

# THE MADNESS OF CROWDS AND THE LIKELIHOOD OF BUBBLES\*

Alex Chinco<sup>†</sup>

November 13, 2018

## Abstract

Market participants are constantly swimming in a sea of psychological biases and trading constraints. And yet, in spite of all these biases and constraints, large pricing errors such as speculative bubbles are rare. Why is this? How often should we expect that some psychological bias will cause some trading constraint to bind? This paper proposes a model to answer this question.

In the model, the number of speculators excited about an asset varies over time due to social interactions. As long as the asset's past returns remain below a critical threshold,  $r < r_*$ , these social interactions disperse any crowd of speculators that happens to get excited. But, as soon as the asset's past returns rise above this critical threshold,  $r > r_*$ , the exact same social interactions suddenly make the speculator population boom. This population explosion amplifies the effects of speculator biases, causing arbitrageur constraints to bind and a speculative bubble to form.

The model predicts that speculative bubbles will be more common in assets where the strength of speculator interactions is more sensitive to short-run fluctuations in past returns. What's more, it's possible to estimate this key sensitivity parameter using data collected during normal times—i.e., when there's no speculative bubble currently taking place. I verify this prediction empirically: a 1%pt increase in the sensitivity of speculator interactions to short-run fluctuations in past returns is associated with a 3.79%pt increase in the likelihood of a speculative bubble following a price run-up.

JEL CLASSIFICATION: G02, G11, G12

KEYWORDS: Speculative Bubbles, Social Interactions, Displacement Events

---

\*I would like to thank Pedro Bordalo, Zhi Da, Xavier Gabaix, Sam Hartzmark, Lawrence Jin, Neil Pearson, Joel Peress, and Josh Pollet for extremely helpful comments and suggestions. The paper has also benefited greatly from presentations at INSEAD, the Wabash River Conference, UIUC, and the NBER Behavioral Finance meetings.

Current Version: <http://www.alexchinco.com/madness-of-crowds.pdf>

<sup>†</sup>University of Illinois at Urbana-Champaign, Gies College of Business. [alexchinco@gmail.com](mailto:alexchinco@gmail.com).

# 1 Introduction

Constrained arbitrageurs might not be able to correct the pricing errors caused by biased speculators. This insight is known as the “limits of arbitrage” (Shleifer and Vishny, 1997). And, over the last 30 years, researchers have tabulated long lists of speculator biases—e.g., overconfidence (Daniel et al., 1998; Scheinkman and Xiong, 2003), heuristic updating (Hong and Stein, 1999; Barberis et al., 2015; Bordalo et al., 2018), sentiment swings (Baker and Wurgler, 2006)—and arbitrageur constraints—e.g., short-sale bans (Miller, 1977; Xiong and Yu, 2011), margin requirements (Gromb and Vayanos, 2002; Garleanu and Pedersen, 2011), bounded rationality (Gabaix, 2014), equity limits (Shleifer and Vishny, 1997; Vayanos, 2004), coordination frictions (Abreu and Brunnermeier, 2003).

According to this limits-to-arbitrage logic, any bias×constraint pair drawn from these two lists has the potential to combine at any moment to produce a pricing error. There are myriad possibilities. But, here’s the thing: large pricing errors such as speculative bubbles are relatively rare. For example, Kindleberger (1978) examined national price indexes since the early 1600s and found only 34 bubble episodes! Similarly, Greenwood et al. (2018) found only 21 episodes since 1928 where the prices of stocks in a particular industry doubled in under two years and then immediately crashed. By any reasonable metric, there are now more academic papers about the limits of arbitrage than speculative bubbles created by them.

Why is this? What pins down the likelihood of a speculative bubble? How often should we expect to find some psychological bias causing some arbitrageur constraint to bind?

Notice that these aren’t questions you can answer within the existing limits-to-arbitrage framework. This framework explains how a pricing error, such as a speculative bubble, can be sustained in equilibrium. But, we’re not asking questions about *how* speculative bubbles can be sustained; we’re asking questions about *how often* we should expect them to occur. And, these are two entirely different things. Discovering *how* Dr. Jekyll turns himself into Mr. Hyde tells you nothing about *how often* you should expect to find a crazed monster terrorizing the streets of Victorian London (Stevenson, 1886).

To answer the above questions, you need to introduce another ingredient to the mix—an on/off switch—something that sporadically amplifies the effects of speculator biases, causing arbitrageur constraints to bind and a speculative bubble to form. This special something is typically called a “displacement event” (Minsky, 1992) in popular accounts of bubble formation. And, in this paper, I propose a theory of displacement events based on a novel recombination of two common elements found in many of these well-known narratives: #1) while speculators get overly excited following good news about fundamentals due to the “madness of crowds”, #2) they only recover their senses “slowly and one by one” (Mackay, 1841).

*Economic Model.* How can a speculative bubble be sustained in equilibrium? How often should we expect a speculative bubble to occur? The limits of arbitrage provide an answer to the first question. However, to answer the second question, we need to add something new to the existing limits-to-arbitrage framework. And, I begin my analysis by proposing a model to do just that. I build on a completely standard limits-to-arbitrage benchmark model. There’s a single risky asset and two kinds of traders: newswatchers and speculators. Newswatchers learn about the asset’s fundamental value and incorporate this information into its price. Speculators only invest in the risky asset when excited, in which case they overvalue the asset and have excess demand. Because the newswatchers are constrained, the excess demand coming from excited speculators can push the risky asset’s price above its fundamental value and create a speculative bubble.

What’s new about the model is that the number of speculators excited about the risky asset varies over time due to social interactions between speculators. There are two key forces at work. The first is the “madness of crowds” (Mackay, 1841). Excited speculators are always trying to get the remaining apathetic speculators excited about the risky asset, too. And, they are always trying to do this using the same kinds of arguments. It’s just that these arguments are more persuasive following good past performance—i.e., when the risky asset’s past returns  $r \in (0, \infty)$  are larger. The interaction between an asset’s past returns and the size of its excited-speculator population is often called “feedback trading” (Cutler et al., 1990).

However, the madness of crowds only provide an *on switch*. And, we need an *on/off switch*. This is an important distinction because speculative bubbles are rare. “Only a relatively small proportion of large shocks leads to a speculative mania.” (Kindleberger, 1978) So, to explain why social interactions only sporadically amplify speculators’ omnipresent biases, I introduce another common element found in popular accounts of bubble formation—namely, that excited speculators “recover their senses slowly and one by one” (Mackay, 1841). This second force is typically used as an excuse for why the madness of crowds can persist for a long time. But, in this paper, it plays an entirely different role, explaining why the madness of crowds only occasionally takes hold in the first place.

Combining these two key forces allows small continuous changes in an asset’s past returns to produce large discontinuous jumps in the number of excited speculators. There exists a critical performance level,  $r_*$ , such that social interactions make the speculator population vanish whenever  $r < r_*$ . This is because, at sufficiently low values of  $r$ , the first speculator excited about an asset will tend to regain his senses before exciting any of his friends, which results in a market where returns are governed by changes in an asset’s fundamental value. However, as soon as  $r > r_*$ , the exact opposite intuition suddenly takes hold. These very same social interactions will now cause the speculator population explode, inducing arbitrageur

constraints to bind and a speculative bubble to form. A positive shock to fundamentals that pushes the asset’s past returns above  $r_*$  represents a “displacement event” (Minsky, 1992) in the model. And, following a displacement event, we say that the market is suddenly governed by “speculative euphoria” (Minsky, 1970), “mob psychology” (Kindleberger, 1978), “irrational exuberance” (Shiller, 2000), etc. . .

For an everyday analogy, think about what happens when you place a glass of water in the freezer, causing it to slowly cool from room temperature to well below freezing. Water molecules in the glass are always trying to attract one another and form ice crystals, even at room temperature. It’s just that the random jiggling of heat energy tends to break up any embryonic two-molecule ice crystals faster than these couplets can attract more of their neighbors at room temperature. So, larger ice crystals can’t form, and the glass of water remains liquid. But, as the temperature steadily drops, this random jiggling gets less and less frenetic. And, a sudden qualitative change occurs when the temperature in the glasses crosses below the critical threshold of  $0^\circ$  Celsius. Below this threshold, embryonic two-molecule ice crystals now tend to attract neighboring molecules faster than Brownian motion can shake them apart. So, larger crystals can form, and the entire glass freezes solid.

*Empirical Implication.* I next develop the main empirical implication of this model: assets where speculator persuasiveness is more sensitive to changes in past returns should experience speculative bubbles more often. For some assets, good performance makes the stories that excited speculators tell much more persuasive. A 10% increase in house prices will generate lots of word-of-mouth buzz, resulting in lots of new second-home buyers. For others, the exact same price run-up will have little effect. A 10% increase in the price of textile stocks will not attract many new traders to this market. I use the parameter  $\theta \in (0, 1)$  to capture the persuasiveness sensitivity for a particular asset. And, I show the model predicts that speculative bubbles will occur more often in real estate than in textile stocks, for example, because the critical performance threshold,  $r_* = 1/\theta$ , is lower in the real-estate market.

More importantly, because social interactions between speculators are taking place both during and between bubble episodes, it’s possible to estimate  $\theta$  using data collected during normal times. This is new and different. If you want to learn about speculative bubbles in the limits-to-arbitrage framework, then you have to wait for arbitrageur constraints to bind. Otherwise, the relevant economic forces simply aren’t operational. You simply can’t observe the phenomenon during normal times. By contrast, in any theory of displacement events, the exact same economic forces must be at work both during and between bubble episodes because these forces must explain not only why speculative bubbles sometimes occur but also why they typically do not. And, as a result, we can learn about the economic forces as captured by  $\theta$  prior to any speculative bubble taking place.


In an ideal world, we'd be able to simply look up the relevant value of  $\theta$  for each asset in an existing database, such as CRSP or Compustat. Unfortunately, this isn't the case. We will have to develop a proxy  $\theta$  using variables that we can observe in the same way that financial economists often use the S&P 500 return as a proxy for the return on total wealth. And, the economic model even provides guidance on how to do this. It suggests estimating  $\theta$  by calculating how much short-run fluctuations in an asset's past returns affect the rate at which its excited-speculator population decays during normal times. By analogy, it's possible to infer the freezing point of water by measuring the speed at which transient two-molecule ice crystals shake apart at room temperature. Thus, if we can find a measure of how long the transient population of excited speculators remains interested during normal times, then the model tells us how to convert this measure into a proxy for the asset's  $\theta$ .

*Econometric Analysis.* Finally, I verify the model's main empirical implication using data on monthly U.S. industry returns. To proxy for each industry's  $\theta$ , I first estimate how long the transient population of excited speculators remains interested in an industry during normal times by counting the number of Wall Street Journal (WSJ) articles referencing that industry. The idea is that news outlets strategically choose which industries to cover so as to maximize total readership (Mullainathan and Shleifer, 2005). Thus, if an industry starts to get more coverage while  $r < r_*$ , then it's likely that the transient population of excited speculators is remaining in the market longer. Then, following the logic outlined above, I estimate  $\theta$  by correlating the number of WSJ articles that reference an industry with short-run fluctuations in that industry's past returns.

I define a speculative bubble as an episode where stock prices in an industry double in under two years and then immediately fall by at least 50%. This definition captures the main industry-level events that many researchers loosely refer to as "speculative bubbles", such as the rise and fall of technology stocks during the late 1990s and early 2000s. But, I want to emphasize that the empirical results do not hinge on whether any particular market episode was or was not a speculative bubble. We can reason about the likelihood of a particular kind of event taking place even if there is sometimes disagreement about whether the event has taken place after the fact. I realize that this might seem counter-intuitive. You're probably thinking to yourself: *'What hope is there of predicting the likelihood of a speculative bubble if we can't even agree about whether a particular market episode was in fact a bubble?'* But, this is a false choice. We do this sort of thing all the time. For example, we know that suicide is the most common form of gun death in America<sup>1</sup> even if experts sometimes disagree about whether a particular gun death should be classified as a suicide or an accident.<sup>2</sup>

<sup>1</sup>"Gun Deaths In America." <https://fivethirtyeight.com/features/gun-deaths/>

<sup>2</sup>"What Counts As An Accident?" <https://fivethirtyeight.com/features/gun-accidents/>



**Figure 1. Neo-Classical Models.** Consider an asset with  $\$v$  payout tomorrow and three units available for purchase today (white circle). Green: speculator demand,  $\{0, 1, \dots, 9\}$ , at price  $p = \$v$ . Red: arbitrageur response,  $\{0, -1, \dots, -9\}$ . **Rational Demand:** Demand when speculators are fully rational. **Excess Demand:** Demand when speculators are biased and arbitrageurs are unconstrained.


Consistent with the prediction that assets with higher  $\theta$ s will experience speculative bubbles more often, I find that a 1%pt increase in an industry's estimated value of  $\theta$  is associated with a 3.79%pt increase in the likelihood of experiencing a speculative bubble. This result holds after controlling for industry-level valuation ratios and return volatility. It also holds when controlling for the overall level of media coverage dedicated to the industry. In addition, I show that these empirical results are robust to omitting particular bubble episodes from the analysis. I'm labeling a particular set of market episodes as speculative bubbles. If you strongly believe that one of these episodes was not in fact a speculative bubble, then feel free to leave it out. The omission will not qualitatively affect the results. Last but not least, I estimate a Cox proportional hazard-rate model to show that, while variation in  $\theta$  across industries predicts the likelihood of a speculative bubble, it does not predict the exact timing of the crash. The factors affecting *how* a speculative bubble can be sustained in equilibrium are different from those affecting *how often* we should expect to encounter one.

## 1.1 Related Literature

The existing literature on speculative bubbles investigates how these large pricing errors can be sustained in equilibrium. This paper is doing something different. So, before moving on, it's helpful to discuss how the current paper relates to and differs from this existing literature.

*Neo-Classical Models.* There's a good reason why the existing literature focuses on the question of existence: there's no such thing as a speculative bubble in neo-classical models. Proving existence is a crucial first step. There's no sense studying non-existent phenomenon. We can still debate whether suicide is the most likely form of gun death if we disagree about the definition of a suicide but not if one of us thinks gun deaths never occur. To illustrate the





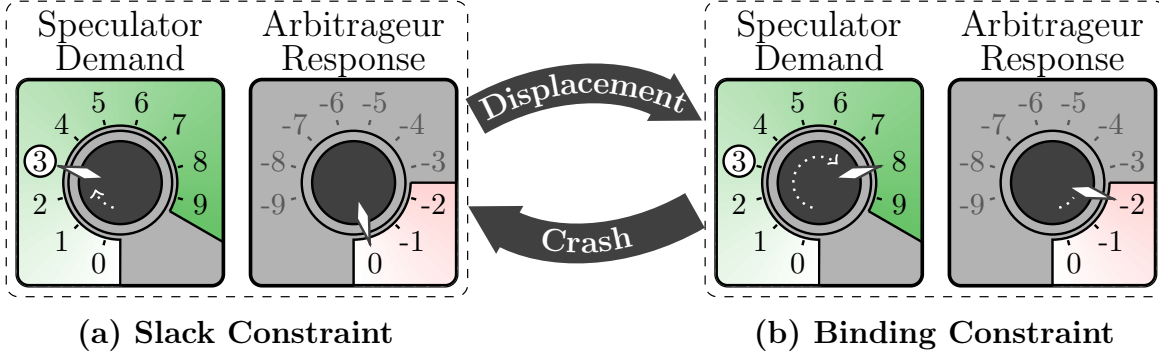
**Figure 2. Limits to Arbitrage.** Consider an asset with a  $\$v$  payout tomorrow and three units available for purchase today (white circle). Green: speculator demand,  $\{0, 1, \dots, 9\}$ , at price  $p = \$v$ . Red: arbitrageur response,  $\{0, -1, -2\}$ , when they can't short more than two units due to some trading constraint.

source of the problem, consider an asset with a  $\$v$  per share payout tomorrow and three units available for purchase today. In a neo-classical model, speculators are assumed to be fully rational. In other words, they demand three units at a price of  $p = \$v$  per share. In Figure 1a, this is shown by turning the green speculator-demand knob, which represents speculator demand when the equilibrium price is  $p = \$v$  per share, to three units.

Neo-classical models also assume that rational arbitrageurs would always be able to undo the effects of any biased speculators in the market (Friedman, 1953; Fama, 1965). Suppose a group of biased speculators wanted to hold an extra five units of the asset as shown by turning the green speculator-demand knob to eight units in Figure 1b,  $3 + 5 = 8$ . This means that rational arbitrageurs could make money by selling five units of the asset at a price above  $p = \$v$  per share to these biased speculators today, as shown by turning the red arbitrageur-demand knob to negative five in Figure 1b. The idea in neo-classical models is that, because the end result of this trade is aggregate demand of three units,  $8 + (-5) = 3$ , this arbitrage activity would effectively correct the pricing error caused by speculators' excess demand. It would result in the same level of aggregate demand as in the original example where speculators were fully rational and demanded only three units themselves.

*Limits to Arbitrage.* The key insight in the limits-to-arbitrage literature (Shleifer and Summers, 1990; Shleifer and Vishny, 1997), however, is that rational arbitrageurs can't always execute this sort of trade in the real world. Instead of shorting the full five units as in Figure 1b, they might only be able to short two units due to various constraints, as shown by limiting the red arbitrageur-demand knob in Figure 2 to values of  $\{0, -1, -2\}$ . Thus, in the presence of both biased speculators and constrained arbitrageurs, the price of an asset can rise above its fundamental value due to the resulting three extra units of aggregate demand,  $(8 - 3) - 2 = 3$ .

The limits-to-arbitrage framework provides necessary conditions for *how* a pricing error, such as a speculative bubble, can be sustained in equilibrium: the green speculator-demand knob needs to be turned too high by some bias; and, some constraint needs to prevent the red arbitrageur-response knob from undoing this error. But, the limits-to-arbitrage framework doesn't tell you *how often* you should expect to observe this configuration. We know that



**Figure 3. Displacement vs. Crash.** Consider an asset with a  $\$v$  payout tomorrow and three units available for purchase today (white circle). Green: speculator demand,  $\{0, 1, \dots, 9\}$ , at price  $p = \$v$ . Red: arbitrageur response,  $\{0, -1, -2\}$ , when they can't short more than two units due to some trading constraint. **Slack Constraint.** Regime where speculator biases don't cause arbitrageur constraints to bind. **Binding Constraint.** Regime where they do.

speculators suffer from a wide range of psychological biases, which make it possible for the green speculator-demand knob to get turned up too high as in Figure 3b. And, we know that arbitrageurs suffer from a wide range of trading constraints, which can prevent the red arbitrageur-response knob from undoing this error as in Figure 3b. So, given all these biases and constraints, why does the world often look like Figure 3a with no pricing error?

