma + \mu_L^P) - p_L \mu_L^P(\gamma + \mu)) - \frac{\pi_{0,W}^t}{\pi_{1,W}^t} (\mu_L^W(1-p_L)(\gamma + \mu_L^P) - p_L \mu_L^P(\gamma + \mu)) > 0.\end{aligned} \quad (4)$$

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

$$\begin{aligned}\frac{\pi_{0,P}^t}{\pi_{1,P}^t} &= \frac{\gamma}{(\gamma + \mu_H^P + \mu_L^P p_L)} \frac{\pi_{0,P}^{t-1}}{\pi_{1,P}^{t-1}} = b_1 \frac{\pi_{0,P}^{t-1}}{\pi_{1,P}^{t-1}} \\ \frac{\pi_{0,W}^t}{\pi_{1,P}^t} &= \frac{\gamma}{(\gamma + \mu_H^P + \mu_L^P p_L)} \frac{\pi_{0,W}^{t-1}}{\pi_{1,P}^{t-1}} = b_1 \frac{\pi_{0,W}^{t-1}}{\pi_{1,P}^{t-1}} \\ \frac{\pi_{0,P}^t}{\pi_{1,W}^t} &= \frac{\gamma}{(\gamma + \mu_H^W + \mu_L^W p_L)} \frac{\pi_{0,P}^{t-1}}{\pi_{1,W}^{t-1}} = b_2 \frac{\pi_{0,P}^{t-1}}{\pi_{1,W}^{t-1}} \\ \frac{\pi_{0,W}^t}{\pi_{1,W}^t} &= \frac{\gamma}{(\gamma + \mu_H^W + \mu_L^W p_L)} \frac{\pi_{0,W}^{t-1}}{\pi_{1,W}^{t-1}} = b_2 \frac{\pi_{0,W}^{t-1}}{\pi_{1,W}^{t-1}}.\end{aligned}$$

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

Suppose $h_t = B$, $s_L^t(B, 1) = 1$ and $s_L^t(B, 0) = 0$ for $S \leq t < S+T+U$, where T takes on the smallest

value such that condition (1) is satisfied. Then condition (8) 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 (1).

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 $s_L^t(NT, 1) = s_L^t(NT, 0) = 1$ and $h_t = NT$, then the market maker updated priors satisfy

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

Consider some $t < S$. If $s_L^t(B, 1) = s_L^t(S, 0) = s_L^{t-1}(B, 1) = s_L^{t-1}(S, 0) = 1$, $h_t = B$ and $h_{t-1} = S$ then the market maker's updated priors satisfy

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

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

In the case where $\pi_{0,W}^0/\pi_{0,P}^0$ is sufficiently close to zero or in the case where $\pi_{0,W}^S/\pi_{0,P}^S$ can be made sufficiently close to zero with only no trades, then herd buying (selling) occurs in a history without any prior buying (selling). \square

Proof of Proposition 17: There was no herding in period t_1 . Therefore, a manipulative buy order in period t_1 increases the market maker's assessment of $P(\tilde{V} = 1)$ and reduces the market maker's assessment of $P(\tilde{V} = 1/2)$. The market maker's assessment of $P(\tilde{V} = 0)$ may increase or decrease as the result of the manipulative buy, though this is not critical to our proof. Let the true trading probabilities at later periods t (for $t_1 < t < t_2$) be designated by the distribution P^t and let the market maker's assessed trading probabilities be designated by the distribution P_m^t .

During the period of herding of IS II with $B^t > 1/2$, the action "no trade" can always be attributed to a noise trader, so that the market maker and the manipulator assess the same probability for "no trade". The market maker has a higher assessed probability for a buy order than does the manipulator and a lower assessed probability for a sell order than does the manipulator, where these assessments differ by the same absolute value. Since a buy order produces a higher expected value than does a sell order, it must be that $E(V_m^{t*+1}|Z_{t*}) \leq E(V_m^{t*+1}|H_{t*}) = V_m^{t*}$

In IS I when there has been an information event (i.e. $V = 1$ or $V = 0$) informed traders are all buyers during the period of herding by the definition of herding. If there has not been an information event, informed traders are either sellers (if $B^t > 1/2$) or they do not trade (if $B^t \leq 1/2$). If $B^t > 1/2$, then the same argument holds as in IS II with $B^t > 1/2$. If $B^t \leq 1/2$, then a sell order can be attributed to a noise trader. Then the market maker has a higher assessed probability for a buy order and a lower assessed probability for a sell order than the manipulator. Once again, since these assessments differ by the same absolute value, V_m^t is a supermartingale with respect to Z_t by the same logic as above. \square

Proof of Corollary 18: In IS I, $V_m(t)$ is a supermartingale in each period between t_1 and t_2^* . In IS II, $V_m(t)$ is a supermartingale in each period between t_1 and t_2^* such that $B_t > 1/2$. But under the conditions of Proposition 13, B_t begins arbitrarily close to $1/2$ and then only increases during the period of herding if there has been an information event. Under these conditions, any nontrivial decline in B_t breaks the period of herding; therefore, $V_m(t)$ is a supermartingale in all periods of herding once prices have moved at all from their initial values.

