Is Beauty Contagious? Higher-Order Uncertainty and Information Aggregation

– Job Market Paper –

Ole Jann*

November 18, 2015

Abstract

I investigate the robustness of information aggregation to higher-order uncertainty. My analysis uncovers a novel mechanism through which a contagion of beliefs can destroy informational efficiency. Consider an asset market where short-lived speculators have information both about the asset's fundamental value and the amount and direction of noise trading. In equilibrium, each speculator's trading takes account of both pieces of information and the market price adjusts to the fundamental information. But if speculators lack common knowledge about noise trading, they worry about other speculators' beliefs, beliefs about beliefs, and so on. Even with minimal higher-order uncertainty, only uninformative (i.e. random) trading is rationalizable. The result can explain observed behavior and has implications for how markets should be organized to make them informationally efficient. In a second application to expert committees, I show that lack of common knowledge among experts with very precise information makes them unable to communicate their information to a decision maker whose interests are aligned with theirs.

Keywords: Speculation, informational efficiency of markets, multiple equilibria, beauty contest, common knowledge, global games, infection, contagion

JEL Classification: D82 (Asymmetric and Private Information), D84 (Expectations; Speculations), G14 (Information and Market Efficiency)

---

*Department of Economics, University of Copenhagen, Øster Farimagsgade 5, DK-1353 Copenhagen K, Denmark. E-mail: ole.jann@econ.ku.dk. I am grateful to Olga Balakina, Amil Dasgupta, Eddie Dekel, Jeppe Druedahl, Nicola Gennaioli, Itay Goldstein, Nenad Kos, Stephen Morris, Michael Møller, Marco Ottaviani, Alessandro Pavan, Sönje Reiche, Jesper Rüdiger, Christoph Schottmüller, Peter Norman Sørensen, Annette Vissing-Jørgensen, Xavier Vives and Asher Wolinsky as well as audiences at Copenhagen University, at the DGPE 2013, EEA 2015 (Mannheim) and EDGE 2015 (Marseille) for helpful comments. An earlier draft of this paper was circulated under the title "When Only the Market Can Vindicate You: Speculation and Inefficient Market Equilibria".

This paper shows that information aggregation in financial markets can be paralyzed by minimal higher-order uncertainty among traders. I demonstrate a novel mechanism by which small doubt about the beliefs of others (and beliefs about beliefs, and so on) leads to a contagion of beliefs which destroys any informationally efficient equilibrium. The result rests mainly on two realistic assumptions: Traders have a short horizon, and they have information about the number and opinion of irrational noise traders.

To understand the mechanism that drives the result, consider the following example: A rational speculator believes that irrationally optimistic traders will drive up the price of a stock. He concludes that he should buy the stock and sell it at a higher price later. A speculator who believes that other speculators believe that optimists will drive up the price will therefore believe that these other speculators will buy the stock and drive up the price, and he will conclude that he should also buy the stock regardless of what he actually believes about the optimists. And so on, to an ever higher degree, until the idea that there are optimists in the market takes on a life of its own without anyone having to believe in it.

If speculators have common knowledge about the actions of noise traders, such contagion of beliefs cannot happen, since everybody knows what everybody knows about the noise traders, and so on. But with the smallest seed of higher-order uncertainty, belief contagion eradicates any connection of price to fundamental value. Note that I am not claiming that the unfounded buying frenzy in the above example is an equilibrium. But under the assumptions mentioned above, this paper argues that a contagion can be unavoidable, and it can destroy all informationally efficient equilibria.

This result offers an explanation for the frequent observation that well-informed speculators seem to trade against their own better knowledge, such as hedge fund managers who bought tech stocks in 1999 or well-connected bankers who did not get out of the market in 1929. A similar contagion can also occur in other institutions for information aggregation, such as expert committees, where it can result in anticipatory obedience to the biases of a decision maker. Understanding the contagion mechanism, and under which conditions and assumptions it emerges, has implications for how financial markets and other institutions for information aggregation should be designed to make them informationally efficient.

I develop a model of an asset market in which speculators have information about the fundamental value v of the asset. v is realized in period 3, but speculators only live until period 2 and can either buy in period 1 and sell in period 2, or the other way around. The speculators therefore want to forecast p_2, the price in period 2, which is determined by the actions of other speculators and noise traders. This problem of trying to forecast a price that is determined by the actions of others who are trying to make the same forecast is what

Keynes (1936) famously called the "beauty contest".[1] I naturally extend Keynes' metaphor by considering the contagion of beliefs if speculators lack common knowledge.[2]

Assume that speculators can observe the order flow from noise traders, x_N. If speculators have common knowledge about v and x_N, there exists an equilibrium in which speculators base their trading on both pieces of information (Proposition 1). In this equilibrium, p_2 is a function of both v and x_N. However, p_2's dependence on v is an equilibrium effect (it only happens if speculators trade on their fundamental information), while noise trading is independent of any strategic reasoning. Only x_N predicts p_2 independently of the strategic decisions of other speculators.

If a speculator believes that x_N is very large (i.e. many noise traders are buying), he expects that p_2 will be high, regardless of what other speculators do. (Even if they all sold, the many buy orders from noise traders would push the price up.) He will therefore condition his trading more on what he knows about noise traders than on what he knows about v. This amplifies the influence that noise traders have on the price, and expectations of strong noise trading can be self-fulfilling. Even a speculator who knows that the actual order flow from noise traders is small can therefore be swayed to base his behavior on that of noise traders by a self-fulfilling worry that others are doing the same. Even small higher-order uncertainty can "infect" the beliefs of speculators, as they worry that others think that there is lots of noise trading, or they worry that others worry about this, and so on.

I show that this contagion of beliefs among speculators is so powerful that there exists no Nash equilibrium in which speculators base their trading on any fundamental information (Proposition 2). The result obtains for the smallest possible amount of higher-order uncertainty. In fact, informative trading is not rationalizable, i.e. any speculator who holds any consistent set of beliefs about the actions of other speculators will conclude that it is never optimal to condition his trading on v. In effect, the contagion of beliefs leads to a "contagion of types," since all informed speculators act like noise traders, and no trades based on fundamental information are made. Prices in periods 1 and 2 are completely independent of v.

Figure 1 illustrates the intuition of the first contagion step. The noise order flow x_N from noise traders determines the sign of the price change only in extreme cases, if x_N is very small or very large. Otherwise, the sign of the price change is determined by the trading

[1]Keynes took the name from a popular newspaper competition where participants had to choose the six most beautiful faces among a hundred photographs, and those who chose the most popular faces could win a prize.

[2]I use the term "common knowledge" in the sense of Aumann (1976), i.e. something is common knowledge if everybody knows it, everybody knows that everybody knows it, and so on. "Higher-order uncertainty" is n-th order uncertainty with an arbitrarily large n, where first-order uncertainty is uncertainty about a variable, second-order uncertainty is uncertainty about someone else's belief about the variable, and so on. Thus higher-order uncertainty implies lack of common knowledge.

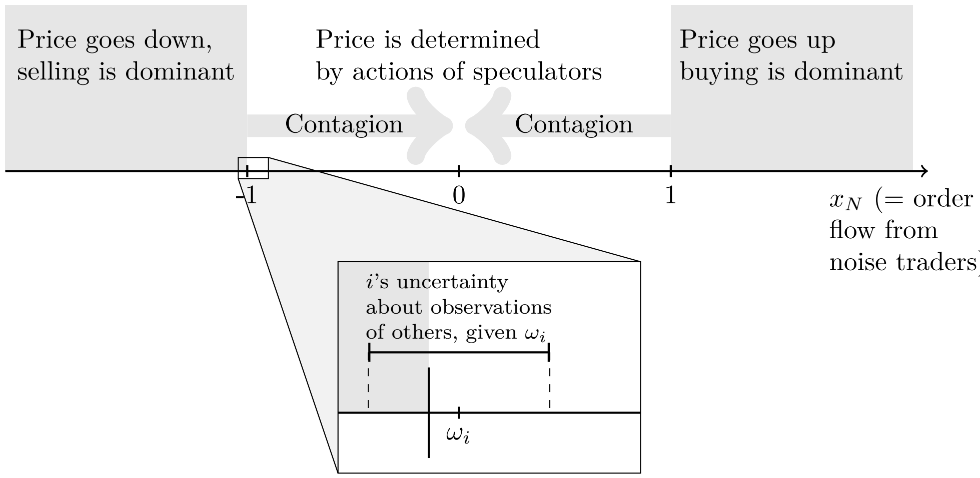

Figure 1: Illustration of the intuition of contagion. The prize is in equilibrium determined by the trading of the speculators, and only in extreme cases by the noise traders. With higher-order uncertainty, a speculator who thinks that the noise trader order flow is close to the cutoff believes that some other speculators believe that it is beyond the cutoff (small, magnified rectangle). It is therefore optimal for him to behave as if x_N was beyond the cutoff.

of speculators, who trade according to their knowledge about fundamental value. But now assume that there is small belief uncertainty, so that speculator i observes signal ω_i, which is a very precise signal of x_N.[3] The small, magnified rectangle shows the problem of a speculator who receives a signal very close to the lower cutoff, below which speculators have no influence on the price change. Because of the small uncertainty, he knows that some other speculators receive a signal that is below the cutoff. These speculators follow their dominant strategy and sell, which effectively turns them into noise traders. This moves the cutoff slightly, since there are fewer speculators "available" that could push the price upwards. It is therefore only rationalizable for i to sell, independently of his knowledge about v.

The main result of this paper is quite stark, and it needs to be qualified to be useful. A general prediction that there is no informative trading in financial markets would quite obviously be at odds with reality. But the results of this paper offer an explanation at what went wrong in times when asset prices were substantially detached from fundamentals, such as just before the Great Crash of 1929 or during the dot-com bubble. In both cases, experts who understood the mispricings refrained from trading on their information.[4] In section 4.2, I discuss several such episodes and relate them to the model.

Even more importantly, understanding under which conditions the result emerges can

[3] I model higher-order uncertainty with the standard global games methodology of Carlsson and van Damme (1993).

[4] There was considerable uneasiness in regulatory circles, the Fed, large banks and the media on the eve of the Great Crash (Galbraith, 1954); fund managers were aware that internet stocks were overpriced but felt they couldn't afford to stay out or short-sell (cf. Abreu and Brunnermeier, 2003).

provide insight into how to design financial markets to make them informationally efficient. While the result is robust to the inclusion of some well-informed and long-lived investors, a market that is made-up mostly of long-term oriented traders would not exhibit beauty-contest features and this model does not apply. This suggests that markets in which there is more short-term trading, and times in which people are looking to make profits quicker, are more susceptible to the belief contagion, and more likely to lead to informationally inefficient prices. The information structure that leads to the contagion is also crucial, since contagion does not occur in models with either more information (i.e. common knowledge) or less information (i.e. no knowledge about noise trading). While the information structure in this model is not implausible, it is clearly important whether traders observe (and think about) the current market sentiment when making their decisions. This offers an explanation for the role that rumors play in financial markets: Rumors are influential not because everybody believes them (most people don't), but because traders don't have common knowledge about the fact that no one believes them. In section 4.1, I discuss which conditions need to be in place to obtain the result, and which tentative policy implications we can draw.

The contagion mechanism is not exclusive to financial markets; it applies in principle to all beauty-contest type models. In section 4.3, I briefly sketch an application to voting in expert committees. Consider an expert committee that has to give a recommendation between two options, where experts receive a payoff from voting for the winning option which is larger than the payoff from voting for the better option if it loses. A contagion of beliefs can make it unrationalizable for the experts to systematically vote for the better option.

The speculators in this paper have two pieces of information on which to base their decision: Fundamental information about the value of the asset, and information about the behavior of noise traders. This has similarities to the studies on private and public information by Morris and Shin (2002) and on beauty contests by Allen, Morris, and Shin (2006), whose central finding is that agents can overweight information that forecasts the actions of others, and underweight information about the state of the world. In this model, the "predictive" information is about the behavior of noise traders. Since noise traders sometimes influence the price to an extreme degree (similar to the "noise trader risk" of De Long et al., 1990), either buying or selling can become the dominant speculator action for some realizations of noise traders behavior.

The central result of theoretical models of contagion is that small higher-order uncertainties can be magnified (Rubinstein, 1989) and select between equilibria (Carlsson and van Damme, 1993). These insights have usually been applied to models in which the contagion occurs on beliefs about a fundamental variable. But this is not necessary. As this paper shows, it can be an otherwise insubstantial variable that completely determines behavior, as long as people find it conceivable that this variable could take values that would make an action strategically dominant. Whether anyone thinks this is likely is not important. That is how it can become uniquely rationalizable for traders to condition their behavior on the irrational ideas of a small group of noise traders.