*Predicting Crashes.* There's been a lot of recent discussion about the importance of being able to predict the end of a speculative bubble. The crash "has to be a predictable phenomena" (Eugene Fama quoted in Cassidy, 2010). The idea is that any complete theory of bubbles should predict when they pop. And, I totally agree. It would be great if someone were able to model the "Crash" arrow in Figure 3. But, the "Crash" arrow is only half the story. Traders obviously care more about timing the crash. But, this isn't true for policymakers. A policymaker wants to know what makes a speculative bubble form in the first place. Whether the resulting crash gets remembered as "Black Monday" or "Black Tuesday" is utterly irrelevant. Policymakers care about displacements. To illustrate, note that Alan Greenspan expressed his concerns about "froth"<sup>3</sup> in the housing market not about the timing of any eventual crash.

*Epidemic Models.* The idea that social epidemics might distort asset prices is not new. For example, Shiller (1984), Shive (2010), and Andrei and Cujean (2017) all use epidemic models to show how social interactions can generate return predictability. Han et al. (2018) uses social dynamics to explain why active investors persist despite losing money. And, Burnside et al. (2016) employs a similar model to study boom-bust cycles in the housing market. Andrei and Cujean (2017) even points out that there exists a threshold contact rate between

<sup>3</sup>"Greenspan Is Concerned About 'Froth' In Housing." The New York Times. May 21, 2005.



traders below which momentum/reversal dynamics disappear.

The current paper clearly borrows from and builds on this exciting earlier work. But, it also applies this idea in a new way. The key insight in this paper is that the questions of ‘*How can speculative bubbles be sustained in equilibrium?*’ and ‘*How often should we expect to see them?*’ are conceptually distinct. You need a different kind of model to answer each one. So, while the earlier literature studies how social interactions can distort speculator beliefs (question #1; intensive margin), this paper shows how social interactions can generate sudden changes in the number of speculators with distorted beliefs (question #2; extensive margin).

The reason for using an epidemic model is also different. If you want to make predictions about the likelihood of a bubble episode in the future (question #2), then it must be possible to observe the relevant economic forces during normal times. And, the non-linearity embedded in epidemic models provides one way to model this connection. The point is not that social interactions can distort prices or that these distortions might only exist above some threshold. The point is that the combination of these two things make it possible to predict the likelihood of a future bubble episode using parameter estimates made during normal times.

*Extrapolative Beliefs.* Finally, I study a model where excited speculators have extrapolative beliefs (Cutler et al., 1990; De Long et al., 1990; Hong and Stein, 1999; Barberis and Shleifer, 2003). But, I want to emphasize that the paper’s results are not specific to this particular bias. A central goal of the model is to distinguish the reason that excited speculators distort prices (in this case extrapolative beliefs) from the mechanism causing excited speculators to suddenly flood the market (the madness of crowds).

I model speculators who suffer from extrapolative beliefs because the price spiral implied by this psychological bias explains key stylized facts about how speculative bubbles evolve once they get started (Barberis et al., 2015; Glaeser and Nathanson, 2017; Barberis et al., 2018). However, the choice of psychological bias only affects how speculative bubbles evolve once excited speculators enter the market. As emphasized in Corollary 3.1, this choice doesn’t affect the ex ante likelihood that a crowd of excited speculators will enter the market.

Furthermore, because they operate entirely within the standard limits-to-arbitrage framework, the proposed mechanism for how speculative bubbles start is different in the existing models. Extrapolative beliefs are always causing arbitrageur constraints to bind. It’s just that, during normal times when fundamental traders are present, the ensuing speculative bubbles are small. Large bubble episodes occur, not when a new population of excited speculators decides to enter the market, but when the existing population fundamental traders decides to exit. Displacement events in these existing models are the result of *rational egress* rather than *irrational exuberance*. This paper pairs the price spiral implied by extrapolative beliefs with a theory of displacement events that’s more faithful to popular accounts.

## 2 Economic Model

This section develops an economic model where the number of speculators excited about an asset varies over time. Nothing about the model’s setup is going to suggest the existence of any sort of sudden qualitative change. Nevertheless, I show that social interactions between speculators are going to allow small continuous changes in an asset’s past returns to generate large discontinuous jumps in the size of the excited-speculator population. This sudden change in population size amplifies the effects of excited speculators’ pre-existing biases, causing arbitrageur constraints to bind and a speculative bubbles to form.

### 2.1 Displacement Events

I begin by describing how the number of excited speculators evolves over time.

*Speculator Population.* Consider a single risky asset in a market where time is continuous,  $\tau \geq 0$ , and there are  $K \gg 1$  speculators. These agents have no private information about the risky asset’s fundamental value. So, when behaving rationally, they sit on the sideline and do not invest. But, speculators don’t always behave rationally. Sometimes a speculator will get overly excited about the asset and enter the market. And, the number of excited speculators in the market will ebb and flow over time due to social interactions. Let  $N_\tau \geq 0$  denote the number of speculators currently excited about the risky asset. Likewise, let  $n_\tau \stackrel{\text{def}}{=} N_\tau/K$  denote the corresponding excited fraction.

*Getting Excited.* Speculators become overly excited about the risky asset in a process resembling a social epidemic. Each excited speculator is always trying to entice the remaining  $(K - N_\tau) = (1 - n_\tau) \cdot K$  apathetic speculators using the same kinds of arguments. However, I assume that the persuasiveness of these arguments is an increasing function of the asset’s past returns,  $r \in (0, \infty)$ . In other words, excited speculators find it easier to persuade their friends to jump on the bandwagon when an asset’s had strong recent performance. This interaction between past returns and the number of excited speculators is often called “feedback trading” (Cutler et al., 1990).

Feedback trading plays a central role in popular accounts of bubble formation. Shiller (2000) describes how “whenever the market reaches a new high, public speakers, writers, and other prominent people suddenly appear, armed with explanations for the apparent optimism seen in the market. . . The new era thinking they promote is part of the process by which a boom may be sustained and amplified—part of the feedback mechanism that. . . can create speculative bubbles”. Similarly, Kindleberger (1978) argues that “speculation often develops in two stages. In the first, sober, stage households, firms and investors, respond to a shock in a limited and rational way; in the second, the anticipations of capital gains play an increasingly

dominant role in their transactions... [and this] analysis in terms of two stages suggests two groups of speculators, the insiders and the outsiders... The outsider amateurs who buy high and sell low are the victims of euphoria that affects them late in the day. After they lose, they go back to their normal occupations to save for another splurge...”

There’s also ample evidence of feedback trading in the academic literature. For example, [Gong et al. \(2016\)](#) and [Pearson et al. \(2017\)](#) both document how good past performance drew in successive rounds of new uninformed investors during the recent Chinese warrants bubble ([Xiong and Yu, 2011](#)). In addition, [Brunnermeier and Nagel \(2004\)](#), [Greenwood and Nagel \(2009\)](#), and [Griffin et al. \(2011\)](#) all give evidence that sophisticated traders exploited a sudden inflow of inexperienced investors during the .com bubble. Likewise, [Chinco and Mayer \(2015\)](#) show that uninformed out-of-town buyers poured into the cities that realized explosive house-price growth during the U.S. housing bubble. There’s a also healthy literature documenting how peer effects influence market participation ([Shiller and Pound, 1989](#); [Duffo and Saez, 2002](#); [Hong et al., 2004](#); [Brown et al., 2008](#); [Engelberg and Parsons, 2011](#); [Kaustia and Knüpfer, 2012](#); [Bursztyn et al., 2014](#); [Li, 2014](#); [Ahern, 2017](#)). Traders are more likely to enter a market if they already know someone who’s done so. The goal of this section is to provide an economic model that explains why these well-documented peer effects only sporadically devolve into feedback trading.

I use  $\Theta(n, r|\Delta\tau)$  to denote the probability that an additional speculator becomes excited about the risky asset during the short interval of time  $(\tau, \tau + \Delta\tau]$  for some small  $\Delta\tau > 0$ :

$$\Theta(n, r|\Delta\tau) \stackrel{\text{def}}{=} \Pr[n_{\tau+\Delta\tau} - n_{\tau} = +1/K \mid n_{\tau} = n, r] \quad (1)$$

And, I use the following functional form to capture the essence of feedback trading:

$$\lim_{\Delta\tau \searrow 0} \Theta(n, r|\Delta\tau) \stackrel{\text{def}}{=} \Theta(n, \theta, r) = \theta \cdot r \cdot (1 - n) \times n \quad (2)$$

In the equation above,  $\Theta(n, \theta, r)$  represents the continuous limit of the population growth rate for the crowd of excited speculators. I’ve broken this growth rate into two multiplicative factors:  $\theta \cdot r \cdot (1 - n)$  and  $n$ . The trailing  $n$  reflects the fact that the rate at which new speculators get excited is determined by the rate at which each individual excited speculator can excite his currently apathetic friends. The leading expression,  $\theta \cdot r \cdot (1 - n)$ , represents this per capita excitation rate, and each of the terms in this expression have a natural interpretation. The  $(1 - n)$  term represents the apathetic fraction—i.e., the fraction of speculators who aren’t currently excited about the asset. This term is included to capture the idea that it’s harder for each excited speculator to entice new speculators when there are fewer remaining apathetic speculators left for him to interact with.

The  $r \in (0, \infty)$  term represents the risky asset’s recent return. Think about  $\$1 \cdot r$  as the amount of money you’d have today if you had invested \$1 in the risky asset several years ago. If

$r = 0$ , then your entire \$1 investment would have been wiped out. Whereas, if  $r = 2$ , then you would have doubled your money. This second term captures the intuition that the arguments made by excited speculators are more persuasive following good past performance. If  $r \gg 1$ , then the remaining apathetic speculators have missed out on an investment opportunity that turned out to be very profitable ex post. And, investors frequently point to the fear of missing out as an important motivation for market participation. “Fear of missing out has become one of the forces spurring many Millennials to finally buy a home, according to a Bank of America survey of 2000 adults early this year.”<sup>4</sup> You can also see evidence of this behavior documented in [Bailey et al. \(2018\)](#).

Finally, the  $\theta \in (0, 1)$  represents how much an asset’s past returns affect the persuasiveness of its excited speculators. The idea is that past returns have a much bigger impact on speculator persuasiveness for some assets than for others. A 10% price increase in the real-estate market would generate a lot of additional word-of-mouth buzz, resulting in many new second-home buyers; by contrast, a 10% increase in the price of textile stocks would do nothing of the sort. I take  $\theta$  as an exogenous asset-specific constant, which encapsulates all of the idiosyncrasies that make one asset more interesting to talk about following good performance than another. See Section 3.2 for further discussion.

*Calming Down.* While popular accounts of bubble formation all emphasize that social interactions play a critical role in exciting new speculators during bubble episodes, they’re also all in agreement that the process by which excited speculators eventually come to their senses is a solo affair. The adage is that speculators “go mad in herds while they only recover their senses slowly and one by one” ([Mackay, 1841](#)). I use  $\Omega(n, r|\Delta\tau)$  to denote the probability that an excited speculator at time  $\tau$  happens to come to his senses during the same short time interval  $(\tau, \tau + \Delta\tau]$  for small  $\Delta\tau > 0$ :

$$\Omega(n, r|\Delta\tau) \stackrel{\text{def}}{=} \Pr[n_{\tau+\Delta\tau} - n_\tau = -1/K \mid n_\tau = n, r] \quad (3)$$

I use the following functional form to capture the essence of the intuition that excited speculators recover their senses slowly and one by one:

$$\lim_{\Delta\tau \searrow 0} \Omega(n, r|\Delta\tau) \stackrel{\text{def}}{=} \Omega(n) = 1 \times n \quad (4)$$

Multiplying  $n$  by the 1 equates the phrase “slowly and one by one” with a constant per capita departure rate from the crowd of excited speculators. In other words, the rate at which each excited speculator comes to his senses is the same regardless of whether 10% or 90% of all speculators are currently excited about the risky asset. And, because Equation (4) does not contain  $r$ , this rate is also independent of the risky asset’s past return. It’s the same if the asset’s recently done well or poorly. Note that the choice of  $\Omega(n) = 1 \times n$  rather than

---

<sup>4</sup>“Instagram, Facebook photos spur Millennials to become homeowners.” USA Today. Apr 11, 2018.

$\Omega(n) = \omega \times n$  for some positive constant  $\omega > 0$  is without loss of generality (see Appendix B.1). The key assumption is that the rate at which each excited speculator calms down is less sensitive to changes in both the size of the excited-speculator crowd,  $n$ , and the risky asset's past returns,  $r$ , than the rate at which currently apathetic speculators get excited. The functional forms in Equations (2) and (4) are a simple way of modeling this assumption.

*Logistic-Growth Model.* If we combine Equations (1) and (3), then we arrive at a master equation governing how the excited-speculator population evolves over time:

$$\frac{dn}{d\tau} = G(n, \theta, r) \stackrel{\text{def}}{=} \Theta(n, \theta, r) - \Omega(n) \quad \text{for } n \in [0, 1] \quad (5)$$

In other words,  $G(n, \theta, r)$  represents the rate at which the excited-speculator population grows when there are currently  $N_\tau = n_\tau \cdot K$  excited speculators in the market and the risky asset's past return is  $r$ . The functional-form choices in Equations (2) and (4), in turn, imply that this law of motion corresponds to what's known as the Velhurst model (Velhurst, 1845) or as the logistic-growth model with proportional harvesting (May, 1974):

$$G(n, \theta, r) = \theta \cdot r \cdot (1 - n) \times n - n \quad (6)$$

It'll sometimes be useful to work with a per capita version of the law of motion when  $n > 0$ :

$$g(n, \theta, r) \stackrel{\text{def}}{=} \frac{1}{n} \cdot G(n, \theta, r) = \theta \cdot r \cdot (1 - n) - 1 \quad \text{for } n \in (0, 1) \quad (7)$$

The quantity  $g(n, \theta, r)$  can be read as the rate at which new speculators get excited about the risky asset per member of the existing crowd of excited speculators at time  $\tau$ .

*Steady State.* We've just seen how the population of excited speculators evolves over time. So, I can now characterize the resulting steady-state behavior of this population. Let  $n_\tau(n_0, \theta, r)$  denote a solution to the initial-value problem associated with Equation (6):

$$n_\tau(n_0, \theta, r) \stackrel{\text{def}}{=} \left\{ n \in [0, 1) : n = \int_0^\tau G(n_{\tau'}, \theta, r) \cdot d\tau', \tau \geq 0 \right\} \quad (8)$$

In other words,  $n_\tau(n_0, \theta, r)$  provides the answer the question: "What fraction of speculators will be excited about the risky asset at time  $\tau \geq 0$  if  $n_0$  were excited at time  $\tau = 0$ ?" Standard texts (see Arnol'd, 2012) show that, if  $G(n, \theta, r)$  is continuously differentiable on an open interval that contains  $[0, 1)$ , then the solution  $n_\tau(n_0, \theta, r)$  will exist and be unique for all times  $\tau \geq 0$  and initial populations  $n_0 \in [0, 1)$ . And, assuming that  $\Theta(n, \theta, r) = \theta \cdot r \cdot (1 - n) \times n$  and  $\Omega(n) = n$  satisfies these requirements.

A steady-state value for the excited-speculator population is a population  $\bar{n} \in [0, 1)$  such that  $n_\tau(\bar{n}, \theta, r) = \bar{n}$  for all  $\tau \geq 0$ :

$$\mathcal{SS}(\theta, r) \stackrel{\text{def}}{=} \left\{ \bar{n} \in [0, 1) : n_\tau(\bar{n}, \theta, r) = \bar{n} \quad \forall \tau \geq 0 \right\} \quad (9)$$

We say that a particular steady-state value,  $\bar{n} \in \mathcal{SS}(\theta, r)$ , is stable if small perturbations away from  $\bar{n}$  die out over time. More precisely,  $\bar{n} \in \mathcal{SS}(\theta, r)$  is stable if for every  $\delta > 0$  there's some  $\epsilon > 0$  such that the solution to the initial-value problem in Equation (8) satisfies

$|n_\tau(n_0, \theta, r) - \bar{n}| < \delta$  for all  $\tau \geq 0$  given any initial population  $n_0 \in (\bar{n} - \epsilon, \bar{n} + \epsilon)$ .

*Sudden Qualitative Change.* Given the functional form of Equation (6), what should we expect the steady-state solution for the excited-speculator population to look like? When  $\Theta(n, \theta, r) = \theta \cdot r \cdot (1 - n) \times n$  and  $\Omega(n) = n$ , there's nothing to suggest a sudden change in the character of the steady-state solution as the asset's past returns change. There are no sudden discontinuous jumps built into the model. Excited speculators are always coming to their senses at the same rate. And, the effect of an asset's past returns on its speculators' persuasiveness is continuous and smooth. Slightly higher past returns always make the current group of excited speculators slightly more persuasive to each of their friends. So, you might expect that a slight increase in the risky asset's past return would always result in a slight increase in the steady-state excited-speculator population.

But, this is not what happens.

Equation (6) actually produces a sudden qualitative change in the behavior of the steady-state solution as an asset's past return crosses a critical threshold. This large qualitative change in the steady-state behavior is called a “bifurcation” (Hirsch et al., 2012; Guckenheimer and Holmes, 2013; Kuznetsov, 2013; Strogatz, 2014). And, the proposition below analytically characterizes the bifurcation that occurs at the critical return threshold,  $r_\star$ , when the population of excited speculators is governed by the logistic-growth model.