The manipulator initiates a buy order at the price A^{t_1} , which is given by the market maker's expected value conditional on the trading history and the current trade. That is, the manipulator transacts at the price $V_m^{t_1+1}$, where $V_m^{t_1+1}$ is an expectation which includes the current buy order (as though it were an ordinary buy order and not a manipulative one). The manipulator then hopes to sell that unit of the asset profitably at the bid price at time t_2 , B^{t_2} , where t_2 is chosen at the manipulator's option prior to the end of herding. Since $B^{t_2} \leq V_m^{t_2}$ by definition, a necessary condition for manipulation to be profitable is $E(V_m^{t_2} | Z_{t_1}) \geq V_m^{t_1+1}$. But this is impossible when V_m^t is a supermartingale with respect to Z_t between periods t_1 and t_2 . \square

References

- [1] Ait-Sahalia, Yacine, (1994), "Nonparameteric Interest Rate Pricing of Derivative Securities," mimeo, Graduate School of Business, University of Chicago.
- [2] Allen and Gale (1992), "Stock-Price Manipulation," in *Review of Financial Studies*, 5(3): 503-529.
- [3] Banerjee, Abhijit, (1992), "A Simple Model of Herd Behavior," *Quarterly Journal of Economics*, 107:797-818.
- [4] Benabou and Laroque (1992) "Using Privileged Information to Manipulate Markets: Insiders, Gurus and Credibility," *Quarterly Journal of Economics*, 107(3):921-958.
- [5] Bikhchandani, Sushil, David Hirshleifer and Ivo Welch, (1992), "A Theory of Fads, Fashion, Custom and Cultural Change as Informational Cascades," *Journal of Political Economy*, 100:992-1027.
- [6] Blume, Lawrence, David Easley and Maureen O'Hara, (1994), "Market Statistics and Technical Analysis: The Role of Volume," *Journal of Finance*, 49:153-181.
- [7] Brown, David and Robert Jennings (1990). "On Technical Analysis," *Review of Financial Studies*, 2:527-552.

- [8] Bulow, Jeremy and Paul Klemperer, (1994), "Rational Frenzies and Crashes", *Journal of Political Economy*, 102:1-24.
- [9] Caplin, Andrew and John Leahy, (1994), "Business as Usual, Market Crashes and Wisdom after the Fact," *American Economic Review*, 84:548-566.
- [10] Caplin, Andrew and John Leahy, (1993), "Sectoral Shocks, Learning and Aggregate Fluctuations," *Review of Economic Studies*, 60:777-794.
- [11] Chamley, Christophe and Douglas Gale, (1992), "Information Revelation and Strategic Delay in a Model of Investment," *Econometrica*, 62:1065-1086.
- [12] Devenow, Andrea and Ivo Welch, 1996. "Rational Herding in Financial Economic," *European Economic Review*, 40, pp.603-615.
- [13] Easley, David and Maureen O'Hara, (1992), "Time and the Process of Security Price Adjustment," *Journal of Finance*, 47(2): 576-605.
- [14] Genotte, Gerard and Hayne Leland (1990) "Market Liquidity, Hedging and Crashes," *American Economic Review*, 80, 999-1021.
- [15] Glosten, Lawrence and Paul Milgrom (1985), "Bid, Ask and Transaction Prices in a Specialist Market with Heterogeneously Informed Traders," *Journal of Financial Economics*, 14, 71-100.
- [16] Grossman, Sanford J. (1988), "An Analysis of the Implications for Stock and Futures Price Volatility of Program Trading and Dynamic Hedging Strategies," *Journal of Business*, 61, 275-298.
- [17] Grossman, Sanford J. and Joseph E. Stiglitz (1980), "On the Impossibility of Informationally Efficient Markets," *American Economic Review*, 70, 393-408.
- [18] Jacklin, Charles J., Allen W. Kleidon and Paul Pfleiderer (1992), "Underestimation of Portfolio Insurance and the Crash of October 1987", *Review of Financial Studies*, 5, 35-63.
- [19] Kirman, Alan (1993), "Ants, Rationality and Recruitment," *Quarterly Journal of Economics*, 108(1): 137-156.
- [20] Kleidon, Allan W. (1992), "Market and Environmental Uncertainty," in *The New Palgrave Dictionary of Money and Finance*,

- [21] Milgrom, Paul and Nancy Stokey, (1982), "Information, Trade and Common Knowledge," *Journal of Economic Theory*, 26:17-27.
- [22] Lee, In Ho (1995). "Market Crashes and Informational Avalanches," mimeo, University of Southampton.
- [23] Romer, David (1993). "Rational Asset-Price Movements without News," i.e. *American Economic Review*, pp. 1112-1130.
- [24] Shiller, Robert J. (1989), *Market Volatility*. Cambridge, MA: MIT Press.
- [25] Shiller, Robert J. (1984), "Stock Prices and Social Dynamics," *Brookings Papers on Economic Activity*, 2, 457-498.
- [26] Shleifer, Andrei and Lawrence H. Summers, (1990), "The Noise Trader Approach to Finance," *Journal of Economic Perspectives*, 4:2, 19-34.
- [27] Vives, Xavier (1995a), "Learning from Others: A Welfare Analysis," mimeo.
- [28] Vives, Xavier (1995b), "The Speed of Information Revelation in a Financial Market Mechanism," *Journal of Economic Theory*, 67, 178-204.