Global games analysis is usually taken to describe selection between two outcomes. In Morris and Shin (1998), for example, the result is (broadly speaking): For some fundamental values, attack occurs, for others it does not. In this model, the important multiple-equilibrium structure is the existence of two intermediate equilibria: "all speculators trade in the right direction" (i.e. buy if value if high and sell if it is low) and "all speculators trade in the wrong direction." It is the ability of speculators to choose one of the equilibria that makes informative trading possible. But through contagion, they are made unable to choose any equilibrium, and informational efficiency is destroyed. Informative trading happens for no realization of noise trading. Thus the result is qualitatively different from global games applications that are outcome-centered and present a selection between the outcomes. The application to voting in committees (section 4.3) presents this insight even more starkly.

Abreu and Brunnermeier (2003) develop a model in which informed arbitrageurs may delay selling for some time despite knowing that an asset is overvalued. Since arbitrageurs are not aware how many others know of the mispricing, they temporarily cannot coordinate on selling. This argument is based on a similar intuition as the contagion which leads to a durable disconnect of prices from value in this paper.

The following section introduces the model under common knowledge and describes the assumptions in some detail. Section 2 describes the equilibrium without higher-order uncertainty; section 3 derives the main result if speculators lack common knowledge. The discussion (section 4) relates the assumptions and result of the model to policy implications (4.1) and to actual events (4.2). Section 4.3 applies the contagion mechanism to voting in expert committees; section 5 concludes. All proofs are in the appendix.

1 The Model

1.1 General Structure

Consider the market for one asset. The asset has a fundamental value v of either either v_H or v_L, where $v_H > v_L$ and both values are equally probable. v is realized in period 3. There is a group of speculators who know v, but they only live until period 2 and can therefore only trade in periods 1 and 2. In period 1, there are also noise traders who buy or sell randomly and have a net order flow of x_N.

Trading occurs in two periods: In the first period, speculators and noise traders post buy or sell orders, which are executed according to a linear pricing function with unknown market depth. In period 2, prices are set by the market (which we can think of as being composed of long-term investors, market makers and others). The market does not know v, but can observe p_1 and therefore make inferences about v from observing p_1.

At $t = 1$, speculators are also informed about x_N, i.e. the direction and size of order flow from noise traders. This seeks to capture the phenomenon that speculators not only have some information about the value of the asset, but that they can also observe the current market sentiment – and hence whether their private information is in line with this sentiment or not. The main result about non-robustness to higher-order uncertainty emerges when we remove common knowledge about x_N by introducing small, idiosyncratic observation noise – see section 3.1 below.

There is no discounting and speculators are risk-neutral, and there are no budget or inventory considerations. However, we restrict the amount that any single trader can trade; such a restriction would otherwise arise naturally from assuming risk-aversion or budget restrictions.

1.2 Assumptions in Detail

The Players There is a continuum of speculators and a continuum of noise traders.[5] The speculators, who are perfectly informed about v, are ordered on the unit interval. They get a utility of $p_2 - p_1$ if they buy in period 1 and $p_1 - p_2$ if they sell in period 1, and 0 if they do nothing.

The uninformed noise traders decide randomly (not necessarily without correlation) whether to buy or sell one unit of the asset (because of liquidity needs, or irrational ideas about v). Denote by x_S the net order flow from the speculators, and by x_N the net order flow from noise traders in period 1. The order flow x_S from speculators is the result of their strategic trading decisions, while x_N is simply the result of some random process. Specifically, we assume that x_N is distributed according to a continous distribution F that is symmetric around 0 (so that $F\left(x'_N\right) = 1 - F\left(n - x'_N\right)$) and has density everywhere on an interval $[-n, n]$ where $n \in \mathbb{R}$. We restrict our attention to cases where $n > 1$, i.e. there could potentially be more trades from uninformed than from informed traders.

[5]The assumption that there are infinitely many traders need not be taken literally - it is simply meant to represent the fact that traders do not take account of the price impact of their trades, or assume that they cannot influence the price. The model works just as well in a discrete setting with finite numbers of speculators and noise traders, see appendix C.

The Pricing Function in Period 1 Let p_t be the market price of the asset in period t. At the beginning, the asset is trading at price p_0, which is the unconditional expected value:

$$p_0 = \frac{v_H + v_L}{2}. \tag{1}$$

Let $x = x_S + x_N$ be the total order flow from speculators and noise traders in period 1. These orders will be cleared by market makers (or a residual market) according to the linear pricing function

$$p_1 = p_0 + \lambda x \tag{2}$$

where λ is an unknown reverse market depth parameter. While λ is similar to Kyle's Lambda (Kyle, 1985) in the role it plays in the pricing function, note that it is exogenously given here. We assume that λ is uniformly distributed on the open interval $\left(0, \hat{\lambda}\right)$. To guarantee existence of well-behaved equilibria, we will impose a maximum condition on $\hat{\lambda}$ (i.e. a minimum condition on market depth) below.

The randomness of λ is mainly a technical assumption to make the mapping from order flow x to price p_1 noisy, so that p_1 is not fully revealing about x. If that were the case, the market would be able to observe p_1 and perfectly infer who had been trading (since speculators and noise traders exist in different measure). With a random λ, the size of the price change is still informative about the trading volume and therefore about whether informed or uninformed traders were trading more, but it is not fully revealing. The fact that speculators don't know λ also precludes the existence of spurious equilibria in which the speculators submit information by precisely encoding it into the price.

The randomness of λ can also be understood differently if we recall that the mass of speculators is normalized to one. Without changing the model, we could fix market depth at a constant, and assume that the mass of noise traders is given by λn and that of speculators by λ, which would leave the pricing function unchanged. Now the market depth would be known, but the mass of speculators would be unknown instead.

The Market in Period 2 In period 2, the market is an intelligent player, who has to set a price p_2 at which it is willing to buy or sell any quantity of the asset. I assume that it gets a utility of $-\left(v - p_2\right)^2$, so that it will always always maximize utility by setting $p_2 = \mathrm{E}\left[v \mid p_1\right]$. We can think of the market in period 2 as a large number of rational long-term investors, market makers and the like, who are in Bertrand-style competition and therefore make zero profit and are willing to buy or sell the asset for its expected value.

Restriction to Trade Size The speculators in period 1 can only buy or sell one unit of the asset each. The main intuition of this assumption is that the market is large compared to any single speculator. In the context of this model it is also a technically desirable assumption, since perfectly informed speculators with no trading or budget restrictions would otherwise have an incentive to trade arbitrarily large quantities and completely correct the price (as there is no fundamental risk for them). Just like in Glosten and Milgrom (1985), our focus is on the informational content of trades, not on their size.

With the introduction of fundamental risk or agency concerns, a similar size restriction would emerge endogenously. It also does in the real world: Even a trader who is absolutely sure of himself will normally not be allowed to trade very large quantities.[6]

All traders (speculators, noise traders and the market) are free of inventory considerations. Speculators can either buy one unit in period 1 and then sell it in period 2, or sell in period 1 and buy back in period 2, or they can abstain from trading at all. Selling and later buying back can also be thought of as a short sale (which has an inherent short horizon, even if we were to assume that speculators were not short-term interested). The market in period 2 is willing to trade any number of units at a fixed price.

Perfectly informed Speculators The speculators in this model are perfectly informed not only about the true value of the asset, but also about how the noise traders are (overall) trading. We could think of x_N being as the market sentiment, i.e. the current (irrational) movement of prices or the direction of the current mispricing.[7] The speculators learn v and x_N at the beginning of period 1.

The speculators do not know the inverse market depth λ, the other source of noise in the model. The noise in λ mostly serves to reduce the informativeness of p_1 such that p_1 doesn't fully reveal who has been trading in which direction (and thus reveal v). The market only knows the probability distributions of v and x_N, and observes p_1 at the beginning of period 2 before deciding which price to offer.

Timing of the Model The timing is shown in figure 2.

While the market behaves rationally in using all information that is contained in p_1, it is conceivable that it could also condition on order flow in period 2 when speculators liquidate their holdings. In particular, it could act similar to the market makers of Glosten and Milgrom

[6] Securities trading companies usually institute rules that limit trading by any one trader, similar to the trade size restriction of this model. Several scandals of "rogue trading" in recent years have highlighted the importance of such restrictions by illustrating the damage than can be done if they are circumvented.

[7] Cf. the discussion on insider trading by Leland (1992), who works with a similar assumption, and the "private learning channel" that speculators have in Cespa and Vives (2015).

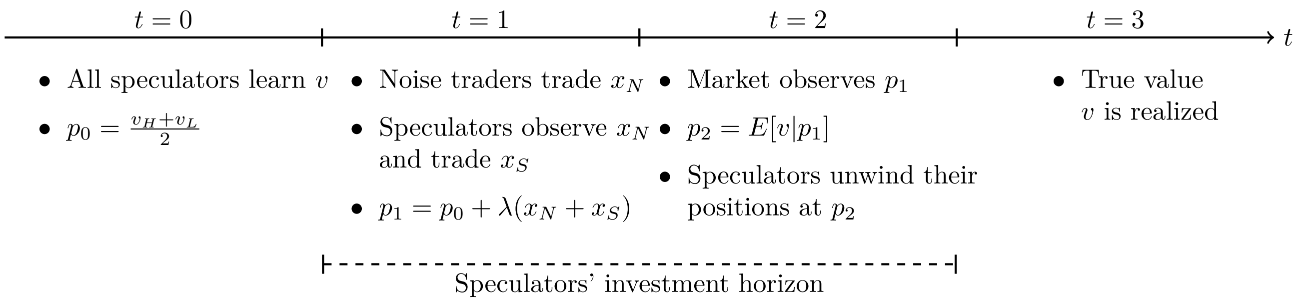

Figure 2: The timing of the model. The dashed line shows the speculators' investment horizon, which does not stretch to the realization of fundamental value in period 3 as speculators need to unwind their position in period 2.

(1985) and adjust p_2 conditional on whether it receives buy or sell orders. But the assumption that all speculators liquidate their holdings in period 2 is merely a simplification. In reality, many or even most traders are short-term oriented not because they have to liquidate their holdings every few days or weeks, but because their holdings get evaluated, by themselves or their superiors, *at market prices* in short time intervals. For their motivation and strategic choice, this is equivalent to a world in which they had to completely sell off and rebuild their portfolio frequently – but it does not per se allow the investors in our model to deduce any information about the order flow from the orders they face in period 2.

2 The Adjustment Equilibrium under Common Knowledge

We start by analyzing the game without higher-order uncertainty, where there exists an equilibrium in which p_1 and p_2 adjust to v on average. In this equilibrium, the market in period 2 assumes that p_1 is informative about v. In particular, it assumes that if $p_1 > p_0$, it is more likely that $v = v_H$ and vice versa. p_2 is set accordingly. If $|x_N| < 1$, the speculators can therefore influence p_2 by their trading, and they buy if v_H and sell otherwise. If $|x_N| \geq 1$, however, whether p_1 is above or below p_0 is determined by the direction of the noise trading, and speculators cannot influence p_1 sufficiently. It is then optimal for them to just trade in the same direction as the noise traders.

The market adjusts its expectation of v according to the function $p_2(p_1)$, which takes the behavior of the speculators and the distribution of x_N into account. If $p_1 > p_0$, for example, they know that overall order flow in the first period was positive, and that therefore either $|x_N| < 1$ and $v = v_H$, or that $|x_N| \geq 1$ and the speculators just followed the herd. The existence of the equilibrium is assured by a maximum condition on inverse market depth,

	v_L	v_H
$x_N \geq 1$	Buy	Buy
$x_N \in (-1, 1)$	Sell	Buy
$x_N \leq -1$	Sell	Sell

Table 1: Equilibrium strategy of the speculators. Only trading for $x_N \in (-1, 1)$ is informative.

which guarantees that it will always be optimal for the speculators to follow their equilibrium strategy.

Proposition 1. *(Adjustment equilibrium) It is an equilibrium if every speculator follows the strategy given by table 1 and the market sets* $p_2 = \pi(p_1)v_H + (1 - \pi(p_1))\, v_L$*, where* $\pi(p_1)$ *is the belief of the market that* $v = v_H$*, given* p_1*. This is under the condition that market depth is sufficient, i.e.*

$$\hat{\lambda} \leq \phi \frac{(v_H - v_L)}{n+1}. \tag{3}$$

The precise expressions of $\pi(p_1)$ *and* ϕ *are given in the proof.*

The intuition of the proof is the following: If speculators follow their equilibrium strategies, p_1 will contain some information about v. The function $\pi(p_1)$, which takes account of the distributions of x_N and λ, gives the probability (and hence the equilibrium belief of the market) that $v = v_H$ for every p_1. Since all possible prices occur in equilibrium, we do not need to consider out-of-equilibrium beliefs.

The speculators, on the other hand, will make an expected profit by following their equilibrium strategies, since the price movement in period 1 is always small enough (if market depth is sufficient, which is where the maximum condition on $\hat{\lambda}$ comes from). In particular, it is always either $p_0 < p_1 < p_2$ or $p_0 > p_1 > p_2$. Because a single speculator has only limited influence on p_1, no single speculator has an incentive to deviate. If a speculator would deviate from his equilibrium strategy, he would make a loss equal to the profit of his equilibrium strategy. The maximum condition on $\hat{\lambda}$ in (3) guarantees that there is no "overshooting" in expectation, i.e. if all speculators buy in period 1, p_1 still doesn't rise above v.