**Proposition 2.1** (*Sudden Qualitative Change*). *Suppose that the excited-speculator population is governed by the law of motion in Equation (6). Define  $r_\star \stackrel{\text{def}}{=} 1/\theta$ .*


1. *If  $r < r_\star$ , there's only one steady-state value for the excited-speculator population,  $\mathcal{SS}(\theta, r) = \{0\}$ . And, this lone steady state,  $\bar{n} = 0$ , is stable.*
2. *If  $r > r_\star$ , there are two steady-state values,  $\mathcal{SS}(\theta, r) = \{0, (r - r_\star)/r > 0\}$ . However, only the strictly positive steady state,  $\bar{n} = (r - r_\star)/r > 0$ , is stable.*

*When the risky asset's past return is low enough,  $r < r_\star$ , an initial population of speculators,  $n_0 > 0$ , that happens to get excited about the risky asset will quickly lose interest and disperse. The excited speculator population will be transient. But, as soon as the risky asset's return crosses a critical threshold,  $r > r_\star$ , that same initial population will suddenly give rise to a persistent crowd of excited speculators.*

*Economic Intuition.* To see where this sudden qualitative change comes from, consider what happens when there's only one excited speculator,  $N_\tau = n_\tau \cdot K = 1$ . In this situation, the entire population of excited speculators will go extinct if its lone member can't excite at least one of his apathetic friends before he himself comes to his senses:

$$\underbrace{\Pr[\Delta N_\tau = +1 \mid N_\tau = 1, r]}_{=\theta \cdot r \cdot (1-1/K) \cdot 1 \cdot \Delta\tau \approx \theta \cdot r \cdot \Delta\tau} < \underbrace{\Pr[\Delta N_\tau = -1 \mid N_\tau = 1]}_{=1 \cdot \Delta\tau} \quad (10)$$





**Figure 4. Sudden Qualitative Change. Top Panel.** Steady-state solutions.  $x$ -axis: past return of the risky asset,  $r \in (0, \infty)$ .  $y$ -axis: steady-state solutions,  $\bar{n} \in \mathcal{SS}(\theta, r)$ , for a population of excited speculators governed by the law of motion in Equation (6). Solid black line reports stable steady states; dashed red line reports unstable ones. Population displays a bifurcation at  $r_* = 1/\theta$  as characterized in Proposition 2.1. **Bottom Panels.** Transition to steady state when  $r < r_*$  vs. when  $r > r_*$ .  $x$ -axis, top: fraction of speculators who are currently excited about the risky asset,  $n$ .  $y$ -axis, top: growth rate of excited-speculator population,  $G(n, \theta, r) = \frac{dn}{d\tau}$ . When  $r < r_*$ , this growth rate is always negative for all  $n > 0$  as indicated by the solid line remaining below the  $x$ -axis. By contrast, when  $r > r_*$ , this growth rate is positive for some population values,  $n > 0$ , as indicated by the solid line arching above the  $x$ -axis.  $x$ -axis, bottom: time since initial group of  $n_0 \geq 0$  speculators got excited about the risky asset at time,  $\tau = 0$ .  $y$ -axis, bottom: number of speculators excited about the risky asset at time  $\tau > 0$ ,  $n_\tau = n_\tau(n_0, \theta, r)$ . Different shades of grey denote different initial population sizes,  $n_0 \in [0, 1]$ . When  $r < r_*$ , any initial population of speculators  $n_0 > 0$  that happens to get excited about the risky asset will quickly lose interest and disperse so that  $n_\tau(n_0, \theta, r) \rightarrow \bar{n} = 0$ . But, when  $r > r_*$ , the excited-speculator population will converge in size to  $\bar{n} = (r - r_*)/r > 0$  whenever a single speculator happens to get excited about the risky asset. Note that, when  $r > r_*$ ,  $\bar{n} = 0$  is still a steady state. But, now this value is unstable as shown by the dashed red line.

The expressions underneath the braces characterize these probabilities when  $N_\tau = 1$  given the functional forms in Equations (2) and (4). By rearranging terms, you can see that there will be no excited speculators left in the market whenever the asset’s past return is sufficiently low,  $r < r_\star = 1/\theta$ , as shown in the left half of Figure 4. However, as soon as the asset’s past return rises above the critical threshold,  $r > r_\star$ , the exact opposite intuition will hold. If a single speculator happens to get excited about the risky asset, then this lone agent will likely be able to excite a friend before he himself comes to his senses. The exact same logic will also then apply to new partner in crime. So, the excited-speculator population will initially grow exponentially fast, remaining above zero in steady state as shown in the right half of Figure 4. Thus, the excited-speculator population can exhibit a sudden qualitative change in behavior as the asset’s past returns cross  $r_\star$  even though this population is always governed by a smoothly changing set of rules.

*Novel Recombination.* Although the notion of speculators “going mad in herds while only recovering their senses slowly and one by one” (Mackay, 1841) has been present in popular accounts dating back to the 1800s, the current model is using this idea in a novel way. In the past, the intuition that speculators recover their senses slowly and one by one was an excuse for why otherwise reasonable people might behave irrationally for a long time during bubble episodes. More is different. A crowd of speculators can be affected by forces that no individual member can feel on his own. So, “slowly and one by one” was a way of taking “the strangeness of collective behavior out of the heads of individual actors and putting it into the dynamics of situations” (Granovetter, 1978).

That’s not what’s going on here. Instead of using the fact that excited speculators come to their senses slowly and one by one to explain the persistence of mass hysteria, this paper uses it as an explanation for why mass hysteria only occasionally takes hold. I’m using it as an off switch rather than an amplifier. If excited speculators are always calming down independently and at the same rate, then the crowd of excited speculators will always decay at the same rate. So, even if social interactions might sometimes lead to the madness of crowds, they will remain irrelevant so long as their effects are weaker than this lower bound.

To illustrate, consider what would go wrong if we ignored the rate at which excited speculators calmed down slowly and one by one. Suppose we only focused on feedback trading and defined  $\tilde{G}(n, \theta, r) \stackrel{\text{def}}{=} \Theta(n, \theta, r) = \theta \cdot r \cdot (1 - n) \times n$  rather than  $G(n, \theta, r) = \Theta(n, \theta, r) - \Omega(n) = \theta \cdot r \cdot (1 - n) \times n - n$ . In this alternative model, the only steady-state solutions for the excited-speculator population would be  $\widehat{\mathcal{S}}(\theta, r) = \{0, 1\}$ . In other words, in a model with only feedback trading, either there are no speculators excited about the risky asset,  $\bar{n} = 0$ , or every single speculator is excited about the risky asset,  $\bar{n} = 1$ . There’s nothing in between. What’s more, in a model without the calming effect, only the steady-state

solution at  $\bar{n} = 1$  is stable—i.e., if a single speculator happens to get excited about the risky asset, then every speculator will quickly join him regardless of an asset’s past performance. And, that’s just not what we see in the data. Large speculative bubbles are rare.

*Displacement Events.* Popular accounts of bubble formation have had an exceptionally hard time defining what constitutes a displacement event. While “only a relatively small proportion of shocks lead to a speculative mania” (Kindleberger, 1978), “an event that is not of unusual size or duration can trigger a sharp financial reaction” (Minsky, 1970). What’s more, the exact same shock—e.g., a change in interest rates—might trigger the madness of crowds for one asset but not another. At first glance, all these irregularities are deeply “unsatisfying to [anyone] seeking scientific certitude” (Shiller, 2000).

By contrast, the analysis above provides a clear definition of what constitutes a displacement event:

**Definition 2.1** (*Displacement Event*). Suppose  $r < r_\star = 1/\theta$ . A displacement event is a positive shock,  $\epsilon > 0$ , that increases the risky asset’s return such that  $r \mapsto r_\epsilon \stackrel{\text{def}}{=} r \cdot (1 + \epsilon) > r_\star$ .

And, this definition naturally resolves all of these apparent irregularities. Only a relatively small proportion of shocks will lead to a speculative mania because even large shocks,  $\epsilon \gg 0$ , will only rarely be large enough to push an asset’s returns above a high threshold level,  $r_\star = 1/\theta$ . But, if an asset’s past returns are already close to this threshold level, then even an arbitrarily small positive shock will be enough to push the asset over the edge. Finally, the exact same shock can have different effects for assets with different values of  $\theta$ —i.e., the parameter reflecting the sensitivity of speculator persuasiveness to an asset’s past returns. Later, we’ll explore this point in much greater detail in the empirical analysis.

## 2.2 Asset Prices

I now embed this time-varying population of excited speculators in an otherwise standard limits-to-arbitrage model to see how their biased demand distorts equilibrium asset prices. The idea will be to study a discrete-time environment where the number of excited speculators in each period,  $n_t$ , is equal to the steady-state population size described in Proposition 2.1 where  $r$  corresponds to the risky asset’s realized return in the previous period.

*Timing and Payouts.* Imagine a market that proceeds in discrete steps indexed by  $t = 1, 2, 3, \dots$ . I will use  $t$  rather than  $\tau \geq 0$  to denote discrete time increments. The market contains a single risky asset with a payout of  $v_t$  dollars per share in period  $t$ . The size of this payout evolves over time as follows:

$$\Delta v_t = \kappa_v \cdot (\mu_v - v_{t-1}) + \sigma_v \cdot \varepsilon_{v,t} \quad (11)$$

In the equation above,  $\mu_v \gg 0$  is the risky asset’s average payout per period,  $\kappa_v \in (0, 1)$  is its

mean-reversion coefficient,  $\sigma_v > 0$  is the volatility of changes in the risky asset's fundamental value each period, and  $\varepsilon_{v,t} \sim N(0, 1)$  is a shock that's independently and identically distributed across time. Assume there are  $\psi \geq 0$  shares of the risky asset available for purchase.

*News watchers.* The market contains two kinds of agents: newswatchers and speculators. Newswatchers incorporate news about the risky asset's fundamental value into its price. These agents live for exactly one period, so every period there's a new unit mass of newswatchers indexed by  $j \in [0, 1]$ . The  $j$ th newswatcher in period  $t$  chooses his demand,  $x_{j,t}$ , so as to maximize his expected end-of-period utility from consuming the risky asset's time- $t$  payout:

$$x_{j,t} = \arg \max_x E_j \left[ -e^{-\gamma \cdot (v_t - p_t) \cdot x} \right] \quad (12)$$

Newswatcher utility displays constant relative risk-aversion. Above,  $E_j[\cdot]$  denotes the expectation of the  $j$ th newswatcher after observing a private signal (see next paragraph), and  $\gamma > 0$  denotes his risk-aversion coefficient, which is the same for every newswatcher in every cohort.

The newswatchers within each cohort have heterogeneous beliefs. Prior to the start of period  $t$ , the  $j$ th newswatcher observes a private signal about the risky asset's time- $t$  payout. After observing this signal, the  $j$ th newswatcher believes the risky asset's payout can be decomposed as follows

$$v_t = s_{j,t} + \varepsilon_{j,t} \quad (13)$$

where  $s_{j,t}$  denotes the  $j$ th newswatcher's beliefs after observing his private signal and  $\varepsilon_{j,t} \sim N(0, 1)$  denotes noise in his beliefs, which is independent and identically distributed across newswatchers within each cohort.

There are two important details to notice about this information structure. First, the heterogeneous beliefs of each cohort of newswatchers' are correct on average:

$$v_t = E[s_{j,t}] = \int_0^1 s_{j,t} \cdot dj \quad (14)$$

Second, although different newswatchers get different private signals, every newswatcher's private signal has the same unit precision. For readers familiar with the literature, this information structure mirrors the one used in [Admati \(1985\)](#).

*Excited Speculators.* In addition to newswatchers, the market also contains  $n_t \geq 0$  excited speculators each period. The role of newswatchers in the model is to cause the price of the risky asset to move in response to changes in its fundamental value. By contrast, the role of excited speculators is to sometimes cause the risky asset's price to move for non-fundamental reasons. The number of excited speculators in the market varies over time depending on the assets' realized return in the previous period. Specifically, suppose that  $n_t$  is given by the formula in Proposition 2.1 where  $r = r_{t-1}$  and  $r_{t-1} \stackrel{\text{def}}{=} p_{t-1}/p_{t-2}$ :

$$n_t = \begin{cases} (r_{t-1} - r_*)/r_{t-1} & \text{if } r_{t-1} > r_* = 1/\theta \\ 0 & \text{otherwise} \end{cases} \quad (15)$$

The idea is to define  $r_{t-1}$  so that  $\$1 \cdot r_{t-1}$  represents the amount of money a speculator would have at the start of time  $t$  if they had invested  $\$1$  in the risky asset at time  $(t-2)$ . If  $r_{t-1} = 0$ , then they'd have lost their entire  $\$1$  initial investment. Whereas, if  $r_{t-1} = 2$ , then they would have doubled their money.

The approach is tantamount to assuming that the continuous-time population dynamics described above play out once per discrete period  $t$  in this asset-pricing model. It's as if at the start of each discrete time period  $t$ , a single speculator gets excited about the risky asset. After he enters, the madness of crowds either takes over or doesn't depending on both the asset's past return,  $r = r_{t-1}$ , and the asset-specific value for  $\theta$ . The excited speculator population that happens to still be there at the end of the discrete time period is given by the steady-state solution in Proposition 2.1. And, the remaining crowd of excited speculators (if any) is responsible for the non-fundamental demand shock the asset realizes in period  $t$ . This model-timing assumption simplifies the analysis by separating population dynamics from price determination and explains why there's no time subscript on  $r$  in Equation (8). And, by doing so, the assumption further emphasizes the distinction between the forces that cause a speculative bubble to form and those that allow a speculative bubble to be sustained. However, I show in Appendix B.3 that the main predictions of the model carry over to a setting where there's continuous feedback between excited-speculator demand and prices.

What does this non-fundamental demand shock look like? Suppose there are  $n_t > 0$  excited speculators in the market. I model the aggregate demand coming from this crowd,  $z_t$ , as proportional to the asset's past returns:

$$z_t = (\lambda \cdot r_{t-1}) \times n_t \quad (16)$$

Thus, whenever a crowd of  $n_t > 0$  speculators gets excited about the risky asset, they choose their demand based on extrapolative beliefs (Cutler et al., 1990; De Long et al., 1990; Hong and Stein, 1999; Barberis and Shleifer, 2003; Barberis et al., 2015; Glaeser and Nathanson, 2017; Barberis et al., 2018). This is a purposeful choice. The goal here is not to break new ground by suggesting a new psychological bias that afflicts speculators. Rather, the goal is to show how social interactions between speculators can sporadically amplify the effects of their omnipresent psychological biases.

*Equilibrium.* I look for a Walrasian equilibrium with private valuations. For markets to clear, the aggregate demand coming from newswatchers and any excited speculators must sum up to the total number of shares:

$$\int_0^1 x_{j,t} \cdot dj + z_t = \psi \quad (17)$$

This market-clearing condition pins down the equilibrium price. So, if newswatchers were fully rational, they could each learn about each others' signals by conditioning their beliefs on the

equilibrium price. However, to allow for equilibrium pricing errors, I assume that newswatchers don't do this; instead, they ignore the information content of prices so that the  $j$ th newswatcher in period  $t$  chooses his demand according to  $x_{j,t} = \arg \max_x E[-e^{-\gamma \cdot (v_t - p_t) \cdot x} | s_{j,t}]$ .

This assumption is not new to this paper. [Hong and Stein \(1999\)](#) motivate this modeling choice as a tractable way to represent bounded rationality on the part of arbitrageurs. You can think about each newswatcher as having his hands full acquiring the private signal,  $s_{j,t}$ . This leaves them without enough mental bandwidth to incorporate any additional information from prices. And, [Eyster et al. \(2018\)](#) show how to relax this assumption so that newswatchers pay too little attention to prices instead of completely ignoring them.

Speculative bubbles are equilibrium pricing errors. And, to have an equilibrium pricing error, a model needs to incorporate some sort of limits to arbitrage. I use this particular form because it's straightforward and clean. Assuming that newswatchers ignore the information content of prices doesn't require me to introduce any additional parameters when modeling the limits to arbitrage. But, the exact way in which newswatchers are constrained is unimportant for our purposes. It's possible, for example, to write down a version of the model where newswatchers are fully rational but face short-sale constraints. Again, the goal is not to add a new constraint to the already voluminous limits-to-arbitrage literature. There are already many different psychological biases and trading constraints to choose from. The goal is to provide a theory of when and where the limit of arbitrage are most likely to bind.

*Asset Prices.* I can now solve the model. Given the payout and information structure, newswatchers use the following demand rule:

$$x_{j,t} = (s_{j,t} - p_t) / \gamma \quad (18)$$

If his private signal results in beliefs that are higher than the price,  $(s_{j,t} - p_t) > 0$ , then he buys; if they result in beliefs that are lower than the price,  $(s_{j,t} - p_t) < 0$ , then he sells. And, given this optimal demand rule, the proposition below characterizes the equilibrium price.

**Proposition 2.2** (*Asset Prices*). *The risky asset's equilibrium price is given by:*

$$p_t = v_t - \gamma \times \psi + \gamma \cdot (\lambda \cdot r_{t-1}) \times n_t \quad (19)$$

*This price is increasing in the fundamental value,  $v_t$ ; it's decreasing in the number of shares,  $\psi$ ; and, it's increasing in the number of speculators excited about the asset,  $n_t \geq 0$ .*

Each term on the right-hand side of Equation (19) has a clear economic interpretation. First, consider  $v_t$ . If the risky asset's fundamental value rises by \$1, then the correct-on-average assumption in Equation (14) implies that the mean newswatcher signal will increase by \$1. Thus, the price of the risky asset will also rise by \$1. Next, consider the term,  $-\gamma \times \psi$ . Because newswatchers are risk averse,  $\gamma > 0$ , they have to be compensated for holding shares of the risky asset in equilibrium whenever the risky asset is in positive supply,  $\psi > 0$ . If the



asset were in zero net supply,  $\psi = 0$ , this term would disappear.

Finally, consider  $\gamma \cdot (\lambda \cdot r_{t-1}) \times n_t$ , which represents the effect of the aggregate demand coming from any excited speculators in the market. Three different forces are at work here. The first force is the limits to arbitrage. If newswatchers were fully rational, then they would condition on the information content of the equilibrium price. And, this price would fully reveal the fundamental value of this risky asset  $v_t$  (Grossman, 1976). So, without the limits of arbitrage, all risk in the model would disappear. Only the  $v_t$  term would remain,  $p_t = v_t$ . The number of speculators who are currently excited about the risky asset would be irrelevant. The limits of arbitrage are what make it possible for the non-fundamental demand coming from excited speculators to affect equilibrium prices. The second force is a behavioral bias. This force is what pins down the size of the non-fundamental demand shock coming from excited speculators whenever they're present in the market. It's what determines the functional form of the parenthetical term,  $\lambda \cdot r_{t-1}$ , in Equations (16) and (19). Both of these first two forces are standard in the behavioral-finance literature.

The third force is not. This force is the madness of crowds as represented by whether  $n_t = 0$  or  $n_t > 0$  in Equation (16). It's a result of allowing the excited-speculator population to vary over time due to social interactions. This force controls whether speculators' omnipresent bias (in this model: extrapolative beliefs) will cause arbitrageurs' omnipresent trading constraint (in this model: not conditioning on prices) to bind as  $n_t$  flips from being precisely zero when  $r_{t-1} < r_\star = 1/\theta$  to being positive whenever  $r_{t-1} > r_\star$ .


*Speculative Bubble.* A speculative bubble occurs in the model when newswatchers happen to push the realized return of the risky asset above the critical threshold,  $r_{t-1} > r_\star$ , causing a non-zero crowd of excited speculators to flood the market. I use the indicator variable

$$B(\theta, r_{t-1}) \stackrel{\text{def}}{=} 1[r_{t-1} > 1/\theta] \quad (20a)$$

$$= 1[n_t > 0] \quad (20b)$$

to denote whether there's a non-zero population of excited speculators in the market at time  $t$ —i.e., to distinguish whether this third force has any impact on equilibrium prices.

*Sample Price Path.* It's easiest to understand the asset-pricing implications of this model by simulating a market and studying the realized outcome. Figure 5 depicts one such sample realization using the parameter values  $\psi = 0$ ,  $\mu_v = 1.0$ ,  $\kappa_v = 0.1$ ,  $\sigma_v = 0.1$ ,  $\theta = 0.4$ , and  $\lambda = 0.5$ . First, take a look at the top panel. The black line depicts the risky asset's equilibrium price each period,  $p$ , while the thin green line depicts this same asset's fundamental value,  $v$ . Because the risky asset is in zero net supply,  $\psi = 0$ , these two lines fall right on top of one another when there are no excited speculators in the market,  $n = 0$ . And, this is exactly what happens most of the time. In the middle panel, the black line represents the risky asset's



**Figure 5. Sample Realization.** Data simulated using parameters  $\psi = 0$ ,  $\mu_v = 1.0$ ,  $\kappa_v = 0.1$ ,  $\sigma_v = 0.1$ ,  $\theta = 0.4$ , and  $\lambda = 0.5$ .  $x$ -axis in all panels represents time,  $t = 1, 2, \dots, 100$ . **Top Panel.** Black line is risky asset's price,  $p$ . Thin green line is its fundamental value,  $v$ . Red shaded regions denote times when the excess demand from excited speculators caused arbitrageur constraints to bind and a speculative bubble to form. **Middle Panel.** Black line is risky asset's realized return,  $r_t = p_t/p_{t-1}$ . Dashed blue line is threshold return level,  $r_* = 1/\theta$ . When  $r < r_*$ , there are no excited speculators in the market,  $n = 0$ . When  $r > r_*$ , there's a non-zero population of excited speculators in the market,  $n > 0$ . **Bottom Panel.** Red vertical bars report the number of excited speculators in the market,  $n$ . The four instances where  $B(\theta, r) = 1$  are labeled on the  $x$ -axis,  $\{t_1, t_2, t_3, t_4\}$ .

realized return in the previous period,  $r_t = p_t/p_{t-1}$ . You can see that this value is typically below the dashed blue line representing  $r_* = 1/\theta$ .

That being said, there are four different points in time where newswatchers pushed the risky asset's return above the threshold level—i.e., where  $B(\theta, r) = 1$ . Each of these instances is denoted by a label,  $\{t_1, t_2, t_3, t_4\}$ , on the  $x$ -axis in the bottom panel. Recall that  $B(\theta, r) = 1$  implies that  $r > r_* = 1/\theta$ , which in turn implies that there's a non-zero population of excited speculators in the market,  $n > 0$ . The height of the red bars in the bottom panel depicts the size of the excited-speculator crowd,  $n$ . The empirical analysis below is going to investigate how often we should expect to see these episodes—i.e., whether we should expect four episodes,  $\{t_1, t_2, t_3, t_4\}$ , or just two,  $\{t_1, t_2\}$ .

## 2.3 Discussion

The existing limits-to-arbitrage framework explains *how* a pricing error such as a speculative bubble can be sustained in equilibrium. But, it does not explain *how often* we should expect to see one of these errors. The model above introduces a mechanism that sporadically amplifies

the effect of speculators’ omnipresent psychological biases. To emphasize this *how-vs.-how often* distinction, I’ve deliberately kept the model as simple as possible. However, there are some natural extensions you might consider. I now examine three of these extensions and show that they do not qualitatively change the implications of the model.

*Functional Forms.* First, the law of motion in Equation (6) takes a very particular functional form. So, you might ask: how specific are the results in this paper to this choice? It turns out that, to a second-order approximation, the growth rate in Equation (6) can be thought of as a stand-in for a broad range of models displaying the same steady-state behavior. It embodies “activity that is self-sustaining once the measure of that activity passes a certain minimum level” (Schelling, 1978). For example, using  $\tilde{G}(n, \theta, r) \stackrel{\text{def}}{=} r \cdot (1 - e^{-\theta \cdot n}) - n$  would deliver qualitatively similar population dynamics as using  $G(n, \theta, r)$  even though  $\tilde{G}(n, \theta, r)$  and  $G(n, \theta, r)$  look superficially very different. In Appendix B.1 I give conditions under which an alternative law of motion,  $\tilde{G}(n, \theta, r)$ , will exhibit feedback trading just like  $G(n, \theta, r)$ .

*Random Fluctuations.* You might also be wondering: what would happen if the excited-speculator population followed a stochastic law of motion? In the presence of random fluctuations, it’s possible that the sudden change in the steady-state excited-speculator population as the risky asset’s past return crosses  $r_\star = 1/\theta$  disappears. Shaking an Etch A Sketch removes sharp lines. I answer this question in Appendix B.2 by redefining the law of motion in Equation (6) as follows:

$$G(n, \theta, r) \mapsto \tilde{G}(n, \theta, r) \stackrel{\text{def}}{=} G(n, \theta, r) + \sigma_n \cdot n \cdot \frac{d\varepsilon_g}{d\tau} \quad (21)$$

In the equation above,  $\sigma_n > 0$  is a positive constant reflecting the instantaneous volatility of the excited-speculator population growth rate, and  $\varepsilon_{g,\tau} \sim N(0, 1)$  is a white-noise process. The stationary distribution for an excited-speculator population governed by this stochastic law of motion displays a sudden change in character as past returns cross the threshold value  $r_\star = 1/\theta$  just as in the deterministic case. When  $r < r_\star$ , any initial population of excited speculators almost surely goes extinct; whereas, when  $r > r_\star$ , this is no longer the case. Adding noise does not eliminate the sharp change in the population dynamics around  $r_\star$ .

*Continuous Feedback.* Finally, in the model above, speculator interactions play out on a much faster timescale than assets are priced. First, speculators observe the risky asset’s return in the previous period. Then, they interact with one another until a steady-state population has been reached. Finally, after this steady-state has been reached, any remaining excited speculators submit their demand to the market. You might wonder: what would happen if short-run changes in the excited-speculator population affected the risky asset’s returns—i.e., what would change if there was continuous feedback between the excited-speculator population and the risky asset’s return?

To answer this question, in Appendix B.3 I study an alternative law of motion where a 1% increase in the population of excited speculators increases the risky asset’s return by a factor of  $\epsilon \in [0, \frac{1}{\theta \cdot r})$ :

$$\tilde{G}(n, \theta, r) \stackrel{\text{def}}{=} \theta \cdot r \cdot (1 + \epsilon \cdot n) \cdot (1 - n) \times n - n \quad (22)$$

The  $(1 + \epsilon \cdot n)$  term in the equation above captures the idea that an inflow of excited speculators at time  $\tau$  will increase the risky asset’s return, which will then make it easier for future excited speculators to recruit their friends. If we set  $\epsilon = 0$ , then we get back the original law of motion in Equation (6). By increasing  $\epsilon$ , we allow transient fluctuations in the excited-speculator population to have a larger and larger effect on speculator persuasiveness via their effect on the asset’s past returns.

While modeling the continuous feedback between population dynamics and asset returns might seem more realistic, it turns out that this extension only affects the size of the steady-state excited-speculator population conditional on entering. It doesn’t affect the threshold,  $r_\star = 1/\theta$ , at which a non-zero population of excited speculators suddenly enters the market. And, this paper is entirely about where this threshold comes from. Understanding the dynamics of the excited-speculator population interacts with an asset’s past returns is very important on the intensive margin. It’s very important if you want to understand how any particular bubble episode will unfold. But, it’s not very important on the extensive margin. It’s not essential if all you want to do is understand the likelihood that a crowd of excited speculators will enter the market in the first place. Allowing the steady-state excited-speculator population and asset prices to be determined on two vastly different timescales is a deliberate modeling choice made to highlight this distinction.

### 3 Empirical Implication

This section explores the key empirical implication of this model: assets with higher values of  $\theta$  should experience speculative bubbles more often. What’s more, it’s possible to estimate  $\theta$  for each asset using data collected during normal times. This is new and different. If you want to learn about speculative bubbles in the limits-to-arbitrage framework, then you have to wait for arbitrageurs’ constraints to bind. However, in this paper, the exact same social interactions are at work both during and between bubble episodes. So, it’s possible to estimate  $\theta$  for each asset before any limits to arbitrage kick in.

#### 3.1 Cross-Sectional Prediction

I begin by developing the model’s main cross-sectional prediction.

*Bubble Likelihood.* Which assets are most likely to experience a speculative bubble follow-

ing good past performance? In the model, a speculative bubble occurs when good news about an asset's fundamental value causes newswatchers to push the asset's return above the critical performance threshold,  $r > r_*$ . When this happens, a crowd of excited speculators enters the market,  $n > 0$ . And, the extrapolative demand coming from these excited speculators causes arbitrageurs' constraint to bind and a speculative bubble to form. Thus, to predict the likelihood of a speculative bubble, we need to answer three questions: i) How is the critical performance threshold,  $r_*$ , determined for each asset? ii) How do newswatchers incorporate information about an asset's fundamental value into its price during normal times? And, iii) how does the risky asset's fundamental value fluctuate over time?

The analysis in the previous section answers all three of these questions: i) Proposition 2.1 tells us how the critical performance threshold is set,  $r_* = 1/\theta$ . ii) Proposition 2.2 tells us how newswatchers link the price of a risky asset to its fundamental value during normal times when there's no crowd of excited speculators in the market,  $p_t = v_t - \gamma \times \psi$ . And, iii) we know how the risky asset's fundamental value fluctuates over time from Equation (11):

$$E_{t-1}[\Delta v_t] = \kappa_v \cdot (\mu_v - v_{t-1}) \quad (23a)$$

$$\text{Var}_{t-1}[\Delta v_t] = \sigma_v^2 \quad (23b)$$

So, the model can be used to compute how changes in  $\theta$  will affect the probability that newswatchers will push the return of the risky asset across an asset's critical return threshold,  $\Pr_{t-1}[r_t > r_* \mid r_{t-1} < r_*] = E_{t-1}[B(\theta, r_t) \mid r_{t-1} < r_*]$ .


**Proposition 3.1** (*Bubble Likelihood*). *The probability that fundamental news will push the return of the risky asset above  $r_* = 1/\theta$  is strictly increasing in  $\theta$ :*

$$\frac{\partial}{\partial \theta} E_{t-1}[B(\theta, r_t) \mid r_{t-1} < r_*] > 0 \quad (24)$$

*If there are two assets with identical fundamental parameters  $(\mu_v, \kappa_v, \sigma_v)$  and past performance,  $r_{t-1} < r_* = 1/\theta$ , then the asset with the higher speculator-persuasiveness sensitivity,  $\theta$ , is more likely to experience a speculative bubble in the future.*

*Economic Intuition.* When house prices rise by 10%, it's national news. Everyone starts talking about housing. It's a lot easier for second-home buyers to convince their friends to speculate in the housing market, too. By contrast, while a 10% increase in the price of textile stocks is a big deal for market participants, it's not going to excite many uninformed agents. Thus, given the same initial conditions, the housing market should be more likely to experience a speculative bubble than the textile industry going forward.

Figure 6 displays this intuition visually. Both panels show the first 25 periods from Figure 5, which was simulated using parameters  $\psi = 0$ ,  $\mu_v = 1.0$ ,  $\kappa_v = 0.1$ ,  $\sigma_v = 0.1$ , and  $\lambda = 0.5$ . In Panel 6a, the sensitivity of speculator persuasiveness to the risky asset's past return is  $\theta = 0.4$  just like in Figure 5. So, the risky asset's critical threshold in the left panel is given




**Figure 6. Bubble Likelihood.** First 25 periods from Figure 5, which was simulated using parameters  $\psi = 0$ ,  $\mu_v = 1.0$ ,  $\kappa_v = 0.1$ ,  $\sigma_v = 0.1$ , and  $\lambda = 0.5$ .  $x$ -axis: time  $t = 1, 2, \dots, 25$ . Black line, top: risky asset's price,  $p$ . Green line, top: risky asset's fundamental value,  $v$ . Black line, middle: risky asset's realized return,  $r_t = p_t/p_{t-1}$ . Dashed blue line, middle: threshold return level,  $r_* = 1/\theta$ . Red bars, bottom: number of excited speculators in the market,  $n$ .  $\theta = 0.4$ . Sensitivity of speculator persuasiveness to past returns is the same as in Figure 5, so there's a speculative bubble at time  $t_1$  just like in the original figure since  $r_{t_1-1} > 1/0.4 = 2.5$ .  $\theta = 0.2$ . Sensitivity of speculator persuasiveness to past returns is lower than in Figure 5, which means that the critical return threshold is higher than in Figure 5. And, no speculative bubble occurs at time  $t_1$  because  $r_{t_1-1} < 1/0.2 = 5.0$ .

by  $r_* = 1/0.4 = 2.5$  just like before. There will only be a crowd of excited speculators in the market at time  $t$  when the risky asset realizes a  $\geq 150\%$  return. This is exactly what happens at time  $t_1$ , which resulted in a speculative bubble.

The only difference between Panel 6a and Panel 6b is that the risky asset's key sensitivity parameter is lower in Panel 6b. It's  $\theta = 0.2$  rather than the original  $\theta = 0.4$  in and Panel 6a. Everything else about Panel 6a and Panel 6b is utterly identical—the underlying parameter values are the same, and the innovations in fundamental value are the same. But, because the key sensitivity is lower, the risky asset's critical return threshold is higher,  $r_* = 1/0.2 = 5.0$ . As a result, in Panel 6b there will only be a crowd of excited speculators in the market when the risky asset realizes a  $\geq 400\%$  return. And,  $r_{t_1-1} < 5.0$  in this simulation. Thus, in Panel 6b, there's no speculative bubble at time  $t_1$  when  $\theta = 0.2$ .

*How vs. How Often.* A key property of the cross-sectional prediction in Proposition 3.1 is that it doesn't depend on the severity of excited speculators' extrapolative beliefs—i.e., it doesn't depend on the size of  $\lambda > 0$ . This observation captures the idea that ‘*How often should we expect to observe a speculative bubble?*’ is a fundamentally distinct question from





**Figure 7. How vs. How Often.** Periods 40,  $\dots$ , 80 from Figure 5, which was simulated using parameters  $\psi = 0$ ,  $\mu_v = 1.0$ ,  $\kappa_v = 0.1$ ,  $\sigma_v = 0.1$ , and  $\theta = 0.4$ .  $x$ -axis: time  $t = 40, 41, \dots, 80$ . Black line, top: risky asset's price,  $p$ . Green line, top: risky asset's fundamental value,  $v$ . Black line, middle: risky asset's realized return,  $r_t = p_t/p_{t-1}$ . Dashed blue line, middle: threshold return level,  $r_\star = 1/\theta$ . Red bars, bottom: number of excited speculators in the market,  $n$ .  $\lambda = 0.50$ . Severity of excited speculators' extrapolation is the same as in Figure 5, so the speculative bubble at time  $t_2$  lasts one period, the speculative bubble at time  $t_3$  lasts three periods, and the speculative bubble at time  $t_4$  lasts one period just like in the original figure.  $\lambda = 0.75$ . Severity of excited speculators' extrapolation is higher than in Figure 5, which means that the speculative bubbles at time  $t_2$  and  $t_3$  now last longer than in Figure 5. The speculative bubble at time  $t_2$  now lasts two periods rather than just one. And, The speculative bubble at time  $t_3$  lasts four periods rather than three. But, this change does not affect the number of speculative bubbles that occur as dictated by Corollary 3.1.

*'How can a speculative bubble be sustained in equilibrium?'*

**Corollary 3.1** (*How vs. How Often*). *The probability that fundamental news will push the return of the risky asset above  $r_\star = 1/\theta$  is independent of  $\lambda$ :*

$$\frac{\partial}{\partial \lambda} \mathbb{E}_{t-1}[\mathbf{B}(\theta, r_t) \mid r_{t-1} < r_\star] = 0 \quad (25)$$

The exact limits-to-arbitrage model you choose will affect how a speculative bubble will unfold once it starts. It will pin down the lifecycle of a speculative bubble. But, it can't tell you how often one of these events will be set in motion. It can't tell you how often market conditions will give birth to a speculative mania. You need a different kind of theory if you want to explain the likelihood that some limit to arbitrage will bind. As illustrated in Figure 7, increasing the severity of excited speculators' extrapolation can cause speculative bubbles to last longer. More severe extrapolation will result in larger excess demand once the return threshold is crossed, and this will result in an even larger excited-speculator population in

the following period. But, this increase doesn't affect the number of bubble episodes.

This is an important distinction to make. Policymakers care about what makes speculative bubbles more or less likely. For example, it's common to see articles discussing whether or not "China's stimulus program is prone to blow more bubbles in the economy next year."<sup>5</sup> The main concern in these articles is not the exact psychological bias or trading constraint that sustains the speculative bubble; nor is it the exact timing of the eventual market crash. The main concern is the likelihood of a speculative bubble occurring the first place; it's the likelihood of a displacement event. In cryptocurrencies?<sup>6</sup> In U.S. stocks?<sup>7</sup> Or, maybe in Chinese real estate?<sup>8</sup> These are questions that policymakers care about. But, they're also questions that are outside the scope of the existing limits-to-arbitrage framework.

### 3.2 Estimation Strategy

To bring the model to the data, we have to find some empirical proxy for  $\theta$ . Fortunately, the economic model above provides guidance about how to do this. If we can find a measure of how long the transient population of excited speculators remains interested in an asset during normal times, then the economic model tells us how to convert this measure into a proxy for that asset's  $\theta$ .

*Normal-Times Estimate.* Suppose that an infinitesimal population of speculators,  $n_0 > 0$ , happens to get excited about the risky asset at time  $\tau = 0$ . Let  $\tau_{1/2}$  denote the time it takes for half of these excited speculators to regain their senses:

$$\tau_{1/2} \stackrel{\text{def}}{=} \min \left\{ \tau > 0 : \frac{n_\tau}{n_0} \leq \frac{1}{2} \right\} \quad (26)$$

If  $r < r_*$ , then Proposition 2.1 dictates that  $\tau_{1/2} < \infty$ ; whereas, if  $r > r_*$ , then  $\tau_{1/2} = \infty$  since the excited-speculator population doesn't decay. In Appendix A, I show that

$$\tau_{1/2}(\theta, r) = \log(2) \cdot (1 - \theta \cdot r) \quad (27)$$

The idea is to learn about  $\theta$  by studying how small changes in past returns affect this half-life during normal times. Let  $\epsilon$  denote a small fluctuation in past returns,  $r \mapsto r_\epsilon \stackrel{\text{def}}{=} r \cdot (1 + \epsilon)$ . Think about  $\epsilon$  as a recent short-run innovation in an asset's cumulative long-run return,  $r$ . For example, in the econometric analysis below,  $r$  will correspond to an asset's cumulative return over the past two years while  $\epsilon$  will correspond to the same asset's monthly return. Let  $\Delta\tau_{1/2}$  denote how this small change in an asset's past returns affects the rate at which its excited speculators calm down:

$$\Delta\tau_{1/2}(\theta, r; \epsilon) \stackrel{\text{def}}{=} \tau_{1/2}(\theta, r_\epsilon) - \tau_{1/2}(\theta, r) \quad (28)$$

<sup>5</sup>*China Blowing Major Bubbles In 2017*. Forbes. Dec 19, 2016.