Figure 3 shows an exemplary price path in the adjustment equilibrium, where noise is small (i.e. $|x_N| < 1$). The speculators then face a coordination problem: They can either all buy or all sell, which will place p_1 either above or below p_0. In both cases they make a profit, and both cases constitute an equilibrium of their coordination game. In the Nash equilibrium, however, the market must optimally extract information from p_1, which is only the case if speculators trade towards the fundamental value v.

If $|x_N| \geq 1$ the speculators cannot influence whether p_2 will be above or below p_1, since they cannot neutralize the noise trading and there is no coordination game among them. It is

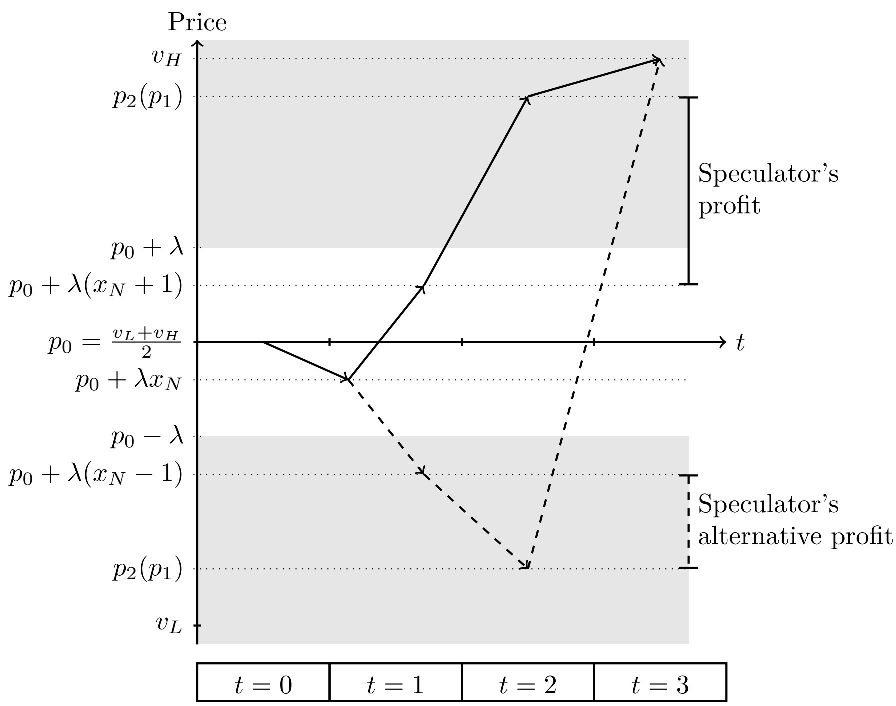

Figure 3: The equilibrium price path in the adjustment equilibrium for a given set of parameters, where $v = v_H$ (asset value high) and $0 > x_N > -1$ (noise traders sell the asset). Noise trading is small, i.e. noise traders do not push the price outside the white area in the center. All speculators buy the asset, thus pushing the price to $p_1 = p_0 + \lambda(x_N + 1)$. The market observes $p_1 > 0$ and sets $p_2(p_1) > p_1$, so that speculators make a profit.

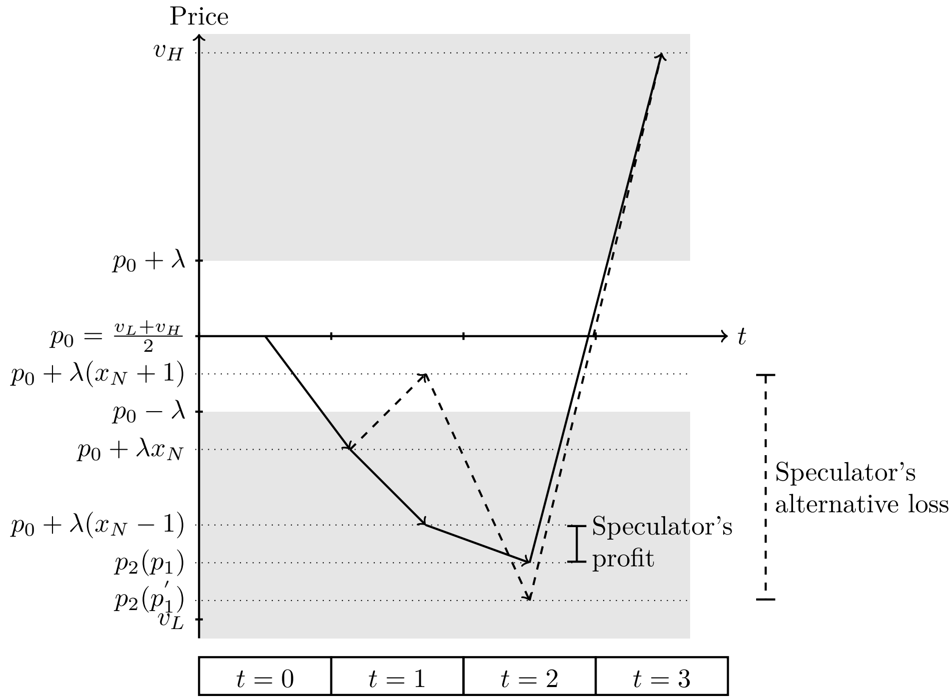

Figure 4: Another equilibrium price path in the adjustment equilibrium. Now $x_N < -1$, i.e. noise traders push the price into an area (given by the gray shade) where it is the dominant action for speculators to sell. Unlike in the previous figure, the coordination game between speculators does not have multiple equilibria and thus no information submission is possible.

dominant for them to follow the herd of noise traders – regardless of whether it is right or wrong. If the noise traders are wrong, that means that the speculators will drive the market price further away from its correct value even though they know better, and even though the investors would gladly enrich them by buying the asset at a more correct price. Figure 4 shows such a price path. The price gets pushed too far away from p_0 (into the grey area), so that the speculators are not able to move it above p_0 again. Once noise trading has pushed the price into the "grey area" of the graphs, p_1 and p_2 will not return to the informative "white area" and are therefore uninformative.

De Long et al. (1990) describe a similar effect when they consider "noise trader risk": In their model, rational and informed arbitrageurs in an overlapping-generations model could correct mispricings that arise through noise trading. But since arbitrageurs are short-lived and the market could get even more irrational (noise trade is randomly distributed), they refrain from fully correcting the mispricings. In my model, the direction of the noise order flow (and therefore also the direction of the mispricing in the next period) is known to the speculators, and they can therefore choose to trade against their information and therefore avoid the noise trader risk. They drive prices further away from fundamentals while doing so, as in models of speculative bubbles such as Abreu and Brunnermeier (2003).

3 Higher-Order Uncertainty

3.1 The Adjustment Equilibrium without Common Knowledge

Now consider the problem of the speculators in the adjustment equilibrium described above and take the market's strategy as given. As we have seen above, for $x_N \in (-1, 1)$ the speculators are playing a coordination game with two equilibria: They can either all sell or all buy; in either case they make a positive profit and no speculator would optimally choose a different action (cf. figure 3 on page 12). Only trading in the correct direction (buying for v_H and selling for v_L) can be part of a Nash equilibrium where the market optimally plays its equilibrium strategy.

This coordination among speculators works under the assumption that x_N is common knowledge, i.e. the ratio between informed and noise traders is common knowledge among the informed traders. Given that both v and x_N are information that is not publicly available to everyone (otherwise the market would be fully informed), it is plausible to consider what happens if x_N is known very precisely to the speculators, but not common knowledge. We will see that in this case, no set of values of x_N remains for which the speculators ever trade on their information. Their worries about other speculators' information about x_N (and their worries about other speculators' worries and so on) lead them to completely disregard any fundamental information, and the adjustment equilibrium collapses.

We loosen the common knowledge assumption in a way that is similar to the canonical models of Carlsson and van Damme (1993) and Morris and Shin (1998). Assume that instead of learning x_N, every speculator i observes some ω_i. All ω_i are independently drawn from a uniform distribution on $[x_N - \varepsilon, x_N + \varepsilon]$, i.e. an interval of length 2ε around the true x_N, with $\varepsilon > 0$ but small. Every single speculator will then know after observing ω_i that $x_N \in [\omega_i - \varepsilon, \omega_i + \varepsilon]$. But about the signal of another speculator j he will only know that $\omega_j \in [\omega_i - 2\varepsilon, \omega_i + 2\varepsilon]$, and he only knows that j believes that $x_N \in [\omega_i - 3\varepsilon, \omega_i + 3\varepsilon]$, and so on. Then, even if the observation of every single speculator is extremely precise, it is only common knowledge that $x_N \in [-n, n]$ – which is identical to the prior. We are interested in the case where $\varepsilon \to 0$, i.e. all speculators are arbitrarily well-informed about x_N, but lack common knowledge of it. In this slightly modified game, we can show the following proposition:

Proposition 2. *Assume that the market follows a strategy where $p_2 > p_1$ if $p_1 > p_0$ and $p_2 < p_1$ if $p_1 < p_0$. Then any rationalizable strategy of the speculators' coordination game has* *the property that all speculators buy if they observe $\omega_i > \varepsilon$ and sell if $\omega_i < -\varepsilon$. For $\varepsilon \to 0$, this gives a uniquely rationalizable equilibrium where all speculators buy if $\omega_i \geq 0$ and sell otherwise.*

Note that this is not a canonical application of the global games refinement, since the speculators' coordination game is not supermodular: Once p_1 is on the right side of p_0, every additional speculator who trades *decreases* the profits of the other speculators.[8] Still, it is possible to show that the above strategy is uniquely rationalizable, as the dominance regions where it is always optimal to buy or sell "infect" the undominated region where there were multiple equilibria in the complete information game.

Intuitively, every single speculator reasons along the following lines:

> I know that x_N is within a small interval around my observation ω_i. If ω_i is in $(-1, 1)$, it is my best guess that all speculators together could overcome the noise so that p_1 correctly reflects our private information. I also know that the other speculators have a very precise idea about x_N—but my knowledge about their knowledge is a little less precise than my own knowledge about x_N. If I consider my knowledge about their knowledge about my knowledge, it gets even less precise.
>
> In particular, if ω_i is very close to 1, I think it is very likely that many other speculators have received a signal above 1 and will therefore play what they believe is the dominant strategy of buying. So I should buy if I observe ω_i very close to but below 1, regardless of what my information about v.
>
> The others will reason the same way, so that if I observe ω_i somewhat less close to 1, I know that many others will observe a ω_j closer to 1, and buy for the reason outlined above. Such contagion carries on, and vice versa from ω_i close to -1. So I will choose to sell if $\omega_i < 0$ and buy if $\omega_i \geq 0$, and disregard my private information about v.

Figure 5 depicts the intuition of the contagion argument in a graph similar to the ones above. The proof formalizes this iterative reasoning by defining an elimination process that starts with the set of all possible strategies and then removes, in each step, strategies that are never a best response to any other strategy in the remaining set. In this way, the proof is in the vein of the original work by Carlsson and van Damme (1993) while taking up and modifying some ideas from Frankel et al. (2003).

[8] As an illustration, consider the case where $v = v_H$ and $x_N = -0.2$. If all other speculators increase their probability of buying from 0.5 to 0.7, buying becomes more attractive. If they increase their probability of buying from 0.8 to 0.9, however, buying becomes *less* attractive. Thus the game is not supermodular.

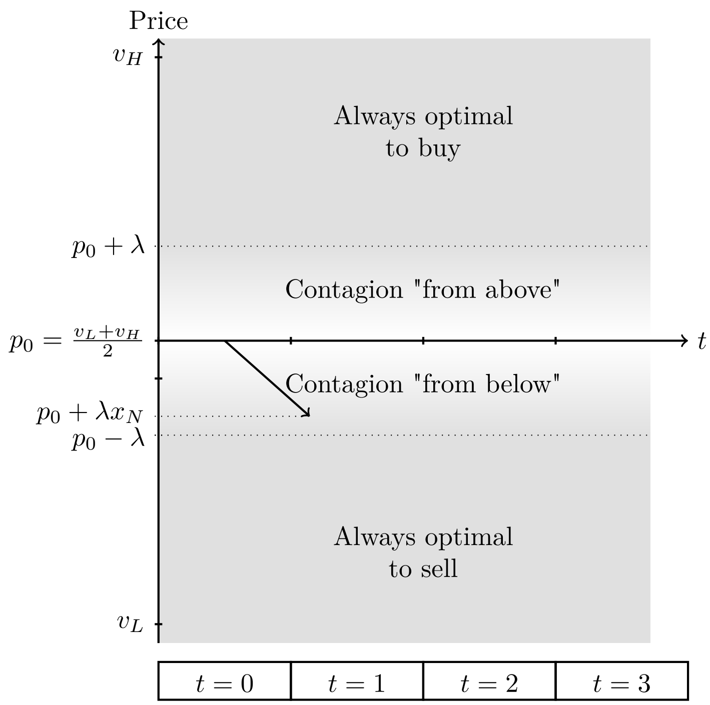

Figure 5: Contagion in the absence of common knowledge. Even though $x_N > -1$, the speculators' coordination game no longer has multiple equilibria. A speculator observing $x_N > -1$ but close to -1 is worried that others might believe that $x_N < -1$, or that others might believe that others believe this, etc. This contagion carries on, such that selling becomes optimal for all $x_N < 0$ and buying for all $x_N > 0$.