<sup>6</sup>"Bitcoin is heading to \$10,000" CNBC. Oct 20, 2017.

<sup>7</sup>"Goldman's Blankfein: Things Have Been Going Up for Too Long" Wall Street Journal. Sep 6, 2017.

<sup>8</sup>"Chinese Efforts to Stem Housing Bubble Show Promise" Bloomberg. Jun 11, 2017.

If a small increase in past returns leads to a large increase in the amount of time it takes for excited speculators to calm down, then the following quantity will be large:

$$C(\theta, r; \epsilon) \stackrel{\text{def}}{=} \Delta\tau_{1/2}(\theta, r; \epsilon) \times \epsilon \quad (29)$$

By contrast, if increasing past returns doesn't affect the rate at which excited speculators calm down, then  $C(\theta, r; \epsilon) = 0$ . The proposition below shows how this quantity is related to the sensitivity of speculator persuasiveness to an asset's past returns,  $\theta$ .

**Proposition 3.2** (*Normal-Times Estimate*). *The quantity  $C(\theta, r; \epsilon)$  as defined in Equation (29) is increasing in  $\theta$ :*

$$\frac{\partial}{\partial \theta} E[C(\theta, r; \epsilon)] > 0 \quad (30)$$

Furthermore, if the small fluctuations in an asset's past returns are drawn from a normal distribution,  $\epsilon \sim N(0, \omega^2)$  for  $\omega > 0$ , then:


$$E[C(\theta, r; \epsilon)] = \text{Cov}[\Delta\tau_{1/2}(\theta, r; \epsilon), \epsilon] \quad (31)$$

Put differently, assets where short-run changes in past returns have a larger impact on the persistence of any transient population of excited speculators have larger values of  $\theta$ .

If we had an empirical proxy for how long excited speculators remained interested in an asset during normal times, then Proposition 3.2 suggests how to use this proxy to estimate  $\theta$ .

Figure 8 depicts the intuition behind this approach. Both panels show the rate at which an initial population of excited speculators decays towards  $\bar{n} = 0$  when an asset's past return is below the critical threshold,  $r < r_*$ . The grey lines represent this transition when the asset's past return is precisely  $r$ ; whereas, the black lines represent this transition when the asset's past return is  $r_\epsilon = r \cdot (1 + \epsilon)$  for  $\epsilon > 0$  and  $r_\epsilon < r_*$ . The half-life of the excited-speculator population corresponds to the point on the  $x$ -axis at which the intersection with the horizontal dotted line at  $\frac{n_0}{2}$  occurs. Comparing Panel 8a to Panel 8b shows how increasing the sensitivity of an asset's speculators to its past returns also increases the sensitivity of the half-life of these speculators during normal times to fluctuations in the asset's past returns—i.e.,  $\Delta\tau_{1/2}(\theta, r; \epsilon)$  is larger when  $\theta = 0.4$  than when  $\theta = 0.2$ .

*What Determines  $\theta$ ?* This paper provides a model of how cross-sectional differences in  $\theta$  will affect asset prices, but it does not explain where these cross-sectional differences come from. This omission might seem problematic at first. But, in truth, it's a feature not a bug. While it would be interesting to understand the sources of variation in  $\theta$ , there's no reason to believe that these micro-foundations are constant across assets and over time. Was the thought process that made residential housing in Las Vegas, NV an exciting proposition to second-home buyers in the mid 2000s the same as the thought process that made dotcom stocks an exciting proposition to day traders in the late 1990s? Were those mid-2000s housing



**Figure 8. Normal-Times Estimate.** Transition of an initial population of  $n_0 > 0$  excited speculators to the steady-state value of  $\bar{n} = 0$  when  $r < r_*$ .  $x$ -axis: time  $\tau \geq 0$ . Grey line: size of excited-speculator population at time  $\tau$  given past return  $r$ ,  $n_\tau(\theta, r)$ . Black line: size of excited-speculator population at time  $\tau$  given past return  $r_\epsilon = r \cdot (1 + \epsilon)$  for  $\epsilon > 0$ ,  $n_\tau(\theta, r_\epsilon)$ . Half-life of excited-speculator population corresponds to the  $x$ -axis coordinate of the intersection with the horizontal dotted line at  $\frac{n_0}{2}$ .  $\Delta\tau_{1/2}$  represents effect of an  $\epsilon$ -fluctuation in past returns on this half-life and corresponds to the horizontal distance between the two vertical dotted lines.  $\theta = 0.4$ . Transition in a market where the risky asset has a parameter value of  $\theta = 0.4$  as in Figure 5.  $\theta = 0.2$ . Transition in an otherwise identical market where  $\theta = 0.2$ .

speculators in Las Vegas thinking along even remotely similar lines to land speculators in 1719 who bought shares in John Law’s Mississippi Company? A key benefit of the estimation strategy embodied by Proposition 3.2 is that you can estimate the sensitivity of speculators to past returns for each of these assets without knowing precisely how to answer these questions—without knowing whether speculators were engaged in “new era thinking” (Shiller, 2000)—without knowing whether “[each] time [was] different” (Reinhart and Rogoff, 2009).

More generally, economic models are useful precisely because they hide lots of unobserved nitty-gritty details about a problem inside a few estimable parameters. Consider the model of liquidity proposed by Kyle (1985). This model predicts that an informed trader’s price impact, which is known as Kyle’s  $\lambda$ , will be proportional to the ratio of an asset’s fundamental volatility to its demand-noise volatility. Thus, assets with more volatile fundamentals will be less liquid because market makers are at a larger informational disadvantage. By contrast, assets with more demand-noise volatility will be more liquid because this noisy order flow makes it easier for informed traders to hide their own demand. What makes this model useful is that you can estimate Kyle’s  $\lambda$  (Hasbrouck, 1991; Amihud, 2002) and build new theories based on this insight (Acharya and Pedersen, 2005; Nagel, 2005) without a general theory for why some assets have more volatile fundamentals or demand noise (Chinco and Fos, 2018). The whole goal in writing down an economic model is to separate the economic insight you want to convey from the hard-to-estimate details about the problem in question.

## 4 Econometric Analysis

This section uses matching analysis to verify the economic model’s main empirical implication. I start by looking for instances where an industry realized explosive price growth followed by an immediate collapse. I call these episodes “speculative bubbles”. Then, I find the point in time prior to each bubble episode where stock prices in that industry first started to rise. This could be two years before the peak. Or, it could be four years before the peak. The exact timing will vary across bubble episodes. And, that’s OK since we are interested in predicting the likelihood of a speculative bubble taking place not the timing of the peak. Finally, I match these pre-bubble industry $\times$ month observations to other similar industry $\times$ month observations that did not subsequently experience a speculative bubble based on observable characteristics such as past returns, return volatility, and price-to-earnings ratios. These matched observations always involve a different industry in a non-overlapping time period. Proposition 3.1 predicts that, holding all else constant, industries with higher  $\theta$ s should be more likely to experience a speculative bubble following good past performance. And, consistent with this prediction, I find that a 1%pt increase in  $\theta$  is associated with a 3.79%pt increase in the likelihood of an industry $\times$ month observation in this matched sample subsequently experiencing a bubble.

### 4.1 Data Description

I start by describing the data. To distinguish between theoretical objects and their empirical counterparts, I use **teletype font** to denote empirically observed variables. For example,  $\text{PastReturn}_{i,t}$  will denote the statistical estimate for the  $i$ th industry’s past return in month  $t$ , which corresponds to the model parameter  $r$ .

*Industry-Specific Data.* I study monthly data on industries defined using the [Fama and French \(1997\)](#) 49-industry classification scheme.<sup>9</sup> I remove the “other” industry because this does not represent a cohesive collection of stocks, and there’s no way for speculators to have

---

<sup>9</sup>I label industries using the abbreviations provided in [Fama and French \(1997\)](#). 1) aero: aircraft; 2) agric: agriculture; 3) autos: automobiles and trucks; 4) banks: banking; 5) beer: beer and liquor; 6) bldmt: construction materials; 7) books: printing and publishing; 8) boxes: shipping containers; 9) bussv: business services; 10) chems: chemicals; 11) chips: electronic equipment; 12) clths: apparel; 13) cnstr: construction; 14) coal: coal; 15) drugs: pharmaceutical products; 16) elceq: electrical equipment; 17) fabpr: fabricated products; 18) fin: trading; 19) food: food products; 20) fun: entertainment; 21) gold: precious metals; 22) guns: defense; 23) hardw: computers; 24) hlth: healthcare; 25) hshld: consumer goods; 26) insur: insurance; 27) labeq: measuring and control equipment; 28) mach: machinery; 29) meals: restaurants, hotels, and motels; 30) medeq: medical equipment; 31) mines: non-metallic and industrial metal mining; 32) oil: petroleum and natural gas; 33) paper: business supplies; 34) persv: personal services; 35) rlest: real estate; 36) rtail: retail; 37) rubbr: rubber and plastic products; 38) ships: shipbuilding and railroad equipment; 39) smoke: tobacco products; 40) soda: candy and soda; 41) softw: computer software; 42) steel: steel works; 43) telcm: communications; 44) toys: recreation; 45) trans: transportation; 46) txtls: textiles; 47) util: utilities; 48) whsl: wholesale.

an exciting conversation about a miscellaneous group of unrelated stocks. I use the value-weighted monthly returns for each industry reported on Ken French’s website. Let  $\text{Return}_{i,t}$  denote the return of the  $i$ th industry in month  $t$ .

I use three different kinds of control variables. The first group of control variables can be computed directly from an industry’s realized returns. Let  $\text{PastReturn}_{i,t}$  denote the cumulative return of the  $i$ th industry over the previous two years:

$$\text{PastReturn}_{i,t} \stackrel{\text{def}}{=} \prod_{t'=t-23}^t (1 + \text{Return}_{i,t'}) \quad (32)$$

This variable represents the amount of extra money that a speculator would have had in their pocket today if they had invested \$1 in the  $i$ th industry two years ago. If  $\text{PastReturn}_{i,t} = 0$ , then they would have lost their entire \$1 investment by today; whereas, if  $\text{PastReturn}_{i,t} = 2$ , then their initial \$1 investment would have doubled in value by today. Note that this definition of  $\text{PastReturn}_{i,t}$  also reflects what happens when speculators study price indexes normalized to one at a previous date (Case and Shiller, 1987). Let  $\text{ReturnVol}_{i,t}$  denote the annualized volatility of the  $i$ th industry’s monthly returns over the same time period.

The second group of control variables measures trading activity in each industry. These data come from the CRSP database. Let  $\text{Turnover}_{i,t}$  denote the two-year moving average of the ratio of annual trading volume to total shares outstanding. So, for example,  $\text{Turnover}_{i,t} = 2$  would imply that each share of the typical stock in industry  $i$  is traded twice each year. Stock-specific observations of share turnover are aggregated up to the industry level each month using a value-weighted average. Let  $\% \text{Issuer}_{i,t}$  denote the two-year moving average of the fraction of stocks in the  $i$ th industry that increased their number of shares outstanding by  $\geq 5\%$  during the previous year.

The third group of control variables characterizes industry-level fundamental values. I download these data from the WRDS Financial Ratios Suite. Let  $\text{CAPE}_{i,t}$  denote the cyclically-adjusted price-to-earnings ratio for the  $i$ th industry in month  $t$ . And, let  $\% \Delta \text{Sales}_{i,t}$  denote the two-year moving average of annual sales growth for the  $i$ th industry in month  $t$ . All industry-level measures of fundamental value are aggregated up from the stock-level each month by computing a value-weighted average.

In my main analysis, I use data for each of the 48 non-other industries in the Fama and French (1997) 49-industry classification scheme from February 1972 to December 2017. To give a sense of what these industry-level returns look like, Figure 9 shows the outcome of continuously re-investing \$1 in each industry starting in February 1972. The binding data constraint comes from the WRDS Financial Ratios Suite. If I only match pre-bubble industry $\times$ month observations to other non-bubble observations using past returns, I can extend my sample back to January 1950. And, I show in Table 4 that the results are similar when looking at this extended sample. The empirical results don’t disappear prior to 1972.



*Bubble Episodes.* As a working definition of what constitutes a speculative bubble, I look for instances where an industry realizes explosive price growth followed by an immediate crash. Let  $t_p$  denote a local maximum in the cumulative returns of the  $i$ th industry, which means that the  $i$ th industry’s cumulative returns in month  $t$  were higher than at any point during the previous or subsequent five years.

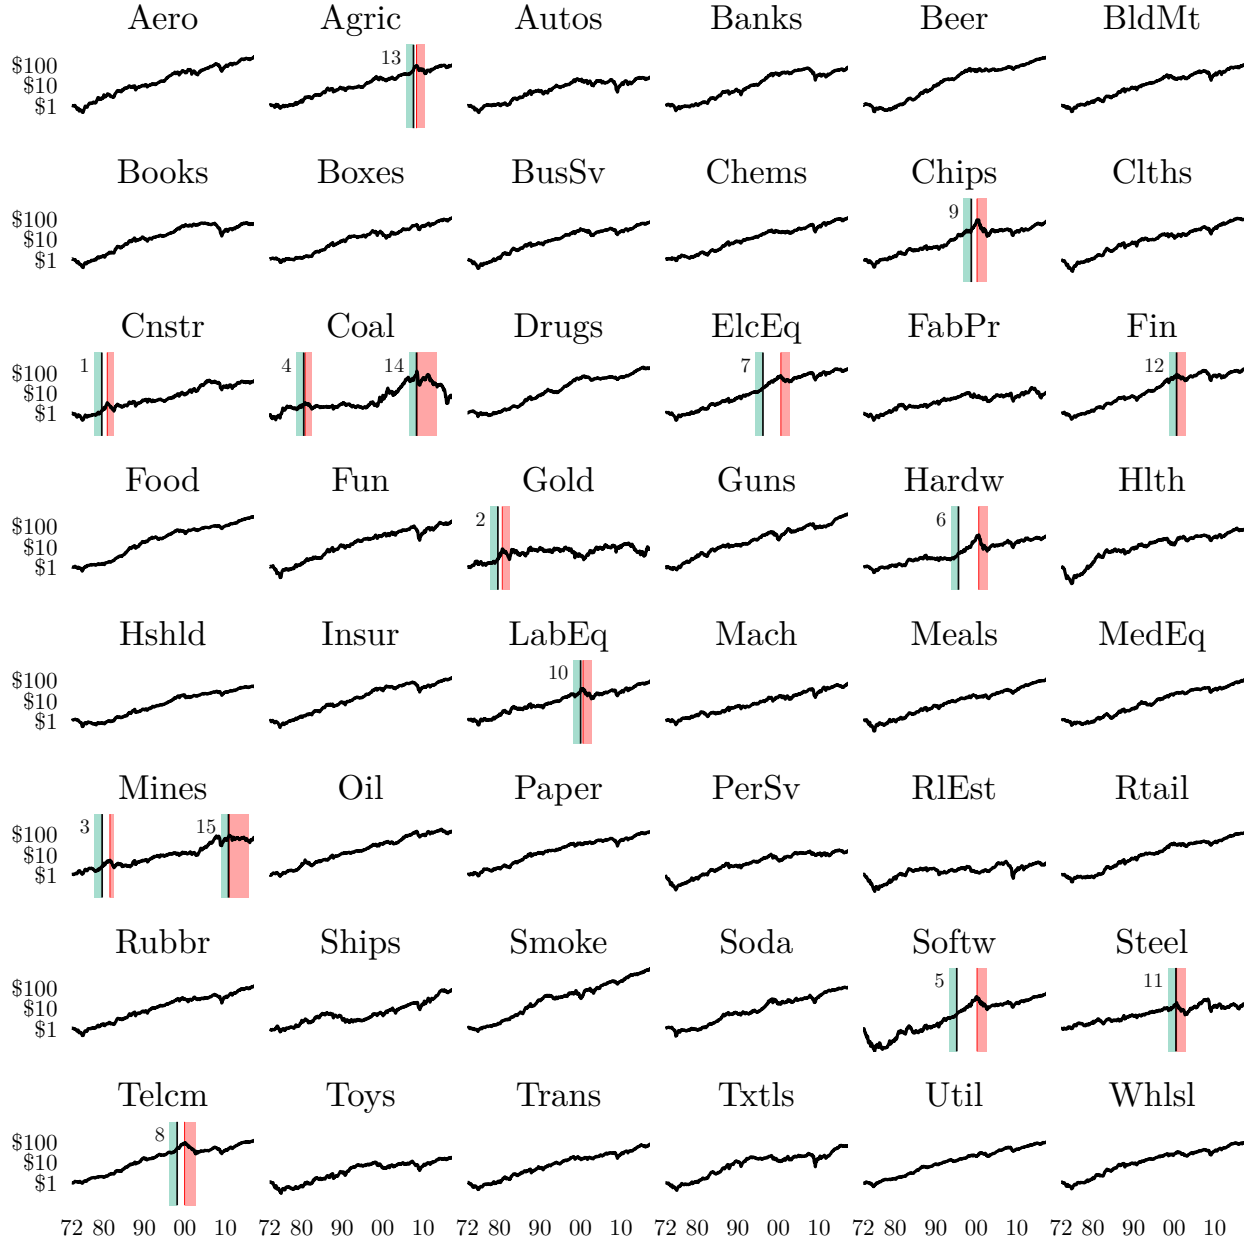
Motivated by the definition used in [Greenwood et al. \(2018\)](#), I say that this local maximum  $t_p$  represents the peak of a bubble episode if the  $i$ th industry’s past returns during the previous two- and five-year periods exceeded 100%, and the industry’s subsequent returns over the following two years were less than  $-50\%$ . See the top panel of Figure 10 for an illustration. I find 15 bubble episodes in my main sample period from February 1972 to December 2017 and 24 such episodes in the extended sample dating back to January 1950. I label each of these speculative bubbles in the plot of industry-level cumulative returns in Figure 9.

This empirical definition a bubble captures the main industry-level events that get called speculative bubbles, such as the rise and fall of technology stocks during the late 1990s and early 2000s. But, I want to emphasize that the empirical results will not hinge on this exact definition. As highlighted in Table 6, we can reason about the likelihood of an event even if there is sometimes disagreement about whether the event has taken place after the fact.


*Matching Analysis.* I next want to find the point in time leading up to each peak where an industry’s price first starts to rise and then match these pre-bubble observations to other similar industry  $\times$  month observations that did not subsequently experience a speculative bubble. Let the start date of a speculative bubble in the  $i$ th industry,  $t_s$ , be the month prior to the peak in which the industry’s two-year cumulative returns over the next two years first exceed 50%. In other words,  $t_s$  is the most recent month prior to the speculative bubble such that  $\text{PastReturn}_{i,t_s+23} < 50\%$  while  $\text{PastReturn}_{i,t_s+24} > 50\%$ . The closest the start date could ever be to the date of the peak is two years since, by definition, any speculative bubble must have two-year cumulative returns greater than 100% during the boom.

I then match these pre-bubble industry  $\times$  month observations to other similar observations based on the information available about the industry in month  $t_0 \stackrel{\text{def}}{=} t_s + 24$ . The idea is to pick an initial date where prices have started to rise but the limits-to-arbitrage dynamics haven’t kicked in yet. For each of the 15 industry  $\times$  month observations that precede a speculative bubble, I find a matched sample of five otherwise similar industry  $\times$  month observations that don’t precede a bubble episode.

Matches are chosen for each treated unit one at a time. At each matching step, I choose the not-yet-matched industry  $\times$  month observation with no subsequent bubble that is most similar to the pre-bubble observation according to **PastReturn**, **ReturnVol**, and **CAPE**. In other words, I match each speculative bubble to other industry  $\times$  month observations based on observable



**Figure 9. Bubble Episodes.**  $x$ -axis: time in months from February 1972 to December 2017.  $y$ -axis: dollar value of a continuously re-invested portfolio that started with \$1 in February 1972 on a logarithmic scale. Each panel represents results for an industry-specific portfolio, which is labeled using the abbreviations provided in [Fama and French \(1997\)](#). Thick black line: portfolio size in month  $t$ . Vertical red line: peak of bubble episode,  $t_p$ . Red shaded region: length of crash following the peak. Green shaded region: first two-year period prior to a bubble episode where an industry's cumulative returns exceeded 50%. Vertical black line: point in time where matching takes place,  $t_0$ . Numbers: total number of bubble episodes since February 1972.



**Figure 10. Bubble Definition and Match Timing.** The thick black line in both panels represents the cumulative returns to investing \$1 in the construction industry, Cnstr, from January 1976 to December 1984. These returns are depicted on a logarithmic scale. Figure 9 indicates that during this time interval the construction industry experienced a speculative bubble that peaked in  $t_p$  = November 1980. **Top Panel.** How this peak is defined. It's a local maximum in the construction industry's cumulative returns with  $> 100\%$  price appreciation during the preceding two years and  $< -50\%$  price depreciation immediately after. **Bottom Panel.** How the start of this speculative bubble is defined. The time interval from August 1977 to July 1979 was the first 24-month period prior to the peak to experience  $> 50\%$  returns. In the matching analysis, I look for five other industry  $\times$  month observations that had similar past returns, return volatility, and price to earnings as the construction industry in  $t_0$  = July 1979 but did not subsequently experience a speculative bubble.

characteristics during the shaded green time interval in the bottom panel of Figure 10—i.e., based on characteristics observable at time  $t_0$ . These matched industry  $\times$  month observations always involve a different industry in a non-overlapping time period. This matching procedure results in a dataset with  $90 = (1 + 5) \times 15$  observations. Let  $\text{WillBeBubble}_{i,t}$  denote an indicator variable that is one if and only if an observation of the  $i$ th industry in month  $t$  precedes a speculative bubble in this matched dataset. By construction, one out of every six observations in this matched dataset will precede a bubble episode. Thus, we have that one out of every six observations in the matched dataset is followed by a bubble episode,  $\Pr[\text{WillBeBubble}_{i,t} = \text{True}] = 1/6 \approx 16.67\%$ .

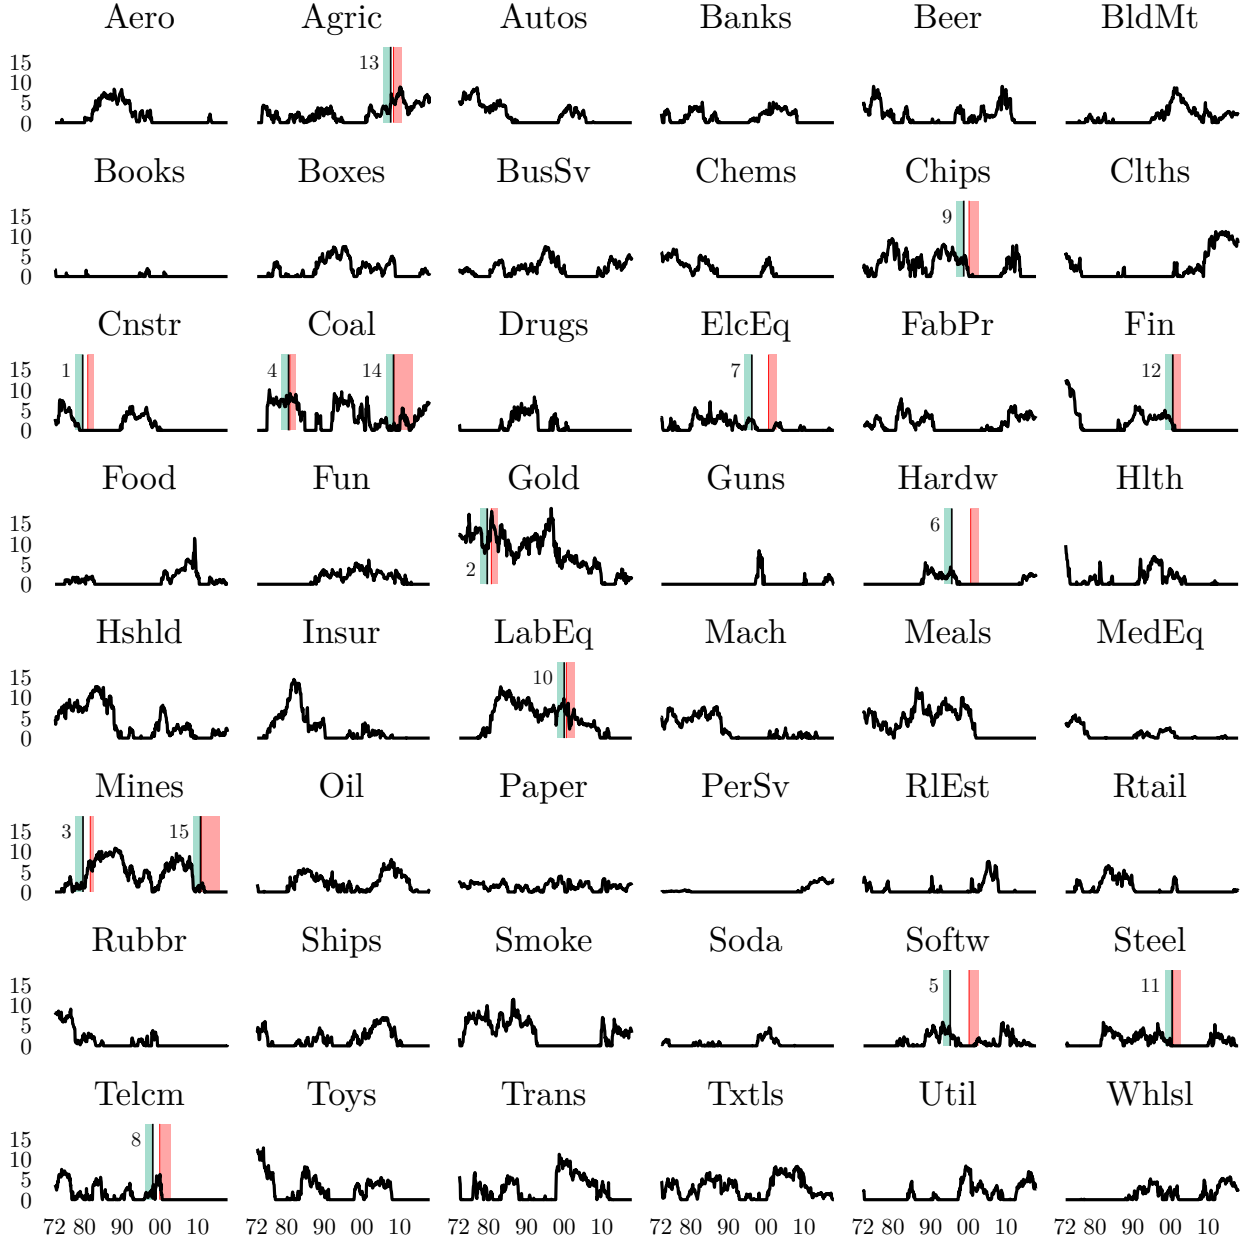
*Estimated Theta.* Our aim is to see if the pre-bubble observations in this matched dataset tend to have higher values of  $\theta$  than their matched counterparts. Unfortunately, we cannot directly observe each industry's value of  $\theta$ . So, we will have to estimate this quantity. Propo-

sition 3.2 details how to convert a measure of how long the transient population of excited speculators remains interested in an industry during normal times into a proxy for that industry’s  $\theta$ . I proxy for how long the transient population of excited speculators remains interested in an industry during normal times with a measure of industry-specific media coverage. The motivation by this approach is that news outlets strategically choose which industries to cover so as to maximize total readership (Mullainathan and Shleifer, 2005). Thus, if an industry starts to get more coverage while  $r < r_*$ , then it’s likely that the transient population of excited speculators is remaining in the market longer.

To compute this measure of industry-specific media coverage, I search the ProQuest historical Wall Street Journal archives for articles that include industry-specific keywords from January 1940 to December 2017. The approach is closely related to several recent applications of textual analysis in finance and economics (Da et al., 2014; Baker et al., 2016; Hansen et al., 2017; Manela and Moreira, 2017; Gentzkow et al., 2017). Let  $\#Articles_{i,t}$  denote the number of articles in month  $t$  that reference the  $i$ th industry’s.

Then, to calculate my empirical proxy for  $\theta$ ,  $\Theta_{i,t}$ , I estimate the covariance between an industry’s media coverage in a given month and its return in the previous month using data covering the previous ten years. This means regressing the log number of Wall Street Journal articles that mention the  $i$ th industry in month  $t$  on the  $i$ th industry’s returns in month  $t - 1$  using data from the 120 months  $t \in \{t - 119, \dots, t\}$ . I estimate separate regressions for each industry in each month. To make sure that the estimate for each industry’s speculator sensitivity only uses data from normal times, I do not compute  $\Theta_{i,t}$  in any month  $t$  where an industry was experiencing a speculative bubble at any point during the estimation window for  $\Theta_{i,t}$ . Figure 11 displays these rolling industry-specific estimates for  $\theta$  during the main sample period from February 1972 to December 2017.

*Summary Statistics.* Tables 1 and 2 provide descriptive statistics about variable used in the empirical analysis. They describe what these variables look like both in general and in the matched sample. The top panel of Table 1 reports summary statistics for each variable using all available industry $\times$ month observations; whereas, the bottom panel of Table 1 reports summary statistics for only those observations used in the matching analysis. This matching is done based on the variables **PastReturn**, **ReturnVol**, and **CAPE**. So, these variables look the same in both the bubble episodes and the matched sample. However, the right-most column shows that meaningful correlations remain even after matching on these three variables. The 15 industry $\times$ month observations with subsequent bubbles have higher values of  $\theta$  and more sales growth than the matched industry $\times$ month observations with no subsequent bubbles. Table 2 provides sample correlations for the variables computed across all available industry $\times$ month observations.



**Figure 11. Estimated Theta.**  $x$ -axis: time in months from February 1972 to December 2017.  $y$ -axis: estimated value for  $\theta$ ,  $\text{Theta}_{i,t}$ , for the  $i$ th industry in month  $t$  in units of percent. Each panel represents results for a single industry, which is labeled using the abbreviations provided in [Fama and French \(1997\)](#). Vertical red line: peak of bubble episode,  $t_p$ . Red shaded region: length of crash following the peak. Green shaded region: first two-year period prior to a bubble episode where an industry's cumulative returns exceeded 50%. Vertical black line: point in time where matching takes place,  $t_0$ . Numbers: total number of bubble episodes since February 1972.

(a) Full Sample					
	Avg	Sd	10%	50%	90%
Theta	1.76	2.56	0.00	0.27	5.43
log(#Articles)	1.50	0.97	0.23	1.45	2.90
%Δ#Articles	0.37	8.76	−9.65	0.18	11.22
PastReturn	26.64	37.09	−17.08	24.91	69.81
ReturnVol	21.24	8.07	12.46	20.08	31.44
CAPE	22.19	12.87	9.74	20.01	38.35
%ΔSales	5.47	5.49	0.13	5.28	10.97
Turnover	1.04	0.92	0.24	0.74	2.31
%Issuer	0.24	0.16	0.05	0.21	0.45

(b) Matched Data					
	Bubble Episodes Avg	Sd	Matched Sample Avg	Sd	Diff
Theta	3.29	2.85	1.70	2.26	1.59*** (0.67)
log(#Articles)	1.53	0.81	1.45	0.98	0.09 (0.27)
%Δ#Articles	2.41	8.24	0.46	7.75	1.95 (2.22)
PastReturn	71.28	31.77	61.30	55.90	9.98 (14.93)
ReturnVol	26.67	10.45	26.26	11.38	0.41 (3.18)
CAPE	37.15	30.02	38.63	20.90	−1.48 (6.39)
%ΔSales	12.02	3.53	8.01	7.90	4.01* (2.09)
Turnover	1.80	1.60	1.40	1.09	0.41 (0.34)
%Issuer	0.32	0.16	0.30	0.20	0.02 (0.06)
#Obs	15		75		

**Table 1. Summary Statistics. Top Panel.** Summary statistics for all industry  $\times$  month observations from February 1972 to December 2017. **Bottom Panel.** Summary statistics for 15 pre-bubble industry  $\times$  month observations and matched sample that did not subsequently experience a speculative bubble. This matching is done based on the variables `PastReturn`, `ReturnVol`, and `CAPE`. Diff: difference between the average value in bubble episodes vs. in matched sample. Numbers in parentheses are standard errors. **Theta**: estimated speculator sensitivity parameter in percent. **log(#Articles)**: log number of WSJ articles per month. **%Δ#Articles**: growth in number of WSJ articles per month over past two years in percent. **PastReturn**: return over past two years in percent. **ReturnVol**: annualized standard deviation of monthly returns in percent. **CAPE**: cyclically adjusted price-to-earnings ratio. **%ΔSales**: annualized sales growth over past two years in percent. **Turnover**: annual trading volume divided by shares outstanding. **%Issuer**: percent of firms that issued  $\geq 5\%$  new shares during past year. Significance: \* = 10%, \*\* = 5%, and \*\*\* = 1%.

---

	<i>Theta</i>	<i>log(#Articles)</i>	<i>%Δ#Articles</i>	<i>PastReturn</i>	<i>ReturnVol</i>	<i>CAPE</i>	<i>%ΔSales</i>	<i>Turnover</i>	<i>%Issuer</i>
Theta	1.00								
log(#Articles)	−0.02	1.00							
%Δ#Articles	0.05	0.11	1.00						
PastReturn	0.00	−0.02	0.00	1.00					
ReturnVol	0.09	−0.02	0.13	−0.20	1.00				
CAPE	−0.02	−0.15	0.02	0.20	−0.09	1.00			
%ΔSales	0.04	−0.03	0.10	0.25	−0.03	0.25	1.00		
Turnover	−0.10	0.00	−0.03	−0.06	0.30	0.18	−0.11	1.00	
%Issuer	0.04	0.09	0.14	0.07	0.18	0.15	0.32	−0.07	1.00

---

**Table 2. Sample Correlations.** Sample correlations computed across all industry  $\times$  month observations from February 1972 to December 2017. **Theta**: estimated speculator sensitivity parameter in percent. **log(#Articles)**: log number of WSJ articles per month. **%Δ#Articles**: growth in number of WSJ articles per month over past two years in percent. **PastReturn**: return over past two years in percent. **ReturnVol**: annualized standard deviation of monthly returns during past two years in percent. **CAPE**: cyclically adjusted price-to-earnings ratio. **%ΔSales**: annualized sales growth over past two years in percent. **Turnover**: annual trading volume divided by shares outstanding. **%Issuer**: percent of firms that issued  $\geq 5\%$  new shares during past year.



## 4.2 Empirical Evidence

Proposition 3.1 predicts that, holding all else constant, industries with higher  $\theta$ s should be more likely to experience speculative bubbles following good past performance. And, consistent with this prediction, I find that a 1%pt increase in an industry’s estimated  $\mathbf{Theta}_{i,t}$  is associated with a 3.79%pt increase in the likelihood that an industry  $\times$  month observation in the matched dataset is one of the 15 pre-bubble observations rather than one of the 75 otherwise similar observations that don’t precede a speculative bubble.

*Bubble Likelihood.* I use the following linear regression specification in the main analysis:

$$\mathbf{WillBeBubble}_{i,t} = \hat{a} + \hat{b} \cdot \mathbf{Theta}_{i,t} + \dots + \mathbf{e}_{i,t} \quad (33)$$

On the left-hand side is an indicator variable,  $\mathbf{WillBeBubble}_{i,t}$ , which signals whether a particular industry  $\times$  month observation in the matched dataset actually precedes a speculative bubble. If  $\mathbf{WillBeBubble}_{i,t} = \text{True}$ , then  $(i, t)$  corresponds to an industry  $\times$  month observation during the run-up to one of the 15 bubble episodes highlighted in Figure 9. By contrast, if  $\mathbf{WillBeBubble}_{i,t} = \text{False}$ , then  $(i, t)$  represents an observation that did not precede a speculative bubble. The key right-hand-side variable is the estimated sensitivity parameter,  $\mathbf{Theta}_{i,t}$ . The “...” term stands for any additional variables that might be added to the regression, and  $\mathbf{e}_{i,t}$  denotes the regression residual. In unreported results, I find that using a Logit or Probit model to account for the categorical nature of the dependent variable delivers qualitatively similar marginal effects.

Table 3 reports the results of this baseline regression specification. Column (1) shows that, all else equal, industry  $\times$  month observations that subsequently lead to speculative bubbles have higher values of  $\theta$ . Columns (2)-(6) then show that this positive relationship is not explained by differences in sales growth, share turnover, or new issuance. The estimate of 3.79%pt in Column (1) represents the change in the bubble likelihood associated with increasing an industry’s estimated  $\mathbf{Theta}$  by 1%pt from its average value.

Here’s one way of thinking about how large an effect this is. Imagine I gave you the six industry  $\times$  month observations associated with a particular bubble episode—i.e., the one pre-bubble observation and its five most similar matches where no speculative bubble later took place. But, suppose that I didn’t tell you which was the pre-bubble observation. If you just picked an industry  $\times$  month observation at random, then the probability that you would pick the pre-bubble observation would be  $1/6 \approx 16.67\%$ . By contrast, if you selected an industry  $\times$  month observation where  $\mathbf{Theta}$  was 1%pt above average, then your probability of picking the pre-bubble observation would rise to  $16.67\% + 3.79\% = 20.46\%$ . It would be as if I had eliminated one of the matched industry  $\times$  month observations from your choice set, increasing your odds of choosing correctly from one in six to better than one in five.