If the trading of all speculators is only dependent on ω_i and independent of v, p_1 will actually be completely uninformative about v. Hence there is no Nash equilibrium of a game in which x_N is not common knowledge where the market treats p_1 as informative.

3.2 A Non-Adjustment Equilibrium

If informative trading is not rationalizable, which equilibrium can we expect the whole market to be in? It depends on the perspective we take on the role the market in period 2. If we see it as a non-thinking actor who simply follows the decision role laid out in proposition 1, the story ends here: Without common knowledge, speculators never trade informatively, and still the market sets p_2 as if p_1 were informative. All price movements in p_1 and p_2 are pure noise.

Almost the same happens if we treat the lack of common knowledge as an unlikely event, or an event that the market does not expect. Since the market cannot observe whether speculators trade informatively or not, it would continue in treating p_1 as informative.

But if we are to take the role of the market as as rational actor seriously, we must assume that in equilibrium it is not "fooled" by p_1 and correctly believes that p_1 is uninformative, in which case it would set $p_2 = p_0$, the prior. Consequently, the trading strategy of the speculators derived in proposition 2 would no longer be optimal. Trading on v, however, does not become optimal. Instead there exists a different equilibrium:

Proposition 3. *(Non-adjustment equilibrium) It is an equilibrium if speculators with probability* $\min\{|\omega_i|, 1\}$ *either buy if* $\omega_i < 0$ *or sell if* $\omega_i \geq 0$ *(and neither buy or sell with the complementary probability), and the market believes that* p_1 *is completely uninformative and therefore sets* $p_2 = p_0$.

If the market believes p_1 to be uninformative, the speculators already know that $p_2 = p_0$ and the only gain they can make is by providing liquidity to noise traders. Since this means they do not act on their information about v, the market is correct to believe that p_1 is uninformative.

This equilibrium actually exists independently of whether there is common knowledge or not, as there is no strategic complementarity in the speculators' actions. No player has an incentive to deviate from their equilibrium strategies: The market would not benefit from assuming that prices contain information, and the speculators cannot gain from unilaterally (or as a group) submitting information (and thereby driving p_1 away from p_0). In this interplay of "not talking" and "not listening", the equilibrium is similar to the "babbling equilibrium" of cheap-talk games (Farrell and Rabin, 1996). There, the sender randomizes between messages such that her message has no correlation to her private information, and the receiver ignores any messages by the sender. This constitutes an equilibrium, albeit (when it comes to everyday communication) perhaps not a plausible one.

The non-adjustment equilibrium is also similar to uninformative equilibrium of Benhabib and Wang (2015), and it is an extreme case of the less informative equilibrium of Cespa and Vives (2015). In both cases, the uninformativeness of the equilibria also emerges through short-term constraints in the models.

3.3 Without Common Knowledge, the Market Cannot Be Informationally Efficient

The rationalization result derived in proposition 2 is clear and general: Consider any equilibrium of the complete-information game in which the market treats p_1 as informative by setting $p_2 > p_1$ if $p_1 > p_0$ and vice versa. Clearly, the adjustment equilibrium and any small perturbation of it belong in this class. Now relax the common knowledge assumption about x_N by introducing the smallest seed of doubt about whether the other speculators are making the same observation as you. For ε very small, which is the case we are interested in, the speculators are still generically 100% sure that x_N is small, they are almost equally sure that everybody else knows that x_N is small, and so on ... but not ad infinitum. They no longer have common knowledge about this fact, and this small seed of doubt means that the strategies that would constitute the adjustment equilibrium are no longer rationalizable – which means that they cannot be part of a Nash equilibrium. Without common knowledge, therefore, no equilibrium can exist in which the market correctly beliefs that p_1 is informative.

Recall also that this result requires no assumptions on the shape of F, the distribution of x_N, except that it is continuous, symmetric and has density everywhere. In particular, this means that F could be shaped such that an arbitrarily large mass of F is inside $[-1, 1]$. Then the noise was almost always small and trading in the adjustment equilibrium would be almost fully revealing. Even in this model, contagion would occur and the adjustment equilibrium would not exist. The frequency of large $|x_N|$ is therefore not important for the relevance of the model – the pure possibility that x_N is outside $[-1, 1]$ is sufficient.

4 Discussion

4.1 When Does Contagion Occur, and What Can We Learn from It?

The main result of this paper is quite stark, as it shows that informative trading only happens if speculators have common knowledge of all model variables. Common knowledge is an exacting requirement that is often unlikely to be met in reality, especially given that the very idea of information aggregation is that the information is not known to everyone. I would therefore like to point out which assumptions of the model are crucial for obtaining the contagion result. That allows us to make predictions about which real-world conditions promote or preclude informational ineffiency through contagion.

In general, it can be said that:

- Contagion does not occur if speculators do not have *short-horizons*, but live until period 3. However, if we endow only a few speculators with long horizons, there is no qualitative change – the remaining speculators are still subject to the same contagion.
- While contagion occurs regardless of any specific assumptions about the likelihood of certain actions by *noise traders*, contagion does not occur if speculators do not consider a large enough set of trading behavior ex ante conceivable.
- If speculators have no knowledge about x_N, contagion does not happen (since there is no information about the beliefs of others that would be reason for worry). Introducing higher-order uncertainty of other variables than x_N does not appear to change anything.

Short horizons A central assumption of the model is that the information about the value v is known only to short-term speculators. This is not to suggest that *all* information arrival at financial markets works in this way, but just that the theory of contagion only applies to situations where this is the case. In general, however, it does not seem a wholly unreasonable assumption that speculators could be better informed than some long-term investors. Just consider that most professional money managers would count as “speculators” in the context of this model if we consider sufficiently long time periods—a few weeks, say, or a quarter. Few of them are allowed and capable of raking up massive losses over such a time frame even if they claim to have superior knowledge that will in the end be vindicated.

Empirical evidence suggests likewise that a large proportion of stock positions are opened for a very limited amount of time, with the expectation of making a profit in less time that it takes to see two quarterly earnings reports. The average holding period of stocks in the

United States is three to four months—not even enough to receive a full dividend payment, let alone profit from long-term business or macroeconomic developments.[9] And even where assets are not bought and sold within days or seconds, those who decide about trading them have their performance evaluated at market prices at very short intervals. If a trader buys an asset at time t for the price p_t, it does not matter to him whether he sells the asset at $t+1$ and it contributes $p_{t+1} - p_t$ to his cash holdings, or whether he still holds it at $t+1$ and it contributes $p_{t+1} - p_t$ to the overall appreciation of his holdings since t.

If all speculators lived until period 3, they would always trade on their information and no contagion would occur. But if we start out with the model in this paper and add a number of long-lived informed investors, the result is robust – up to a point. Consider, for example, a modified model in which there is a measure $\mu < 1$ of informed investors, who always buy if $v = v_H$ and sell otherwise. This would be akin to shifting the distribution of x_N by μ, so that noise trading is given by $\hat{x_N} = x_N \pm \mu$, depending on v. As long as the distribution of $\hat{x_N}$ has density both below -1 and above 1 so that it reaches into the dominance regions, contagion occurs.

In general, the contagion result is remarkably robust to small changes in the payoffs of the speculators. This matters, for example, if we assume that speculators get a small additional payoff from trading in the "right" direction, because there was an exogenous chance that they could live longer. To see why this is the case, note that the payoff structure of a speculator looks like this (+ denotes positive profits, − negative profits):

Speculator's action at $t=1$:	Result: $p_2 > p_1$	$p_2 < p_1$
Buy	+	-
Sell	-	+

A speculator that decides whether to buy or to sell will only ever compare two values in the same colum, since there is no uncertainty in the rationalization argument as to which way the price will move. (Assuming that a speculator lives until period three with a certain probability would mean that he plays the game given by the matrix above with a certain probability, and another game otherwise.) The chain of rationalizability arguments that led to the contagion result therefore only relies on the fact that the values in the main diagonal are larger than the other values in the same column. As long as the intrinsic payoff of trading on v, and the probability of being long-lived, are small enough, the contagion result obtains.

[9]The "World Bank Financial Development Indicators" show stock market turnover ratios, which is the inverse of average holding period. In the United States in 2008, for example, trade volume was 4.35 times as high as total market capitalization. Since this is the mean holding period and the distribution is truncated at 0, the median holding period is probably much lower.

Noise trading Some authors (e.g. Dow and Gorton, 1994, p. 825) argue that the presence of noise traders has to be explained. But the absence of noise traders would mean that all traders, at all times, act rationally to maximize their expected payoff from trading. There are two main types of traders for whom that does not apply. Firstly, substantial research on behavioral finance has shown that traders, institutional or private, fall prey to a large number of irrational biases. Secondly, even a rational trader might find it optimal to sell an asset (whose price he expects to rise) for liquidity reasons – for example when he needs to access his savings to retire or pay unforeseen expenses.

Once we accept the assumption that there are noise traders in the market, the question naturally arises whether additional assumptions about the actions of noise traders are necessary. Note, however, that the only two assumptions about the distribution F of x_N that are used in the proof of proposition 2 are (a) that the probability density function of F is continuous and (b) that F has density everywhere on $[-n, n]$. It is therefore only required that speculators consider any order flow from noise traders conceivable – they don't have to think it likely. In fact, if we assume that x_N was normally distributed around 0, we could make the standard deviation of this distribution arbitrarily small without in any way containing the contagion. The distribution of the order flow from noise traders could be so concentrated that speculators were almost sure that x_N was in $(-1, 1)$. In that case, trading in the common-knowledge equilibrium (proposition 1) would almost always be informative and $\mathbb{E}[(v - p_2)^2]$ would get arbitrarily small in this equilibrium. Yet as soon as we introduce the smallest higher-order uncertainty about x_N, contagion carries through all the way and informative trading is not rationalizable.

The information available to speculators When I have considered higher-order uncertainty in this paper, I have limited this uncertainty to the realization of x_N and continued to assume common knowledge of v. This begs the question of what would happen if there was also higher-order uncertainty of v, so that every speculator would worry also about other speculators' belief about v. Could there be a similar contagion of beliefs that might even restore dependence of the speculators' actions on v?

The answer is no, at least in a setup like in this paper where there are no possible values of v for which any action by the speculators would be dominant. There is simply no possible belief about v to "start" a contagion of beliefs. In the case of uncertainty about x_N, this is the belief that another speculator might think that another speculator might think etc. that x_N is so large or so small that buying or selling is the dominant action. Without such "dominance regions", there can be no contagion. In the teminology of Weinstein and Yildiz (2007), the "richness assumption" fails on v, since the parameter space of v is not rich enough to contain dominance regions.

It is possible to think of situations where there are conceivable fundamental values that make buying or selling dominant. If, for example, speculators know that they live until period 3 with some probability, and v is extremely large with some probability, they might find it dominant to buy the asset. This reinforces the point (made above) that a sufficient long-term orientation of the speculators can break the chain of contagion.

Finally, a crucial requirement of contagion is that speculators actually have an observation of x_N, since it is the worry about other speculators' beliefs of x_N that keeps the contagion alive. If speculators are completely unaware of x_N and only observe v, their only consideration is whether x_N is outside the dominance regions with enough probability to make informative trading profitable. We therefore have the seemingly curious result that contagion fails both if speculators have less and more information (i.e. no information or common knowledge of x_N). If speculators fall prey to contagion, the market would function much better if they did not have access to information about the market sentiment. The sort of coverage that is most beloved by newspapers and tv stations the world over – "Panic at NYSE! Euphoria as Asian Markets Open!" – can thus have a hugely detrimental effect by giving informed speculators information about the noise in the market without generating common knowledge about it. Common knowledge would only be generated if all speculators followed the same news sources, had common knowledge about this fact, and also had common knowledge about the fact that they all understand the news in the same way – a tall order. Ultimately, the contagion argument rationalizes a folk argument among economists: The hype and sensationalist coverage surrounding financial markets can magnify the "psychological moods" of the market and eradicate cool-headed, rational trading – and everybody would be better off without it.

4.2 Examples of the Mechanism at Work

As an example of the paralysis of informative trading described in this paper, consider the so-called Dot-com bubble in the late 1990s and early 2000s. In the context of this model, we could think of internet stocks as being worth either v_L ("most of these companies will never make a profit") or v_H ("they will change the economy forever and be hugely profitable"). Many market participants did not know which was the case, but because v_H was extremely large their unconditional prior $\frac{v_L+v_H}{2}$ was also large. The uncertainty was large enough to make it plausible that it would only be resolved quite far into the future (what if internet companies needed to grow for a decade before turning huge profits?), far beyond the investment horizon of most investors.