Dependent Variable: WillBeBubble					
	(1)	(2)	(3)	(4)	(5)
Intercept	9.22* (4.97)	-0.39 (6.68)	0.78 (7.18)	5.34 (8.03)	-5.51 (9.19)
Theta	3.79** (1.60)	3.96** (1.57)	4.17** (1.60)	3.88** (1.61)	4.25*** (1.58)
% $\Delta$ Sales		1.07** (0.51)			1.08* (0.55)
Turnover			5.27 (3.26)		4.80 (3.27)
%Issuer				0.12 (0.20)	-0.08 (0.21)
Matched on...					
PastReturn	✓	✓	✓	✓	✓
ReturnVol	✓	✓	✓	✓	✓
CAPE	✓	✓	✓	✓	✓

**Table 3. Bubble Likelihood.** Each column reports the results of a separate regression estimated using data on a sample of industry $\times$ month observations matched on **PastReturn**, **ReturnVol**, and **CAPE**. Sample period covers February 1972 to December 2017 and contains 15 speculative bubbles. Matched dataset contains five similar industry $\times$ month observations per speculative bubble and 90 observations in total. **WillBeBubble**: dependent variable is an indicator variable that is one if an observations is one of the 15 industry $\times$ month observations that precedes a speculative bubble. **Theta**: estimated speculator sensitivity parameter in percent. **PastReturn**: return over past two years in percent. **ReturnVol**: annualized standard deviation of monthly returns in percent. **CAPE**: cyclically adjusted price-to-earnings ratio. **% $\Delta$ Sales**: annualized sales growth over past two years in percent. **Turnover**: annual trading volume divided by total shares outstanding. **%Issuer**: percent of firms that issued  $\geq 5\%$  new shares during past year. All coefficients represent percent point changes in the likelihood of a speculative bubble. Numbers in parentheses are standard errors. Significance: \* = 10%, \*\* = 5%, and \*\*\* = 1%.

*Sub-Sample Analysis.* The WRDS financial-ratios suite only reports data on cyclically-adjusted price-to-earnings ratios for each industry since January 1970. And, we are using CAPE as our measure of industry fundamentals when finding matches for our 15 pre-bubble industry $\times$ month observations. This is why the sample period used in the main analysis starts in February 1972—the extra 24 months are needed to compute past returns and return volatility for each industry. But, if we matched each pre-bubble industry $\times$ month observation to otherwise similar industry $\times$ month observations using only on past returns, then we could extend our sample period all the way back to January 1950.

Table 4 shows the results when we adopt this relaxed matching procedure. Column (1) reports the point estimate for the coefficient on the estimated value of  $\Theta_{i,t}$  using the original sample period from February 1972 to December 2017 when matching only on `PastReturn` rather than on `PastReturn`, `ReturnVol`, and `CAPE`. Columns (2)-(4) then extend this analysis to cover the entire sample period since January 1950, which contains an additional 9 speculative bubbles. Column (2) reports results for the full sample while columns (3) and (4) report results for the first and second halves of this extended sample period. Although the power shrinks when we cut our sample period in half, the point estimate for the coefficient on  $\Theta$  is positive and statistically significant for all regression specifications. This suggests the results in Table 3 are not specific to the post-1972 time period.

*Total Media Coverage.* Table 3 shows that industry $\times$ month observations in the matched dataset with higher  $\Theta$  values are more likely to later experience bubble episodes. But, maybe there is some other omitted variable driving the results? It’s hard to rule this out entirely, but there is one glaring possibility that needs to be considered. I am using a measure of media coverage to estimate  $\Theta$ . Perhaps it’s the total amount of media coverage or the growth in media coverage that’s actually driving the results in Table 3? It seems plausible that industries with more media coverage—i.e., industries that get mentioned more often in WSJ articles—might be more likely to experience bubble episodes.

The results in Table 5 show that this alternative hypothesis is not supported by the data. Columns (1)-(3) replicate the analysis in Table 3 using the level of media coverage, `log(#Articles)`, and the growth in media coverage over the past two years, `% $\Delta$ #Articles`, rather than  $\Theta$  as the key right-hand-side variable. They reveal that there’s only a weak relationship between total media coverage and the likelihood of a speculative bubble. Columns (4) and (5) of Table 5 then re-introduce  $\Theta$  into the regression specification and confirm that the point estimate for  $\Theta$  is unchanged when controlling for the level and growth of total media coverage. In short, the main results in this section are being driven by cross-sectional differences in  $\Theta$  rather than cross-sectional differences in total media coverage.

*Ex Post Disagreement.* There isn’t general consensus about what constitutes a speculative

Dependent Variable: WillBeBubble				
	1972-2017	1950-2017	1950-1983	1984-2017
	(1)	(2)	(3)	(4)
Intercept	9.81** (4.76)	11.30*** (3.78)	11.85** (4.91)	11.33** (5.46)
Theta	3.70** (1.51)	2.70** (1.11)	2.47* (1.32)	3.11* (1.77)
Matched on...				
PastReturn	✓	✓	✓	✓

**Table 4. Sub-Sample Analysis.** Each column reports the results of a separate regression estimated using data on a sample of industry×month observations matched only on **PastReturn**. Full sample period now covers January 1950 to December 2017 and contains 24 speculative bubbles. Each column reports results for a different subset of this full sample period, but the matched data contains five similar industry×month observations per speculative bubble. **WillBeBubble**: dependent variable is an indicator variable that is one if an observations is one of the industry×month observations that precedes a speculative bubble. **Theta**: speculator sensitivity parameter in percent. **PastReturn**: return over past two years in percent. All coefficients represent percent point changes in the likelihood of a speculative bubble. Numbers in parentheses are standard errors. Significance: \* = 10%, \*\* = 5%, and \*\*\* = 1%.

bubble. So, perhaps you looked at Figure 9 and thought to yourself: “Some of these other events seem like speculative bubbles, but episode  $x$  certainly wasn’t. I have a perfectly good fully rational explanation for why prices boomed and then crashed in that industry at that time.” Table 6 shows that the previous results are not affected by this sort of ex post disagreement. In other words, the empirical results in Table 3 don’t hinge on whether or not any particular market episode was or wasn’t a speculative bubble.

Each row in the table summarizes the results of re-estimating the specification in Equation (33) after omitting some combination of  $k \in \{0, 1, 2, 3, 4\}$  episodes from the 15 total speculative bubbles from February 1972 to December 2017. Think about each of these omitted episodes as a market event where you and I disagree. I think it should be called a speculative bubble. You don’t. After you choose some combination of  $k$  speculative bubbles to omit, I then match the remaining  $(15 - k)$  pre-bubble industry×month observations on **PastReturn**, **ReturnVol**, and **CAPE** just as in Table 3. The resulting matched dataset will then contain 90 observations when  $k = 0$  episodes are omitted, 84 observations when  $k = 1$  episode is omitted, 78 observations when  $k = 2$  episodes are omitted, etc. . .

The first row in Table 6 reports results when no episodes are dropped. So, the point estimate corresponds to the output in Column (1) of Table 3. Moving down the rows of the table corresponds to omitting more and more bubble episodes from the analysis. It’s clear

Dependent Variable: WillBeBubble					
	(1)	(2)	(3)	(4)	(5)
Intercept	14.68** (7.31)	16.31*** (3.98)	14.81** (7.32)	8.34 (7.70)	-5.25 (11.21)
Theta				3.68** (1.62)	4.20** (1.62)
log(#Articles)	1.36 (4.20)		1.04 (4.22)	0.57 (4.13)	0.10 (4.10)
%Δ#Articles		0.45 (0.51)	0.44 (0.51)	0.34 (0.50)	0.14 (0.50)
%ΔSales					1.06* (0.57)
Turnover					4.71 (3.33)
%Issuer					-0.09 (0.22)
Matched on...					
PastReturn	✓	✓	✓	✓	✓
ReturnVol	✓	✓	✓	✓	✓
CAPE	✓	✓	✓	✓	✓

**Table 5. Total Media Coverage.** Each column reports the results of a separate regression estimated using data on a sample of industry×month observations matched on **PastReturn**, **ReturnVol**, and **CAPE**. Sample period covers February 1972 to December 2017 and contains 15 speculative bubbles. Matched dataset contains five similar industry×month observations per speculative bubble and 90 observations in total. **WillBeBubble**: dependent variable is an indicator variable that is one if an observations is one of the 15 industry×month observations that precedes a speculative bubble. **log(#Articles)**: log of the number of WSJ articles per month referencing an industry. **%Δ#Articles**: growth rate of number of WSJ articles per month referencing an industry during past two years. **Theta**: speculator sensitivity parameter in percent. **PastReturn**: return over past two years in percent. **ReturnVol**: annualized standard deviation of monthly returns in percent. **CAPE**: cyclically adjusted price-to-earnings ratio. **%ΔSales**: annualized sales growth over past two years in percent. **Turnover**: ratio of annual trading volume to shares outstanding. **%Issuer**: percent of firms that issued  $\geq 5\%$  new shares during past year. All coefficients represent percent point changes in the likelihood of a speculative bubble. Numbers in parentheses are standard errors. Significance: \* = 10%, \*\* = 5%, and \*\*\* = 1%.

#Omitted	#Obs	$\binom{15}{k}$	Pr[ <i>signif</i> ]	Avg[ $\hat{b}$ ]	Sd[ $\hat{b}$ ]	Pr[ $\hat{b} \leq 0$ ]
$k = 0$	90	1	100.0%	3.79	.	0.0%
1	84	15	60.0%	3.71	1.07	0.0%
2	78	105	62.9%	3.87	1.18	0.0%
3	72	455	58.5%	3.92	1.43	0.0%
4	66	1365	53.6%	4.01	1.57	0.2%

**Table 6. Ex Post Disagreement.** Each row summarizes the results of re-estimating the specification in Equation (33) after omitting  $k \in \{0, 1, 2, 3, 4\}$  of the 15 speculative bubbles used in Table 3 from the analysis. After choosing which  $k$  speculative bubbles to omit, the  $(15 - k)$  remaining industry $\times$ month observations are then matched to five other observations with no subsequent bubble based on **PastReturn**, **ReturnVol**, and **CAPE** just as in Table 3. The resulting matched dataset contains 90 observations when  $k = 0$  episodes are omitted, 84 observations when  $k = 1$  episode is omitted, 78 observations when  $k = 2$  episodes are omitted, and so on... The first row is the same specification as in Column (1) of Table 3. #Omitted: number of speculative bubbles omitted from the analysis. #Obs: number of industry $\times$ month observations in the matched dataset when omitting  $k$  speculative bubbles.  $\binom{15}{k}$ : number of ways to omit  $k$  episodes from among the 15 speculative bubbles in the sample period from February 1972 to December 2017. Pr[*signif*]: percent of the matched datasets that omit  $k$  episodes which yield statistically significant point estimates for the slope coefficient at the 5% level. Avg[ $\hat{b}$ ]: average value of the slope coefficient across all matched datasets omitting  $k$  episodes. Sd[ $\hat{b}$ ]: standard deviation of the slope coefficient computed across all matched datasets omitting  $k$  episodes. Pr[ $\hat{b} \leq 0$ ]: percent of all matched datasets omitting  $k$  episodes that yield a negative slope coefficient.

from the column labeled Pr[*signif*] that dropping bubble episodes from the sample reduces the power of the test. It becomes less and less likely that the slope coefficient on **Theta** will be statistically significant at the 5% level. This is mechanical. However, the column labeled Sd[ $\hat{b}$ ] reveals that, while the statistical significance does drop, the point estimates don't fluctuate wildly in response to dropping particular episodes. In fact, the final column shows that you would have to drop more than three different episodes to produce a point estimate that's negative. The findings in Table 6 underscore the fact that we don't need complete after-the-fact agreement about which events were in fact speculative bubbles to draw conclusions about what makes these kinds of events more likely.

*Predicting The Crash.* Last but not least, the key insight in this paper is that predicting the likelihood of a speculative bubble occurring is conceptually distinct from predicting the timing of the eventual crash. Different economic forces are at work in each case. And, Table 7 shows that this idea bears itself out in the data.

Dependent Variable: SurvTime							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Theta	0.00 (0.13)			−0.05 (0.16)	−0.04 (0.16)	−0.07 (0.14)	−0.13 (0.15)
log(#Articles)		0.85* (0.45)					
%Δ#Articles			0.05 (0.04)				
PastReturn				0.01 (0.01)	0.01 (0.01)	0.05 (0.02)	0.02 (0.02)
ReturnVol				0.13*** (0.04)	0.14*** (0.05)	0.29*** (0.09)	0.17*** (0.06)
CAPE				0.02 (0.02)	0.02 (0.02)	0.06*** (0.02)	0.04** (0.02)
%ΔSales					0.02 (0.10)		
Turnover						−1.28** (0.53)	
%Issuer							−0.05* (0.03)

**Table 7. Predicting The Crash.** Each column reports the results of a separate Cox proportional-hazard model using data on the 15 bubble episodes from February 1972 to December 2017. **SurvTime**: dependent variable is the number of months  $t_0$  to a bubble’s peak,  $t_p$ . **Theta**: speculator sensitivity parameter in percent. **log(#Articles)**: log of the number of WSJ articles per month referencing an industry. **%Δ#Articles**: growth rate of number of WSJ articles per month referencing an industry during past two years. **PastReturn**: return over past two years in percent. **ReturnVol**: annualized standard deviation of monthly returns in percent. **CAPE**: cyclically adjusted price-to-earnings ratio. **%ΔSales**: annualized sales growth over past two years in percent. **Turnover**: ratio of annual trading volume to shares outstanding. **%Issuer**: percent of firms that issued  $\geq 5\%$  new shares during past year. Coefficients represent log hazard ratios. Positive coefficients are associated with shorter survival times—i.e., fewer months between  $t_0$  and  $t_p$ . Numbers in parentheses are standard errors. Significance: \* = 10%, \*\* = 5%, and \*\*\* = 1%.

Each column reports the coefficient estimates of a separate Cox proportional hazard model using data on the 15 pre-bubble industry×month observations since February 1972. All right-hand-side variables are sampled at date  $t_0$  prior to the peak. The dependent variable is the number of months between  $t_0$  and the peak of the bubble,  $\text{SurvTime}_{i,t} \stackrel{\text{def}}{=} t_p - t_0$ . A positive coefficient estimate in Table 7 implies that speculative bubbles with higher values reached their peak sooner—i.e., that the time gap between  $t_p$  and  $t_0$  was smaller.

The table confirms that, while variables such as share turnover and new issuance do



predict the timing of the crash as reported in [Greenwood et al. \(2018\)](#), the sensitivity of speculator persuasiveness to past returns does not. The coefficient on **Theta** is always small and statistically insignificant, which is consistent with the idea that different economic forces determine the likelihood of a bubble and the timing of the crash. And, this isn't just because news coverage is unrelated to the timing of the crash. High overall levels of media coverage tend to be associated with earlier crashes. It's only that speculator-persuasiveness sensitivity is unrelated to crash timing.

## 5 Conclusion

By the logic of limits to arbitrage, any bias-constraint pair might potentially combine at any moment to produce an equilibrium pricing error. And, over the past several decades, researchers have cataloged a wide variety of psychological biases and arbitrageur constraints at work in the market. But, when you look at the data, large pricing errors such as speculative bubbles are relatively rare. Why is this? What pins down the likelihood of a speculative bubble? How often should we expect to find some bias causing some constraint to bind?

The limits of arbitrage explain how an equilibrium pricing error can be sustained in equilibrium. But, these aren't questions about *how* a speculative bubble can be sustained; they're questions about *how often* we should expect one to occur. So, to answer them, you need to introduce another ingredient—an on/off switch—something that sporadically amplifies the effect of speculators' omnipresent biases, causing arbitrageurs' constraints to bind and a speculative bubble to form. This special something is typically called a “displacement event” ([Minsky, 1992](#)). And, this paper proposes a theory of displacement events that recombines two common elements found in popular accounts of bubble formation in a new way: #1) while speculators get overly excited following good news about fundamentals due to the “madness of crowds”, #2) they only recover their senses “slowly and one by one” ([Mackay, 1841](#)).

The model predicts that speculative bubbles will be more common in assets where speculator persuasiveness is more sensitive to fluctuations in past returns. Moreover, it should be possible to estimate this key sensitivity parameter for each asset using data collected during normal times—i.e., when there's no speculative bubble currently taking place. I verify this prediction empirically: after similar price run-ups, a 1%pt increase in  $\theta$  is associated with a 3.79%pt increase in the likelihood of a subsequent speculative bubble. This relationship does not appear to be explained by an industry's level of media coverage or other variables that predict the likelihood of an immediate crash, such as sales growth and share turnover ([Greenwood et al., 2018](#)). In short, these results give empirical evidence that predicting the likelihood of a speculative bubble requires a different kind of model than the one needed to predict the timing the eventual crash.

## References

- Abreu, D. and M. Brunnermeier (2003). Bubbles and crashes. *Econometrica*.
- Acharya, V. and L. Pedersen (2005). Asset pricing with liquidity risk. *Journal of Financial Economics*.
- Admati, A. (1985). A noisy rational-expectations equilibrium for multi-asset securities markets. *Econometrica*.
- Ahern, K. (2017). Information networks: evidence from illegal insider-trading tips. *Journal of Financial Economics*.
- Amihud, Y. (2002). Illiquidity and stock returns: cross-section and time-series effects. *Journal of Financial Markets*.
- Andrei, D. and J. Cujean (2017). Information percolation, momentum, and reversal. *Journal of Financial Economics*.
- Arnol'd, V. (2012). *Geometrical methods in the theory of ordinary differential equations*. Springer Science & Business Media.
- Bailey, M., R. Cao, T. Kuchler, and J. Stroebel (2018). The economic effects of social networks: evidence from the housing market. *Journal of Political Economy*.
- Baker, M. and J. Wurgler (2006). Investor sentiment and the cross-section of stock returns. *Journal of Finance*.
- Baker, S., N. Bloom, and S. Davis (2016). Measuring economic policy uncertainty. *Quarterly Journal of Economics*.
- Barberis, N., R. Greenwood, L. Jin, and A. Shleifer (2015). X-CAPM: an extrapolative capital asset-pricing model. *Journal of Financial Economics*.
- Barberis, N., R. Greenwood, L. Jin, and A. Shleifer (2018). Extrapolation and bubbles. *Journal of Financial Economics*.
- Barberis, N. and A. Shleifer (2003). Style investing. *Journal of Financial Economics*.
- Bordalo, P., N. Gennaioli, and A. Shleifer (2018). Diagnostic expectations and credit cycles. *Journal of Finance*.
- Brown, J., Z. Ivković, P. Smith, and S. Weisbenner (2008). Neighbors matter: causal community effects and stock-market participation. *Journal of Finance*.
- Brunnermeier, M. and S. Nagel (2004). Hedge funds and the technology bubble. *Journal of Finance*.
- Burnside, C., M. Eichenbaum, and S. Rebelo (2016). Understanding booms and busts in housing markets. *Journal of Political Economy*.