Many sophisticated fund managers knew that internet stocks were overvalued, i.e. that

$v = v_L$.[10] But to coordinate on an informative sell-off of internet stocks, they would need common knowledge about the fact that there were enough informed traders. As we have seen, it does not matter how large the number of informed speculators is in relation to the number of noise traders. Without common knowledge, the sheer possibility that there could be many noise traders infects everyone's beliefs, despite the fact that all speculators know this not to be the case. So even a well-informed and sophisticated fund manager who knew that stocks were overvalued, and who knew that there were enough others to support a sufficiently large sell-off, feared that others would not sell because they feared that still others would not sell, and he would therefore not sell himself.

A similar pattern emerges when we consider what is perhaps the most notorious market movement in history, the "Great Crash" of 1929. The crash was by no means unexpected, as many experts had come to realize throughout 1929 that stock prices were unsustainably high. Galbraith (1954, ch. 2) describes the uneasiness in regulatory circles and the various attempts to deflate the bubble, and also documents prescient warnings by well-known bankers, financial services and the *New York Times*. But without common knowledge about the fact that informed traders could outnumber noise traders, there was no informative sell-off.

1929 also offers a glimpse into how an equilibrium shift from the non-adjustment to the adjustment equilibrium can occur when common knowledge is generated. On October 24 ("black thursday"), prices fell suddenly and violently by nearly 13%. They swiftly recovered (the closing was only 2.1% down that day), but the event had shown that there were many traders in the market willing to sell. What was even more important was that, since everybody could reasonably assume that everybody else would follow the market closely enough to notice such an event, the preponderance of informed traders was now also common knowledge. In the following days, despite no substantial economic news (cf. Shiller, 2000, p. 94), informed market participants could now coordinate on selling, and the Dow fell over 23% in two days.

As an example of (unprofitable) out-of-equilibrium behavior, consider the spread between Royal Dutch and Shell stocks in the late 1990s. The stocks were trading at different exchanges, but prices should have been at a fixed proportion, because cash flows were paid in a fixed proportion. Instead, there was a spread that was quite stable around 8% (cf. Froot and Dabora, 1999). It appears that the market was in a non-adjustment equilibrium where the many traders who were aware of the unreasonable spread could not coordinate on trading to narrow it, since they didn't have common knowledge about their on combined strength in the market. When the hedge fund Long-Term Capital Management (LTCM) began to trade against the spread in 1997, there was sufficient trading in the opposite direction to maintain

[10] See for example the discussion in Abreu and Brunnermeier (2003, p. 175). Brunnermeier and Nagel (2004) document that hedge funds were heavily invested in tech stocks, and argue that this was not because they believed prices to be reasonable.

	A is chosen	B is chosen
Vote for A	1000	0
Vote for B	0	1

Table 2: Payoffs of committee members, assuming that A is the better option. (If B is the better option, the payoffs 1 and 1000 change places.)

the spread – as we would expect from the model, as informed speculators had settled on trading against changes in the spread instead of betting on it to close.[11]

4.3 An Application to Voting in Committees

The main theoretical insight of this paper can be applied to other settings besides financial markets. This section sketches an application to voting in committees.

A decision maker has to decide between two options, A and B. One of them is better than the other; the decision maker gets a payoff of 1 for choosing the better option and 0 otherwise. The decision maker does not know which option is better, but he gets help from a committee of experts, who all know which option is better. As an example, consider a department head who has to choose between two applicants for an academic position, or an authority that has to decide whether to approve a new drug.

Committee members (experts) are on the unit interval. They strongly prefer that the better option gets chosen, but they also get a small payoff if they back the worse option and it gets chosen. They get nothing if they vote for a losing option. Their payoffs, assuming that A is the better option, are shown in table 2.

The timing is as follows:

1. All experts observe whether A or B is better.

2. All experts decide simultaneously whether to vote for A or B.

3. The decision maker observes the proportion a of experts that voted for A and makes a choice.

There is a simple and robust Nash equilibrium where all experts vote for the better option, and the decision maker implements the option that gets the majority of the votes. Now assume that there are some outside forces that influence the vote count or its transmission, so that the decision maker observes $\hat{a} = a + \theta$ instead of a, where $\theta \sim \mathrm{N}(0, \sigma)$. Let σ be small. This could also be thought of as influencing the decision maker himself, for example a bias in his perception or another source of information that he has besides the expert committee.

[11] The managers at LTCM, however, were “mystified” – cf. Lowenstein (2000, p. 148).

One could think of a pharmaceutical company lobbying for or against the approval of a new medication, or an unknown bias on the side of the department head for one candidate or another. For σ sufficiently small, the equilibrium strategies remain the same, and the worse option is only chosen in very few cases (only if $\theta \notin (-0.5,\, 0.5)$).

Now assume that all experts observe θ at the same time that they observe which option is better. The equilibrium of this game is almost the same: Experts vote for the better option if $\theta \in (-0.5, 0.5)$, vote for A if $\theta \geq 0.5$ and for B if $\theta \leq -0.5$. The decision maker chooses A if $\hat{a} \geq 0.5$ and B otherwise, and the better option is chosen if $\theta \in (-0.5, 0.5)$, which is usually the case since σ is small by assumption.

Now, however, assume that instead of observing θ perfectly, each expert i observes signal ω_i which is i.i.d. uniformly distributed on $[\theta - \epsilon,\, \theta + \epsilon]$, and we are interested in the case where $\epsilon \to 0$. Experts are still very precisely informed, but lack common knowledge about θ. They still have common knowledge about which option is better.

We can now use the same technique as in the proof of proposition 2 to show that the dominance regions infect the multiple-equilibria region. Taking the decision maker's strategy as given, there is no rationalizable strategy for any expert that conditions voting on which option is best. The informationally efficient equilibrium gets destroyed completely, and instead we get an equilibrium in which each expert votes A if $\omega_i > 0$ and B otherwise, and the decision maker follows their recommendation. The options are chosen randomly depending on the realization of θ, and the better option gets chosen only half of the time. This is despite the fact that it is common knowledge among the experts which option is better, the "transmission noise" $|\theta|$ is small with very high probability, and experts as well as the decision maker prefer the better option. It is enough that the experts consider it remotely conceivable that the department head would ignore their recommendation (i.e. that the distribution of θ has density outside of $(-0.5,\, 0.5)$) to make them use their very precise observation of her bias to always choose the candidate to which she is leaning. This anticipatory obedience even occurs if her bias is almost always negligibly small.

The implications of this model are similar to recommendations for financial markets above. If experts derive sufficient intrinsic motivation from voting for the correct option, the contagion collapses. No contagion occurs if we simply place 2 in the upper right, and -2 in the lower left field of table 2. Furthermore, experts should, if possible, be kept in ignorance of the biases that influence the decision maker: As we have seen above, the informationally efficient equilibrium continues to exists if we introduce θ but keep it hidden from the experts.

5 Conclusion

Perhaps the main reason for the triumph of market-based economic systems is that no other mechanism can transmit information about scarcity, efficiency and ability as realiably, fast and cheaply as the price mechanism (cf. Hayek, 1945). We live in a system of financial capitalism because financial markets are the ultimate way of transmitting information: Financial assets are standardized and fungible, all information other than prices is stripped away, information flow is immediate and transaction costs minimal. But the well-functioning of financial markets requires that they actually incorporate the information that is held by market participants.

This paper describes a mechanism that can destroy informational efficiency if traders only care about the short run and have knowledge about the irrational moods and passions of the market. Both assumptions are compatible with empirical observations. The effect of the latter assumption also supports the conclusion that the spread of rumors and ideas can be highly destructive even in a market that is mainly populated by rational traders. Rumors need a medium to spread, and accordingly Shiller (2000) hast pointed out that "the history of speculative bubbles begins roughly with the advent of newspapers."

The concrete uses of the model lie in providing advice on how to prevent belief contagion in financial markets (section 4.1) and explaining observed behavior (section 4.2). But the theoretical contribution goes beyond. As we have seen in section 4.3, contagion can destroy information aggregation in other settings if actions are strategic complements. Ultimately, the role of contagion in magnifying noise trading and detaching financial markets from fundamentals is only one application, if perhaps the most important, of the general theoretical insight. Contagion only requires that actions are strategic complements, and that people find it conceivable that the world would be in a state where each of them had a uniquely optimal action. Then, with even minimal higher-order uncertainty, contagion guarantees that for any state of the world, there is a uniquely rationalizable action. And crucially, as this paper argues, the signal that they condition their actions on need not be fundamental. It can be irrational ideas about the prospects of dot-com companies or the biases of a decision maker, but it might just as well be any other idea that is not ruled out by prior beliefs. A general pattern emerges by which higher-order uncertainty can detach outcomes from the fundamental variables that actually matter. Instead, behavior is determined by the spurious realizations of meaningless signals, purely out of the self-fulfilling belief that others are following these signals. There are connections to the theories of groupthink (Janis, 1972) and preference falsification (Kuran, 1997), which suggest other applications to political behavior, decision making in groups and the collection of knowledge in organizations.

Appendix

A Proofs

Proof of Proposition 1. **Part 1: The market has no incentive to deviate (and $\pi(p_1)$ is obtained by Bayes' rule).**

Assume that the speculators follow their equilibrium strategies and consider the case where $p_1 > p_0$. The market can then, from observing p_1, draw conclusions about v. Let $\pi(p_1)$ be the conditional probability that $v = v_H$ after observing a certain p_1, $\Pr(v_H|p_1)$.

It is

$$
\begin{aligned}
\pi(p_1) &= \Pr(v_H|p_1) = \frac{\Pr(p_1 \cap v_H)}{\Pr(p_1)} \\
&= \frac{\Pr(p_1 \cap v_H \cap |x_N| < 1) + \Pr\left(p_1 \cap v_H \cap x_N \geq 1\right)}{\Pr\left(p_1 \cap |x_N| < 1\right) + \Pr\left(p_1 \cap x_N \geq 1\right)} \\
&= \frac{\int_{-1}^{n} g\left(\frac{p_1 - p_0}{x_N + 1}\right)\, dF(x_N)}{\int_{-1}^{n} \left(1 + \mathbf{1}_{x_N > 1}\right) g\left(\frac{p_1 - p_0}{x_N + 1}\right)\, dF(x_N)}
\end{aligned}
$$

where g is the density of λ. Since $g\left(\frac{p_1-p_0}{x_N+1}\right) = \frac{1}{\hat{\lambda}}$ if if $0 < \frac{p_1-p_0}{x_N+1} < \hat{\lambda}$ and 0 otherwise, we can rewrite this as

$$
\begin{aligned}
\pi(p_1) &= \frac{\int_{\frac{p_1-p_0}{\hat{\lambda}}-1}^{n} dF(x_N)}{\int_{\frac{p_1-p_0}{\hat{\lambda}}-1}^{n} \left(1 + \mathbf{1}_{x_N > 1}\right)\, dF(x_N)} \\
&= \frac{1 - F\left(\frac{p_1-p_0}{\hat{\lambda}} - 1\right)}{2 - F\left(\frac{p_1-p_0}{\hat{\lambda}} - 1\right) - F\left(\max\left\{\frac{p_1-p_0}{\hat{\lambda}} - 1,\, 1\right\}\right)}.
\end{aligned}
$$

$\Pr(v_L|p_1)$ is the complementary probability $1 - \pi(p_1)$, so that the expected value of v given p_1 is $E\left[v|p_1\right] = \pi(p_1)v_H + (1 - \pi(p_1))\, v_L$. A similar argument applies to the case where $p_1 < p_0$. If $p_1 = p_0$, the price contains no information and p_2 should be set equal to the prior.

p_1 is between $p_0 - \hat{\lambda}(n+1)$ and $p_0 + \hat{\lambda}(n+1)$. For $x_N \in \{-n, n\}$, all possible p_1 occur with positive density, so that in equilibrium all possible p_1 occur with positive probability and there can be no out-of-equilibrium beliefs.

Part 2: Speculators make a positive profit in equilibrium.

Now assume that the market follows its equilibrium strategy. Consider the case where

$p_1 > p_0$, meaning that either $|x_N| < 1$ and $v = v_H$ or simply $x_N \geq 1$.[12] If they follow their equilibrium strategies, the speculators' buy orders will drive the price to $p_0 + \lambda(x_N + 1) > p_0$, and in period 2 all speculators will be able to sell their holdings at $\pi(p_1)v_H + (1 - \pi(p_1))\, v_L$. Their profit is then $\pi(p_1)v_H + (1 - \pi(p_1))\, v_L - p_0 - \lambda(x_N + 1)$, which can also be written as

$$\frac{F\left(\max\left\{\frac{p_1-p_0}{\hat{\lambda}} - 1,\, 1\right\}\right) - F\left(\frac{p_1-p_0}{\hat{\lambda}} - 1\right)}{4 - 2F\left(\max\left\{\frac{p_1-p_0}{\hat{\lambda}} - 1,\, 1\right\}\right) - 2F\left(\frac{p_1-p_0}{\hat{\lambda}} - 1\right)} (v_H - v_L) - \lambda\,(x_N + 1)\,.$$

Since every speculator knows x_N, the expected profit is

$$\mathbb{E}\left[\left.\frac{F\left(\max\left\{\frac{p_1-p_0}{\hat{\lambda}} - 1,\, 1\right\}\right) - F\left(\frac{p_1-p_0}{\hat{\lambda}} - 1\right)}{4 - 2F\left(\max\left\{\frac{p_1-p_0}{\hat{\lambda}} - 1,\, 1\right\}\right) - 2F\left(\frac{p_1-p_0}{\hat{\lambda}} - 1\right)}\right| x_N\right] (v_H - v_L) - \frac{\hat{\lambda}}{2}\,(x_N + 1)\,.$$

Since p_1 is increasing in x_N, $F\left(\max\left\{\frac{p_1-p_0}{\hat{\lambda}} - 1,\, 1\right\}\right)$ and $F\left(\frac{p_1-p_0}{\hat{\lambda}} - 1\right)$ are also (weakly) increasing in x_N. The expression therefore becomes minimal for $x_N = n$. If at this minimal point it is still non-negative, speculators make a positive expected profit in equilibrium; this is the case if

$$\begin{aligned}\hat{\lambda} \;&\leq\; \mathbb{E}\left[\left.\frac{F\left(\max\left\{\frac{p_1-p_0}{\hat{\lambda}} - 1,\, 1\right\}\right) - F\left(\frac{p_1-p_0}{\hat{\lambda}} - 1\right)}{4 - 2F\left(\max\left\{\frac{p_1-p_0}{\hat{\lambda}} - 1,\, 1\right\}\right) - 2F\left(\frac{p_1-p_0}{\hat{\lambda}} - 1\right)}\right| x_N = n\right] \frac{(v_H - v_L)}{n+1}. \\ &\leq\; \phi\frac{(v_H - v_L)}{n+1}\end{aligned}$$

This gives a minimum condition for market depth, which is simply given by the spread between high and low value, adjusted for the number of market participants and some adjustment factor ϕ that depends on the precise shape of F. If this minimum condition is fulfilled, speculators make a non-negative expected profit in equilibrium. Note that $\phi \in (0, 1/2)$ since the expression in the expectation is at least 0 (if $\frac{p_1-p_0}{\hat{\lambda}} > 2$) and at most $1/2$ (if $\frac{p_1-p_0}{\hat{\lambda}} \leq 2$), and both cases occur.

Part 3: No single speculator has an incentive to deviate from his equilibrium strategy.

Part 2 shows that every speculator has, after having observed v and x_N, a non-negative expected profit from following his equilibrium strategy. If his equilibrium action is to buy, then $p_1 - p_0 \geq 0$, and $p_0 - p_1 \geq 0$ if his equilibrium action is to sell. If he were to do nothing instead, his profit would be 0, which is not better. If he were to do the opposite, his profit would be non-positive, which is also not an improvement. All speculators hence optimally follow their equilibrium strategies. □

[12] An analogous argument applies where $p_1 < p_0$.

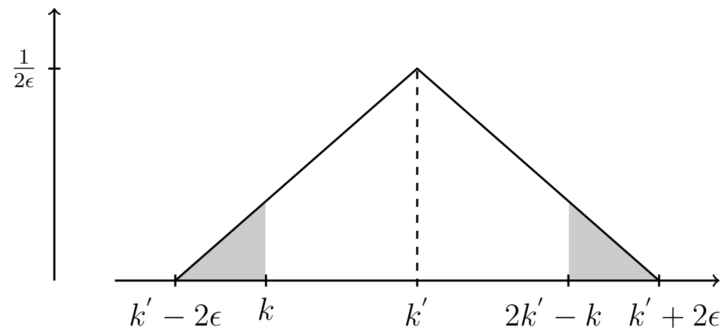

Figure 6: Illustration of the proof for the lemma. For small ϵ, the distribution of the signals of the other speculators is a symmetric triangular distribution around the own signal. Given that all speculators that receive a signal lower than k sell, a mass of speculators that is given by the shaded area on the left will always sell. Their sell orders will at least cancel out the buy orders by a mass of speculators given by the shaded area on the left, so that the maximum number of buy orders is given by the white area between the two shaded areas. If it should be undominated to buy with positive probability after receiving signal k', the white area would have to be larger than $-k' - \epsilon$. If $k' - k$ is below the upper bound given by the definition of $B(k, \epsilon)$, that is not possible.

Proof of Proposition 2. A strategy is a function $s(\omega)$, where $s : [-n-\epsilon, n+\epsilon] \to [0,1]$ gives the probability of buying, given any observation ω. Let Σ be the set of all strategies. Define the iterative-dominance function $\rho : \mathcal{P}(\Sigma) \to \mathcal{P}(\Sigma)$ where $\mathcal{P}$ is the power set. Given a set of strategies Σ', ρ returns a set of strategies $\rho(\Sigma')$ that is identical to the first one except that all strategies in Σ' that are never a best-reponse to any strategy in Σ' have been removed. Let $\rho^2(\Sigma) = \rho(\rho(\Sigma))$ and so on; a strategy s is rationalizable if $\forall n \in \mathbb{N} : \; s \in \rho^n(\Sigma)$.

What does $\rho(\Sigma)$ look like, where Σ is the set of all strategies? Clearly, no strategy that puts probability higher than 0 on buying for any $\omega_i \in [-n-\epsilon, -1-\epsilon]$ is in $\rho(\Sigma)$, since otherwise the speculator would be buying with positive probability even though he knows for sure that $p_1 > p_2$.

Let $B(k,\epsilon) = \left\{k' \in \mathbb{R} \middle| \left|k' - k\right| < 2\epsilon\left(1+\epsilon-\sqrt{(1+\epsilon)^2 + k + \epsilon}\right)\right\}$ be an open ball around k with a size that depends on k and ϵ. Note that the size of $B(k,\epsilon)$ is always below 4ϵ if $k \geq -1-\epsilon$. The following lemma establishes that we can use this ball $B(k,\epsilon)$ to exclude elements from $\rho(\Sigma')$ if no strategy that buys for k is in Σ'. The proof is illustrated in figure 6.

Lemma 1. *If Σ' contains no strategy that puts positive probability on buying for any ω_i with $-1-\epsilon \leq \omega_i < k < -\epsilon$, then $\rho(\Sigma')$ contains no strategy that puts positive probability on buying for any $\omega_i \in B(k,\epsilon)$.*

Proof. Consider the reverse, i.e. there exists a $k' \in B(k,\epsilon)$ such that there is a strategy

$s \in \rho(\Sigma')$ with $s(k') > 0$. (Only $k' > k$ is possible by assumption.) If speculator i gets the signal $\omega_i = k'$, he knows that $\int_{k'-2\epsilon}^{k} dH$ speculators will sell, with H being the distribution of the signals of other speculators conditional on receiving signal $\omega_i = k'$. For ϵ very small, this conditional distribution is approximately a symmetric triangular distribution on $[k' - 2\epsilon,\, k' + 2\epsilon]$, and therefore $\int_{k'-2\epsilon}^{k} dH \approx \frac{(2\epsilon-(k'-k))^2}{8\epsilon^2}$. If a mass $\frac{(2\epsilon-(k'-k))^2}{8\epsilon^2}$ of speculators is always selling, the maximum mass of net buy orders is (since every sell order cancels one buy order)

$$1 - \frac{(2\epsilon - (k' - k))^2}{4\epsilon^2} = \frac{4\epsilon(k' - k) - (k' - k)^2}{4\epsilon^2}.$$

Since the signal is $\omega_i = k'$, the minimum number of buy orders to make $x_N + x_S$ positive and therefore make buying profitable is $-k' - \epsilon$ (remember that $k' < 0$). Buying can therefore only make sense after receiving $\omega_i = k'$ if the maximum number of buy orders is larger than the minimum number of buy orders required to make buying profitable, i.e.

$$\begin{aligned}
-k' - \epsilon &< \frac{4\epsilon(k' - k) - (k' - k)^2}{4\epsilon^2} \\
(k' - k)^2 - 4\epsilon(k' - k) - 4\epsilon^2(k' + \epsilon - k + k) &< 0 \\
(k' - k)^2 - (4\epsilon + 4\epsilon^2)(k' - k) - 4\epsilon^2(k + \epsilon) &< 0
\end{aligned}$$

The last inequality is not true for k' very large or very small, so that it must be true between the two solutions for the corresponding equality (since these solutions exist) and we get that buying can only be profitable for k' if

$$2\epsilon\left(1 + \epsilon + \sqrt{(1+\epsilon)^2 + k + \epsilon}\right) > k' - k > 2\epsilon\left(1 + \epsilon - \sqrt{(1+\epsilon)^2 + k + \epsilon}\right).$$

But that is incompatible with $k' \in B(k, \epsilon)$. □

Using this lemma, we can show that there is no $k < -\epsilon$ such that there exists a strategy s with $s(k) > 0$ and $s \in \rho^m(\Sigma)$ for all $m \in \mathbb{N}$. Assume to the contrary that there exists a non-empty set of such ks and let $\hat{k}$ be the infimum of this set. Then $\forall k < \hat{k} : s(k) = 0$, and we can pick a $\bar{k}$ that is arbitrarily close to but below $\hat{k}$. Since $\hat{k} < -\epsilon$, there exists a $\bar{k}$ such that $\hat{k} \in B(\bar{k}, \epsilon)$. Then it follows from the lemma that there cannot exist a strategy that has positive probability of buying anywhere in an open ball around $\bar{k}$.

We can show analogously that there is no $k > \epsilon$ such that there is a strategy s with $s(k) < 1$ and and $s \in \rho^m(\Sigma)$ for all $m \in \mathbb{N}$. Hence the only rationalizable strategies are those that sell with probability 1 for all $\omega_i < -\epsilon$ and buy for all $\omega_i > \epsilon$. □

Proof of proposition 3. First, I show existence if there is common knowledge of x_N. Assume that the market follows its equilibrium strategy, so that $p_2 = p_0$. It is then profitable for any speculator to buy at $p_1 < p_0$ or sell at $p_1 > p_0$. Speculators therefore trade against the noise traders until either all of them have posted an order or $x = 0$ and $p_1 = p_0$. No speculator has any incentive to deviate: Those who post orders either make a positive profit (if $|x_N| > 1$) or no profit (otherwise), and those who do not post orders (since other informed speculators have already driven the price back to p_0) would lose money by trading (since they would move the price above p_0 if they bought or below p_0 if they sold).

Now assume that the speculators follow this strategy. Then p_1 contains absolutely no information about v, since the speculators only either do nothing or counteract the noise traders (whose actions are independent of v), and none of their behavior is conditional on v. The market can therefore only follow its prior and set $p_2 = p_0$.

In the game without common knowledge, consider the following argument: If each speculator can only observe his signal ω_i, it is still optimal to buy if $\omega_i \leq -1$, because in expectation $p_1 < p_0$ regardless of the behavior of other speculators. Now consider the case where $\omega_i \in (-1, 0)$. If all other speculators buy with probability $-\omega_j$ upon observing $\omega_j \in (-1, 0)$, they will on average buy with probability $-x_N$, which means that p_1 will be 0 in expected terms. Every single speculator is then indifferent between buying or selling or doing nothing. Therefore, it is an equilibrium if all speculators buy for $\omega_i \leq -1$, buy with probability $-\omega_i$ for $\omega_i \in (-1, 0)$, sell with probability ω_i if $\omega_i \in (0, 1)$ and always sell if $\omega_i \geq 1$. The non-adjustment equilibrium remains completely undisturbed if x_N is no longer common knowledge.

A brief remark on out-of-equilibrium beliefs: In this equilibrium, total order flow x will be between $1-n$ and $n-1$, meaning that $p_1 \in \left[p_0 + \hat{\lambda}(1-n),\, p_0 + \hat{\lambda}(n-1)\right]$. Out-of-equilibrium beliefs are what the market thinks if p_1 should lie outside that interval. But it is clearly not optimal for the market to assume that prices outside this interval are informative. If it did, and accordingly set some $p_2 > p_0 + \hat{\lambda}(n-1)$ after observing $p_1 > p_0 + \hat{\lambda}(n-1)$, the speculators would have an incentive to try to push p_1 above $p_0 + \hat{\lambda}(n-1)$ regardless of whether $v = v_H$ or $v = v_L$, so that p_1 would not be any more informative than it was before. If, on the other hand, they were to set p_2 with $p_0 < p_2 < p_0 + \hat{\lambda}(n-1)$ after observing $p_1 > p_0 + \hat{\lambda}(n-1)$, the speculators would have no incentive to drive prices out of equilibrium range at all, even if they could submit information in this way. □

B Which Equilibrium is Pareto-Preferred?

Proposition 4. *If f (the density of x_N) is single-peaked, speculators prefer the adjustment to the non-adjustment equilibrium.*

To simplify notation, let $p_2^H(p_1)$ be the expected value of v given p_1 if $p_1 > p_0$, and $p_2^L(p_1)$ the expected value of v given p_1 if $p_1 < p_0$. We make use of the following lemma:

Lemma 2. *If $2 > \frac{p_1-p_0}{\hat{\lambda}}$ it is $\frac{\partial p_2^H(p_1)}{\partial p_1} < 0$ (and hence also $\frac{\partial p_2^L(p_1)}{\partial p_1} > 0$). If $2 \leq \frac{p_1-p_0}{\hat{\lambda}}$, then $p_2^H = p_2^L = p_0$ and consequentially $\frac{\partial p_2^H(p_1)}{\partial p_1} = \frac{\partial p_2^L(p_1)}{\partial p_1} = 0$.*

Proof. It is $p_2^H(p_1) = \pi(p_1)v_H + (1-\pi(p_1))\, v_L$, or

$$p_2^H(p_1) = \frac{1-F\left(\frac{p_1-p_0}{\hat{\lambda}}-1\right)}{2-F\left(k_H\right)-F\left(\frac{p_1-p_0}{\hat{\lambda}}-1\right)}v_H + \frac{1-F\left(k_H\right)}{2-F\left(k_H\right)-F\left(\frac{p_1-p_0}{\hat{\lambda}}-1\right)}v_L.$$

Since $k_H = \max\left\{\left(\frac{p_1-p_0}{\hat{\lambda}}-1\right), 1\right\}$, there are two possible cases:

1. $2 > \frac{p_1-p_0}{\hat{\lambda}}$. Then $k_H = 1$ and

$$p_2^H(p_1) = \frac{1-F\left(\frac{p_1-p_0}{\hat{\lambda}}-1\right)}{2-F\left(1\right)-F\left(\frac{p_1-p_0}{\hat{\lambda}}-1\right)}v_H + \frac{1-F\left(1\right)}{2-F\left(1\right)-F\left(\frac{p_1-p_0}{\hat{\lambda}}-1\right)}v_L.$$

   As $F\left(\frac{p_1-p_0}{\hat{\lambda}}-1\right)$ is monotonously growing in p_1, and since $v_H > v_L$, it is then $\frac{\partial p_2^H(p_1)}{\partial p_1} < 0$.

2. $2 \leq \frac{p_1-p_0}{\hat{\lambda}}$. Then $k_H = \frac{p_1-p_0}{\hat{\lambda}} - 1$ and

$$p_2^H(p_1) = \frac{v_H+v_L}{2} = p_0.$$

□

Proof of Proposition 4. Speculators' expected profit from the efficient equilibrium is the sum of expected profits if $|x_N| < 1$ and $|x_N| \geq 1$. More precisely, it is

$$\Pr\left(|x_N|<1\right)\left(\mathbb{E}\left[p_2^H(p_1)||x_N|<1\right] - p_0 - \frac{\hat{\lambda}}{2}\left(\mathbb{E}\left[x_N||x_N|<1\right]+1\right)\right)$$

$$+\Pr(|x_N|=1)\left(\mathbb{E}\left[p_2^H\left(2\lambda\right)\right] - p_0 - \frac{\hat{\lambda}}{2}\left(2\right)\right) \tag{4}$$

$$+\Pr\left(|x_N|>1\right)\left(\mathbb{E}\left[p_2^H(p_1)|x_N>1\right]-p_0-\frac{\hat{\lambda}}{2}\left(\mathbb{E}\left[x_N|x_N>1\right]+1\right)\right)$$

(Note that, because of symmetry, we can restrict ourselves to the expected prices if $p_1 > p_0$.) All three summands are clearly positive, as we can see from lemma 2 and the proof of proposition 1.

In the inefficient equilibrium, expected profit for any speculator is positive only if $|x_N| > 1$, so that overall expected profit from the inefficient equilibrium is

$$\Pr\left(|x_N|>1\right)\frac{\hat{\lambda}}{2}\left(\mathbb{E}\left[x_N|x_N>1\right]-1\right).$$

If the expression "(Expected profit from efficient equilibrium)−(Expected profit from inefficient equilibrium)" is positive, speculators prefer the efficient equilibrium. We can write this expression as the sum of some positive terms and the term

$$\Pr\left(|x_N|>1\right)\left(\mathbb{E}\left[p_2^H(p_1)|x_N>1\right]-p_0-\frac{\hat{\lambda}}{2}\left(\mathbb{E}\left[x_N|x_N>1\right]+1\right)-\frac{\hat{\lambda}}{2}\left(\mathbb{E}\left[x_N|x_N>1\right]-1\right)\right). \tag{5}$$

From the proof of proposition 1 we know that $\mathbb{E}\left[p_2^H(\lambda(n+1))\right]-p_0-\frac{\hat{\lambda}}{2}(n+1)>0$. From lemma 2, it follows that then also $\mathbb{E}\left[p_2^H(\lambda(x_N+1))|x_N>1\right]>p_0+\frac{\hat{\lambda}}{2}(n+1)$. That means that if

$$\frac{\hat{\lambda}}{2}(n+1)-\frac{\hat{\lambda}}{2}\left(\mathbb{E}\left[x_N|x_N>1\right]+1\right)-\frac{\hat{\lambda}}{2}\left(\mathbb{E}\left[x_N|x_N>1\right]-1\right) \tag{6}$$

is positive, then expression (5) is also positive. (6) simplifies to $n+1-2\mathbb{E}\left[x_N|x_N>1\right]$, which is positive if $\frac{n+1}{2}>\mathbb{E}\left[x_N|x_N>1\right]$. If $f(x)$ is falling in $|x|$, that is the case. □

It should be noted that this is a sufficient, but not a necessary condition: The difference between expected payoffs from the efficient and the inefficient equilibrium can well be positive even if $\frac{n+1}{2}<\mathbb{E}\left[x_N|x_N>1\right]$. But it can be shown that the efficient equilibrium is not always preferred: If f is not falling in its argument, it is possible that speculators actually prefer the inefficient equilibrium. Intuitively, that is the case if f has a lot of mass towards n and $-n$, so that large bubbles (which are profitable for rational speculators in the inefficient equilibrium) become very likely. In the efficient equilibrium, the market adjusts $\pi(p_1)$ accordingly, and speculators' expected profit margins in the efficient equilibrium (which is now not very efficient) become very low. In the inefficient equilibrium, on the other hand, speculators could now make large expected gains, since their profit is higher the further noise traders drive p_1 away from p_0.

Corollary. *There are distributions of x_N so that the efficient equilibrium exists, but speculators* *ex ante prefer the inefficient equilibrium.*

Proof. Consider the the case where $\Pr(x_N = 1) = \Pr(x_N = -1) = \varepsilon$ and $\Pr(x_N = n) = \Pr(x_N = -n) = \frac{1}{2} - \varepsilon$. Then the expected payoff in the inefficient equilibrium is $(1-2\varepsilon)\frac{\hat{\lambda}}{2}(n-1)$, while the expected payoff from the efficient equilibrium is

$$\mathbb{E}\left[\varepsilon p_2^H(p_0 + \lambda(1+x_N))\Big| x_N = 1\right] + \mathbb{E}\left[\varepsilon p_2^H(p_0 + \lambda(1+x_N))\Big| x_N = -1\right]$$

$$+\mathbb{E}\left[(1-2\varepsilon) p_2^H(p_0 + \lambda(1+x_N))\Big| x_N = n\right] - \frac{\hat{\lambda}}{2} - (1-2\varepsilon)\frac{\hat{\lambda}}{2}n - p_0.$$

Let $D_i = \mathbb{E}\left[p_2^H(p_0 + \lambda(1+x_N))\Big| x_N = i\right] - p_0$. Then the difference between profits from the efficient and inefficient equilibrium is

$$\varepsilon D_1 + \varepsilon D_{-1} + (1-2\varepsilon)D_n - \hat{\lambda}\left(\varepsilon + (1-2\varepsilon)n\right). \tag{7}$$

If we take the maximal $\hat{\lambda}$ such that the efficient equilibrium still exists,[13] we have $\hat{\lambda} = 2\frac{D_n}{1+n}$, and (7) becomes

$$\varepsilon D_1 + \varepsilon D_{-1} + \frac{(1-4\varepsilon) - (1-2\varepsilon)n}{1+n} D_n.$$

For this always to be positive, it would have to be

$$\frac{D_1 + D_{-1}}{2} / D_n > \frac{4\varepsilon - 2\varepsilon n - 1 + n}{2\varepsilon(1+n)}.$$

Intuitively, this means that as ε gets arbitrarily small, the prices that result in period 1 from $x_N = 1$ and $x_N = -1$ would have to become infinitely more informative than the prices that result from $x_N = n$ and $x_N = -n$. But a price p_1 that results from $x_N = n$ lies within the price range $\left[p_0 + \hat{\lambda}(-2),\, p_0 + 2\hat{\lambda}\right]$ with constant probability $\frac{2}{1+n}$ because of the price formation process through noisy λ. Therefore, the prices resulting from $x_N = 1$ and $x_N = -1$ can never be infinitely more informative than the prices resulting from $x_N = n$. Therefore, there exists a distribution for x_N so that for a large enough $\hat{\lambda}$ speculators prefer the inefficient to the efficient equilibrium. □

These conditions on the shape of f might seem rather abstract, but they have an intuitive interpretation. f is falling in distance from 0 if the correlation between noise traders' decisions is sufficiently small (they might make their decisions independently, or their actions might even be negatively correlated). In these cases, speculators will always prefer the efficient equilibrium. But high correlation between the decisions of the noise traders means nothing

[13] For very small $\hat{\lambda}$, speculators always prefer the efficient equilibrium.

else than strong herding. If noise traders are sufficiently prone to strong herding, all rational market participants weakly prefer an equilibrium in which no information is transmitted to a partially revealing equilibrium.

C A Discrete Model where Speculators have Market Power

The model can also be written with a finite number of speculators and noise traders, such that signle speculators actually have market power and can influence the price. While this makes some of the expressions less tractable and slightly changes the proofs, the main theorems remain intact and the two equilibria still exist. Assume in the following that there is a finite number n of noise traders and m of speculators.

Proposition 5. *(Efficient equilibrium) It is an equilibrium if every speculator follows the strategy "If $x_N \leq -1$, sell and if $x_N \geq 1$ buy. If $|x_N| < 1$, buy if $v = v_H$ and sell if $v = v_L$." and the market sets*

$$p_2 = p_2^H(p_1) = \pi(p_1)v_H + (1 - \pi(p_1))\, v_L \quad \text{if} \quad p_1 > p_0 \tag{8}$$

$$p_2 = p_2^L(p_1) = (1 - \pi(p_1))\, v_H + \pi(p_1)v_L \quad \text{if} \quad p_1 < p_0 \tag{9}$$

$$p_2 = p_0 \quad \text{if} \quad p_1 = p_0, \tag{10}$$

where

$$\pi(p_1) = \begin{cases} \frac{1-F\left(\left\lfloor \frac{p_1-p_0}{\hat{\lambda}} - m \right\rfloor\right)}{2-F(k_H)-F\left(\left\lfloor \frac{p_1-p_0}{\hat{\lambda}} - m \right\rfloor\right)} & \text{if } p_1 > p_0 \\ \\ \frac{1-F\left(\left\lceil \frac{p_1-p_0}{\hat{\lambda}} + m \right\rceil\right)}{2-F(k_L)-F\left(\left\lceil \frac{p_1-p_0}{\hat{\lambda}} + m \right\rceil\right)} & \text{if } p_1 < p_0 \end{cases}$$

where $\pi(p_1)$ is the market's belief that $v = v_H$ if $p_1 > p_0$ or that $v = v_L$ if $p_1 < p_0$, respectively, with $k_H = \max\left\{\left\lfloor \frac{p_1-p_0}{\hat{\lambda}} - m \right\rfloor, m-1\right\}$ and $k_L = \min\left\{\left\lceil \frac{p_1-p_0}{\hat{\lambda}} + m \right\rceil, -m+1\right\}$, if and only if

$$\hat{\lambda} \leq \mathbb{E}\left[\frac{F(k_H) - F\left(\left\lfloor \frac{p_1-p_0}{\hat{\lambda}} - m \right\rfloor\right)}{2 - F\left(k_H\right) - F\left(\left\lfloor \frac{p_1-p_0}{\hat{\lambda}} - m \right\rfloor\right)} \middle| x_N = n \right] \frac{v_H - v_L}{m+n}. \tag{11}$$

Proof (similar to the continuous case). **Part 1: The market has no incentive to deviate (and $\pi(p_1)$ is the correct belief).**

Assume that the speculators follow their equilibrium strategies and consider the case where $p_1 > p_0$. The market can then, from observing p_1, draw conclusions about v. Let $\pi(p_1)$ be the conditional probability that $v = v_H$ after observing a certain p_1, $\Pr(v_H|p_1)$.

It is

$$\pi(p_1) = \Pr(v_H|p_1) = \frac{\Pr(p_1 \cap v_H)}{\Pr(p_1)} = \frac{\Pr(p_1 \cap v_H \cap |x_N| < m) + \Pr\left(p_1 \cap v_H \cap x_N \geq m\right)}{\Pr\left(p_1 \cap |x_N| < m\right) + \Pr\left(p_1 \cap x_N \geq m\right)}$$

since $\Pr(p_1 \cap x_N \leq -m) = 0$.

If g is the probability density function of λ, we can express this as

$$\pi(p_1) = \frac{\frac{1}{2} \sum\limits_{y=-m+1}^{m-1} f(y) g\left(\frac{p_1 - p_0}{y+m}\right) + \frac{1}{2} \sum\limits_{y=m}^{n} f(y) g\left(\frac{p_1 - p_0}{y+m}\right)}{\frac{1}{2} \sum\limits_{y=-m+1}^{m-1} f(y) g\left(\frac{p_1 - p_0}{y+m}\right) + \sum\limits_{y=m}^{n} f(y) g\left(\frac{p_1 - p_0}{y+m}\right)}.$$

The product in all the sums, $f(y)g\left(\frac{p_1-p_0}{y+m}\right)$, gives the probability that $x_N = y$ and $\lambda = \frac{p_1-p_0}{y+m}$, in which case the parameters would lead to the given p_1 if speculators always bought in period 1. The first sum in the numerator is hence the overall probability that p_1 would be observed as a result of some $x_N \in [-m+1, m-1]$ if speculators always bought the asset. Since, if $x_N \in [-m+1, m-1]$, speculators buy the asset only if $v = v_H$, this probability has to be multiplied by $\frac{1}{2}$ to give the probability $\Pr(p_1 \cap v_H \cap |x_N| < m)$. The second sum in the numerator gives the probability that p_1 would be observed as the result of some $x_N \geq m$. Since $v = v_H$ in only half of these cases, we again need to multiply with $\frac{1}{2}$ (albeit for different reasons) to get the unconditional probability that p_1 would happen as the result of some $x_N > m$ *and* that also $v = v_H$. In the numerator, therefore, we have the overall probability that a given p_1 is observed and is informative.

In the denominator, we then have the overall probability that a given p_1 is observed. This is given by the expression from the numerator, only that now *all* cases in which $x_N > m$ are considered (since they all lead to $p_1 > p_0$), whereas only half of them are informative. The fraction therefore gives the ratio between the number of cases in which p_1 is observed and it is $v = v_H$ and the overall number of cases in which p_1 is observed. This is the conditional probability $\Pr(v_H|p_1)$.

We can simplify the expression: Since λ is uniformly distributed on the interval $\left(0, \hat{\lambda}\right)$, $g\left(\frac{p_1-p_0}{y+m}\right) = \frac{1}{\hat{\lambda}}$ if $0 < \frac{p_1-p_0}{y+m} < \hat{\lambda}$ and 0 otherwise. For any $p_1 > 0$, it is $0 < \frac{p_1-p_0}{y+m}$, but $g\left(\frac{p_1-p_0}{y+m}\right)$

is nonzero only for $y > \frac{p_1 - p_0}{\hat{\lambda}} - m$. We can write

$$\begin{aligned}\pi(p_1) &= \frac{\sum_{y=\left\lceil \frac{p_1-p_0}{\hat{\lambda}} - m \right\rceil}^{k_H} f(y) + \sum_{y=k_H+1}^{n} f(y)}{\sum_{y=\left\lceil \frac{p_1-p_0}{\hat{\lambda}} - m \right\rceil}^{k_H} f(y) + 2\sum_{y=k_H+1}^{n} f(y)} \\ &= \frac{1 - F\left(\left\lfloor \frac{p_1-p_0}{\hat{\lambda}} - m \right\rfloor\right)}{2 - F\left(\left\lfloor \frac{p_1-p_0}{\hat{\lambda}} - m \right\rfloor\right) - F\left(k_H\right)}\end{aligned}$$

where $k_H = \max\left\{\left\lfloor \frac{p_1-p_0}{\hat{\lambda}} - m \right\rfloor, m-1\right\}$. Therefore, given the speculators' strategies, $\pi(p_1)$ gives the correct beliefs in equilibrium.

$\Pr(v_L|p_1)$ is the complementary probability $1 - \pi(p_1)$, so that the expected value of v given p_1 is $E[v|p_1] = \pi(p_1)v_H + (1 - \pi(p_1))\, v_L$. A similar argument applies to the case where $p_1 < p_0$. If $p_1 = p_0$, the price contains no information and p_2 should be set equal to the prior.

p_1 is between $p_0 - \hat{\lambda}(m+n)$ and $p_0 + \hat{\lambda}(m+n)$. For $x_N \in \{-n, n\}$, all possible p_1 occur with positive probability, so that in equilibrium (where $x_N \in [-n, n]$) all possible p_1 occur with positive probability and there can be no out-of-equilibrium beliefs.

Part 2: Speculators make a positive profit in equilibrium.

Now assume that the market follows its equilibrium strategy. Consider the case where $p_1 > p_0$, meaning that either $|x_N| < m$ and $v = v_H$ or simply $x_N \geq m$. If they follow their equilibrium strategies, the speculators' buy orders will drive the price to $p_0 + \lambda(m + x_N) > p_0$, and in period 2 all speculators will be able to sell their holdings at $p_2^H = \pi(p_1)v_H + (1 - \pi(p_1))\, v_L$. Their profit is then $p_2^H - p_1$, or $\pi(p_1)v_H + (1 - \pi(p_1))\, v_L - p_0 - \lambda(m + x_N)$, which can also be written as

$$\begin{aligned}&\left[\frac{1 - F\left(\left\lfloor \frac{p_1-p_0}{\hat{\lambda}} - m \right\rfloor\right)}{2 - F\left(k_H\right) - F\left(\left\lfloor \frac{p_1-p_0}{\hat{\lambda}} - m \right\rfloor\right)} - \frac{1}{2}\right] v_H + \left[\frac{1 - F\left(k_H\right)}{2 - F\left(k_H\right) - F\left(\left\lfloor \frac{p_1-p_0}{\hat{\lambda}} - m \right\rfloor\right)} - \frac{1}{2}\right] v_L - \lambda(m + x_N) \\ &= \frac{F(k_H) - F\left(\left\lfloor \frac{p_1-p_0}{\hat{\lambda}} - m \right\rfloor\right)}{4 - 2F\left(k_H\right) - 2F\left(\left\lfloor \frac{p_1-p_0}{\hat{\lambda}} - m \right\rfloor\right)}(v_H - v_L) - \lambda(m + x_N)\end{aligned} \tag{12}$$

x_N is known to the speculators. Then we can write expression 12 in expected terms (given x_N):

$$\mathbb{E}\left[\left.\frac{F(k_H) - F\left(\left\lfloor \frac{p_1-p_0}{\hat{\lambda}} - m \right\rfloor\right)}{4 - 2F\left(k_H\right) - 2F\left(\left\lfloor \frac{p_1-p_0}{\hat{\lambda}} - m \right\rfloor\right)}\right| x_N\right](v_H - v_L) - \frac{\hat{\lambda}}{2}(m + x_N).$$

Since p_1 is monotonically increasing in x_N, and therefore $F\left(\left\lfloor \frac{p_1-p_0}{\hat{\lambda}} - m \right\rfloor\right)$ and $F(k_H - 1)$ are weakly increasing in x_N, the whole expression becomes minimal for $x_N = n$, where it is

$$\mathbb{E}\left[\left.\frac{F(k_H) - F\left(\left\lfloor\frac{p_1-p_0}{\hat{\lambda}} - m\right\rfloor\right)}{4 - 2F\left(k_H\right) - 2F\left(\left\lfloor\frac{p_1-p_0}{\hat{\lambda}} - m\right\rfloor\right)}\right| x_N = n\right](v_H - v_L) - \frac{\hat{\lambda}}{2}(m+n).$$

If this is positive, then speculators will make an expected profit by following their equilibrium strategies for all x_N (the case where x_N is negative is analogous and leads to the same result). We can reformulate the condition as

$$\hat{\lambda} \leq \mathbb{E}\left[\left.\frac{F(k_H) - F\left(\left\lfloor\frac{p_1-p_0}{\hat{\lambda}} - m\right\rfloor\right)}{2 - F\left(k_H\right) - F\left(\left\lfloor\frac{p_1-p_0}{\hat{\lambda}} - m\right\rfloor\right)}\right| x_N = n\right]\frac{v_H - v_L}{m+n}$$

which is simply the spread between high and low value, adjusted for the number of market participants and some adjustment factor that depends on the precise shape of f.

Part 3: No single speculator has an incentive to deviate from his equilibrium strategy.

As speculators always make a profit in equilibrium, it would not be profitable for any speculator to deviate by not trading at all. But what if a speculator decided to sell if his equilibrium action would be to buy? We have to distinguish three cases (note that "buy" would never be an equilibrium action if $x_N \leq -m$):

1. $x_N = -(m-1)$. In this case it is $x = 1$ in equilibrium, and if a single speculator decided to sell instead of buying, x would be -1. Since $p_2(p_1)$ is point-symmetric around (p_0, p_0) (i.e. $p_2(p_1) - p_0 = p_0 - p_2(p_0 - (p_1 - p_0))$ because of the symmetry assumption on f), the speculator who sold would gain just as much in expectation as he would have by buying. Since he is thus indifferent, there is no incentive to deviate from equilibrium strategies.

2. $x_N = -(m-2)$. Then $x = 2$ in equilibrium, but if a single speculator sold instead of buying, the resulting net order flow would be 0, so that $p_1 = p_0$. Then it would also be $p_2 = p_0$, so that the speculator would make no gain at all by selling, whereas he could have made a positive profit by buying.

3. $x_N > -(m-2)$. Then $x > 2$ in equilibrium, and a single speculator can only lower x to some slightly lower, but still positive number. Then $p_2 = p_2^H(p_1) > p_1$, so that the speculator would actually make a loss by selling in period 1.

We can therefore conclude that no speculator has an incentive to deviate from his equilibrium strategy if $p_1 > p_0$. A similar argument applies where $p_1 < p_0$ (i.e. if speculators bought instead of selling). □

References

Abreu, D. and M. Brunnermeier (2003). Bubbles and crashes. *Econometrica 71*(1), 173–204.

Allen, F., S. Morris, and H. Shin (2006). Beauty contests and iterated expectations in asset markets. *Review of Financial Studies 19 (3)*, 719–752.

Aumann, R. (1976). Agreeing to disagree. *Annals of Statistics 4*(6), 1236–1239.

Benhabib, J. and P. Wang (2015). Private information and sunspots in sequential asset markets. *Journal of Economic Theory 158*, 558–584.

Brunnermeier, M. and S. Nagel (2004). Hedge funds and the technology bubble. *Journal of Finance 59 (5)*, 2013 – 2040.

Carlsson, H. and E. van Damme (1993). Global games and equilibrium selection. *Econometrica 61*(5), 989–1018.

Cespa, G. and X. Vives (2015). The beauty contest and short-term trading. *Journal of Finance 70(5)*, 2099–2153.

De Long, J., A. Shleifer, L. Summers, and R. Waldmann (1990). Noise trader risk in financial markets. *Journal of Political Economy 98*(4), 703–738.

Dow, J. and G. Gorton (1994). Arbitrage chains. *Journal of Finance 49*(3), 819–849.

Farrell, J. and M. Rabin (1996). Cheap talk. *Journal of Economic Perspectives 10*(3), 103–118.

Frankel, D., S. Morris, and A. Pauzner (2003). Equilibrium selection in global games with strategic complementarities. *Journal of Economic Theory 108*, 1–44.

Froot, K. and E. Dabora (1999). How are stock prices affected by the location of trade? *Journal of Financial Economics 53*(2), 189–216.

Galbraith, J. (1954). *The Great Crash 1929.* Mariner Books (reprinted 1997).

Glosten, L. and P. Milgrom (1985). Bid, ask and transaction prices in a specialist market with heterogeneously informed traders. *Journal of Financial Economics 14*(1), 71–100.

Hayek, F. (1945). The use of knowledge in society. *American Economic Review 35(4)*, 519–530.

Janis, I. (1972). *Victims of Groupthink: a Psychological Study of Foreign-Policy Decisions and Fiascoes.* Houghton Mifflin.

Keynes, J. (1936). *The General Theory of Employment, Interest, and Money.* Harcourt (reprinted 1965).

Kuran, T. (1997). *Private truths, public lies: The social consequences of preference falsification.* Harvard University Press.

Kyle, A. (1985). Continuous auctions and insider trading. *Econometrica 53*(6), 1315–1336.

Leland, H. (1992). Insider trading: Should it be prohibited? *Journal of Political Economy 100(4)*, 859–887.

Lowenstein, R. (2000). *When genius failed: The rise and fall of Long-Term Capital Management.* Random House Trade Paperbacks.

Morris, S. and H. Shin (1998). Unique equilibrium in a model of self-fulfilling currency attacks. *American Economic Review 88(3)*, 587–597.

Morris, S. and H. Shin (2002). Social value of public information. *American Economic Review 92*(5), 1521–1534.

Rubinstein, A. (1989). The electronic mail game: Strategic behavior under "almost common knowledge". *American Economic Review 79(3)*, 385–391.

Shiller, R. (2000). *Irrational Exuberance.* Princeton University Press.

Weinstein, J. and M. Yildiz (2007). A structure theorem for rationalizability with application to robust predictions of refinements. *Econometrica 75 (2)*, 365–400.