- Bursztyn, L., F. Ederer, B. Ferman, and N. Yuchtman (2014). Understanding the mechanisms underlying peer effects: evidence from a field experiment on financial decisions. *Econometrica*.
- Case, K. and R. Shiller (1987). Prices of single-family homes since 1970: new indexes for four cities. *New England Economic Review*.
- Cassidy, J. (2010). Interview with Eugene Fama. *The New Yorker*.
- Chinco, A. and V. Fos (2018). The sound of many funds rebalancing. SSRN.
- Chinco, A. and C. Mayer (2015). Misinformed speculators and mispricing in the housing market. *Review of Financial Studies*.
- Cutler, D., J. Poterba, and L. Summers (1990). Speculative dynamics and the role of feedback traders. *American Economic Review*.
- Da, Z., J. Engelberg, and P. Gao (2014). The sum of all FEARS: investor sentiment and asset prices. *Review of Financial Studies*.
- Daniel, K., D. Hirshleifer, and A. Subrahmanyam (1998). Investor psychology and security-market under and overreactions. *Journal of Finance*.
- De Long, B., A. Shleifer, L. Summers, and R. Waldmann (1990). Positive-feedback investment strategies and destabilizing rational speculation. *Journal of Finance*.
- Dufló, E. and E. Saez (2002). Participation and investment decisions in a retirement plan: the influence of colleagues' choices. *Journal of Public Economics*.
- Engelberg, J. and C. Parsons (2011). The causal impact of media in financial markets. *Journal of Finance*.
- Eyster, E., M. Rabin, and D. Vayanos (2018). Financial markets where traders neglect the informational content of prices. *Journal of Finance*.
- Fama, E. (1965). The behavior of stock-market prices. *Journal of Business*.
- Fama, E. and K. French (1997). Industry costs of equity. *Journal of Financial Economics*.
- Friedman, M. (1953). The case for flexible exchange rates. In *Essays in Positive Economics*.
- Gabaix, X. (2014). A sparsity-based model of bounded rationality. *Quarterly Journal of Economics*.
- Garleanu, N. and L. Pedersen (2011). Margin-based asset pricing and deviations from the law of one price. *Review of Financial Studies*.
- Gentzkow, M., B. Kelly, and M. Taddy (2017). Text as data. NBER.
- Glaeser, E. and C. Nathanson (2017). An extrapolative model of house-price dynamics. *Journal of Financial Economics*.

- Gong, B., D. Pan, and D. Shi (2016). New investors and bubbles: an analysis of the Baosteel call-warrant bubble. *Management Science*.
- Granovetter, M. (1978). Threshold models of collective behavior. *American Journal of Sociology*.
- Greenwood, R. and S. Nagel (2009). Inexperienced investors and bubbles. *Journal of Financial Economics*.
- Greenwood, R., A. Shleifer, and Y. You (2018). Bubbles for Fama. *Journal of Financial Economics*.
- Griffin, J., J. Harris, T. Shu, and S. Topaloglu (2011). Who drove and burst the tech bubble? *Journal of Finance*.
- Gromb, D. and D. Vayanos (2002). Equilibrium and welfare in markets with financially constrained arbitrageurs. *Journal of Financial Economics*.
- Grossman, S. (1976). On the efficiency of competitive stock markets where traders have diverse information. *Journal of Finance*.
- Guckenheimer, J. and P. Holmes (2013). *Nonlinear oscillations, dynamical systems, and bifurcations of vector fields*. Springer Science & Business Media.
- Han, B., D. Hirshleifer, and J. Walden (2018). Social-transmission bias and investor behavior. NBER.
- Hansen, S., M. McMahon, and A. Prat (2017). Transparency and deliberation within the FOMC: a computational linguistics approach. *Quarterly Journal of Economics*.
- Hasbrouck, J. (1991). Measuring the information content of stock trades. *Journal of Finance*.
- Hirsch, M., S. Smale, and R. Devaney (2012). *Differential equations, dynamical systems, and an introduction to chaos*. Academic Press.
- Hong, H., J. Kubik, and J. Stein (2004). Social interaction and stock-market participation. *Journal of Finance*.
- Hong, H. and J. Stein (1999). A unified theory of underreaction, momentum trading, and overreaction in asset markets. *Journal of Finance*.
- Horsthemke, W. and R. Lefever (2006). *Noise-induced transitions: theory and applications in physics, chemistry, and biology*. Springer Berlin Heidelberg.
- Kaustia, M. and S. Knüpfer (2012). Peer performance and stock market entry. *Journal of Financial Economics*.
- Kindleberger, C. (1978). *Manias, panics, and crashes*. Macmillan.
- Kuznetsov, Y. (2013). *Elements of applied bifurcation theory*. Springer Science & Business Media.

- Kyle, A. (1985). Continuous auctions and insider trading. *Econometrica*.
- Li, G. (2014). Information sharing and stock-market participation: evidence from extended families. *Review of Economics and Statistics*.
- Mackay, C. (1841). *Extraordinary popular delusions and the madness of crowds*. Richard Bentley, London.
- Manela, A. and A. Moreira (2017). News-implied volatility and disaster concerns. *Journal of Financial Economics*.
- May, R. (1974). *Stability and complexity in model ecosystems*. Princeton University Press.
- Miller, E. (1977). Risk, uncertainty, and divergence of opinion. *Journal of Finance*.
- Minsky, H. (1970). *Financial instability revisited: the economics of disaster*. Board of Governors of the Federal Reserve System.
- Minsky, H. (1992). The financial instability hypothesis. SSRN.
- Mullainathan, S. and A. Shleifer (2005). The market for news. *American Economic Review*.
- Nagel, S. (2005). Short sales, institutional investors, and the cross-section of stock returns. *Journal of Financial Economics*.
- Pearson, N., Z. Yang, and Q. Zhang (2017). Evidence about bubble mechanisms: precipitating events, feedback trading, and social contagion. SSRN.
- Reinhart, C. and K. Rogoff (2009). *This time is different: eight centuries of financial folly*. Princeton University Press.
- Safuan, H., Z. Jovanoski, I. Towers, and H. Sidhu (2013). Exact solution of a non-autonomous logistic population model. *Ecological Modeling*.
- Scheinkman, J. and W. Xiong (2003). Overconfidence and speculative bubbles. *Journal of Political Economy*.
- Schelling, T. (1978). *Micromotives and macrobehavior*. Norton.
- Shiller, R. (1984). Stock prices and social dynamics. *Brookings Papers on Economic Activity*.
- Shiller, R. (2000). *Irrational exuberance*. Princeton University Press.
- Shiller, R. and J. Pound (1989). Survey evidence on the diffusion of interest and information among investors. *Journal of Economic Behavior & Organization*.
- Shive, S. (2010). An epidemic model of investor behavior. *Journal of Financial and Quantitative Analysis*.
- Shleifer, A. and L. Summers (1990). The noise-trader approach to finance. *Journal of Economic Perspectives*.

- Shleifer, A. and R. Vishny (1997). The limits of arbitrage. *Journal of Finance*.
- Stevenson, R. (1886). *The strange case of Dr. Jekyll and Mr. Hyde*. New York: Scribner.
- Strogatz, S. (2014). *Nonlinear dynamics and chaos: applications to physics, biology, chemistry, and engineering*. Westview Press.
- Vayanos, D. (2004). Flight to quality, flight to liquidity, and the pricing of risk. NBER.
- Velhurst, P. (1845). Recherches mathématiques sur la loi d'accroissement de la population. *Proceedings of the Royal Academy of Brussels*.
- Xiong, W. and J. Yu (2011). The Chinese warrants bubble. *American Economic Review*.

## A Technical Appendix

**Proof** (Proposition 2.1). Suppose the excited-speculator population obeys the law of motion  $G(n, \theta, r) = \theta \cdot r \cdot (1 - n) \times n - n$ . We can rewrite this law of motion as

$$G(n, \theta, r) = (\theta \cdot r - 1) \times n - \theta \cdot r \times n^2$$

1. If  $r < 1/\theta$ , then  $(\theta \cdot r - 1) < 0$ . So, the only way for the right-hand side of the above equation to equal zero when  $r < 1/\theta$  is for  $n = 0$ . Thus, when  $r < 1/\theta$ ,  $\mathcal{SS}(\theta, r) = \{0\}$ . And, this unique steady state is stable since

$$\frac{\partial}{\partial n}[G(n, \theta, r)]_{n=0, r < 1/\theta} = \theta \cdot (r - 1/\theta) - 2 \cdot \theta \cdot r \cdot 0 < 0$$

2. If  $r > 1/\theta$ , then  $(\theta \cdot r - 1) > 0$ . So, there are now two ways for the right-hand side of the above equation to equal zero:  $n = 0$  and  $n = (r - 1/\theta)/r$ . Thus, when  $r > 1/\theta$ ,  $\mathcal{SS}(\theta, r) = \{0, (r - 1/\theta)/r\}$ . And, only the strictly positive steady state is stable since

$$\begin{aligned} \frac{\partial}{\partial n}[G(n, \theta, r)]_{n=0, r > 1/\theta} &= \theta \cdot (r - 1/\theta) - 2 \cdot \theta \cdot r \times 0 > 0 \\ \frac{\partial}{\partial n}[G(n, \theta, r)]_{n=(r-1/\theta)/r, r > 1/\theta} &= \theta \cdot (r - 1/\theta) - 2 \cdot \theta \cdot r \times (r - 1/\theta)/r < 0 \end{aligned}$$

□

**Proof** (Proposition 2.2). Newswatchers have demand given by:

$$x_{j,t} = (s_{j,t} - p_t)/\gamma$$

Market clearing then implies that:

$$\psi = \int_0^1 (s_{j,t} - p_t)/\gamma \cdot dj + \lambda \cdot r_{t-1} \times n_t$$

And, since the newswatcher signals are correct on average, we can conclude that:

$$\psi = (v_t - p_t)/\gamma + \lambda \cdot r_{t-1} \times n_t$$

Rearranging this equation so that the price is on the left-hand side gives the desired result. □

**Proof** (Proposition 3.1). The probability of realizing a speculative bubble at time  $t$  given that  $r_{t-2} < r_*$  can be written as

$$\mathbb{E}_{t-2}[\mathbb{B}(\theta, r_{t-1}) \mid \mathbb{B}(\theta, r_{t-2}) = 0] = \Pr_{t-2}[r_{t-1} > r_* \mid r_{t-2} < r_*]$$

And, given the stochastic process governing fundamentals, we know that

$$\begin{aligned} \mathbb{E}_{t-2}[\Delta v_{t-1}] &= \kappa_v \cdot (\mu_v - v_{t-2}) \\ \text{Var}_{t-2}[\Delta v_{t-1}] &= \sigma_v^2 \end{aligned}$$

Thus, given knowledge of  $v_{t-2}$ ,  $p_{t-2}$  and  $r_{t-2} < r_*$ , we can write the probability density function (PDF) for the price of the risky asset at time  $(t - 1)$  as

$$F_{t-2}(p) = \frac{1}{\sigma_v \cdot \sqrt{2 \cdot \pi}} \cdot e^{-\frac{1}{2 \cdot \sigma_v^2} \cdot (p - v_{t-2} - \mathbb{E}_{t-2}[\Delta v_{t-1}] + \gamma \cdot \psi)^2} \quad (\text{A1})$$

This PDF can be used to write down an integral expression for the probability of a speculative bubble at time  $t$  since  $r_{t-1} \stackrel{\text{def}}{=} p_{t-1}/p_{t-2}$ :

$$\mathbb{E}_{t-2}[\mathbb{B}(\theta, r_{t-1}) \mid r_{t-2} < r_*] = \int_{p_{t-2}/\theta}^{\infty} F_{t-2}(p) \cdot dp \quad (\text{A2})$$

And, the desired results follow from two facts about this integral expression. Fact #1: because  $F_{t-2}(p)$  is a PDF, it's a strictly positive function. Fact #2:  $\theta$  plays no part in  $F_{t-2}(p)$  itself; it only enters into Equation (A2) as a boundary condition. Combining these two facts implies that increasing  $\theta$  simply increases the size of the interval over which a strictly positive function is being integrated. Thus,  $\mathbb{E}_{t-2}[\mathbb{B}(\theta, r_{t-1}) \mid r_{t-2} < r_*]$  must be strictly increasing in  $\theta$ . □



**Proof** (Corollary 3.1). This corollary follows from the fact that excited speculators' extrapolative beliefs only affect equilibrium asset prices during a speculative bubble—i.e., when  $n_t > 0$ . However, the likelihood of entering into a speculative bubble is based on considerations made when there are currently no excited speculators in the market—i.e., when  $n_t = 0$ . More formally, the strength of excited speculators' extrapolative bias,  $\lambda$ , does not show up in either the PDF in Equation (A1) or the boundary conditions in Equation (A2).  $\square$

**Derivation** (Equation 27). Suppose an infinitesimal population of speculators,  $n_0 > 0$ , gets excited about the risky asset at time  $\tau = 0$ . The time it takes for half of this population to lose interest,  $\tau_{1/2}(\theta, r)$ , can be expressed as follows:

$$\begin{aligned}\tau_{1/2}(\theta, r) &= \int_0^{\tau_{1/2}} d\tau \\ &= \int_0^{\tau_{1/2}} \frac{dn}{dn} \cdot d\tau \\ &= \int_{n_0}^{\frac{n_0}{2}} \frac{d\tau}{dn} \cdot dn \\ &= \int_{n_0}^{\frac{n_0}{2}} G(n, \theta, r)^{-1} \cdot dn\end{aligned}$$

Since the initial population is infinitesimal,  $n_0 \approx 0$ , second-order terms will have a negligible impact on the law of motion governing the excited-speculator population:

$$G(n_0, \theta, r) = (\theta \cdot r - 1) \cdot n_0 + O[n_0^2]$$

In the equation above,  $O[n_0]$  represents ‘Big O’ notation. We say that  $f(x) = O[x]$  as  $x \rightarrow 0$  if there exists a positive constant  $C > 0$  such that  $|f(x)| \leq C \cdot x$  for all  $|x| < x_{\max}$ .

Equation (27), which characterizes the half-life of small populations of excited speculators when  $r < 1/\theta$ , follows from evaluating the integral expression for  $\tau_{1/2}(\theta, r)$  using only the first-order terms in  $G(n_0, \theta, r)$ :

$$\begin{aligned}\tau_{1/2}(\theta, r) &= \int_{n_0}^{\frac{n_0}{2}} \frac{1}{(\theta \cdot r - 1) \cdot n} \cdot dn \\ &= \int_{n_0}^{\frac{n_0}{2}} \frac{1}{\theta \cdot r - 1} \cdot \frac{1}{n} \cdot dn \\ &= (\theta \cdot r - 1)^{-1} \times \int_{n_0}^{\frac{n_0}{2}} n^{-1} \cdot dn \\ &= (\theta \cdot r - 1)^{-1} \times -\log(2)\end{aligned}$$

$\square$

**Proof** (Proposition 3.2). Suppose we look at small fluctuations,  $\epsilon \approx 0$ , in an asset's past return,  $r \mapsto r_\epsilon = r \cdot (1 + \epsilon)$  and define the following measure of co-movement between  $\epsilon$  and the half-life of small excited-speculator populations when  $r < r_\star$ :

$$C(\theta, r; \epsilon) = [\tau_{1/2}(\theta, r_\epsilon) - \tau_{1/2}(\theta, r)] \times \epsilon$$

The functional form for  $\tau_{1/2}(\theta, r)$  given by Equation (27) implies that for small fluctuations

in returns we have that both  $\partial_{\theta,r}[\tau_{1/2}] = \log(2) \cdot (1 - \theta \cdot r)^{-2}$  and

$$\begin{aligned} C(\theta, r; \epsilon) &= [\tau_{1/2}(\theta, r_\epsilon) - \tau_{1/2}(\theta, r)] \times \epsilon \\ &= \epsilon \cdot (\theta \cdot r) \cdot \partial_{\theta,r}[\tau_{1/2}(\theta, r)] \times \epsilon \\ &= \log(2) \cdot \epsilon^2 \cdot (1 - \theta \cdot r)^{-2} \cdot (\theta \cdot r) \end{aligned}$$

Thus, this measure of co-movement is strictly increasing in  $\theta$  when  $r < 1/\theta$  since:

$$\begin{aligned} \frac{\partial}{\partial \theta}[C(\theta, r; \epsilon)] &= \log(2) \cdot \epsilon^2 \cdot (-2) \cdot (1 - \theta \cdot r)^{-3} \cdot (-r) \cdot (\theta \cdot r) \\ &\quad + \log(2) \cdot \epsilon^2 \cdot (1 - \theta \cdot r)^{-2} \cdot r \\ &= \log(2) \cdot \underset{>0}{\epsilon^2} \cdot \underset{>0}{(1 - \theta \cdot r)^{-3}} \cdot \underset{>0}{r} \cdot \underset{>0}{(1 + \theta \cdot r)} \end{aligned}$$

If we further assume that  $\epsilon \sim N(0, \omega^2)$  for some small  $\omega > 0$ . Then, expectation of  $C(\theta, r; \epsilon)$  taken with respect to these small short-run fluctuations in returns will be equal to the following covariance

$$E[C(\theta, r; \epsilon)] = \text{Cov}[\tau_{1/2}(\theta, r_\epsilon) - \tau_{1/2}(\theta, r), \epsilon]$$

since both  $\epsilon$  and  $\tau_{1/2}(\theta, r_\epsilon) - \tau_{1/2}(\theta, r)$  will be mean zero ignoring higher-order terms.  $\square$

## B Some Extensions

This appendix examines three extensions of the economic model in Section 2 and shows that none of them qualitatively changes the implications of the model.

### B.1 Functional Forms

The law of motion in Equation (6) takes a very particular functional form. So, you might ask: how specific are the results in this paper to this choice? I now show that, to a second-order approximation, the growth rate in Equation (6) can be thought of as a stand-in for a broad range of models displaying the same steady-state behavior.

*Feedback Trading.* Feedback trading occurs when an initial positive shock generates excess media coverage and word-of-mouth buzz, which attracts new speculators to the market, which generates even more media coverage and word-of-mouth buzz, which excites still more speculators, which. . . . This narrative incorporates four key elements.

1. First, there must be some notion of a typical size for the excited-speculator population. Without loss of generality, let's normalize this size to  $n = 0$ .
2. Second, the excited-speculator population dynamics should reflect the fact that their population grows due to social interactions. "It is a very strong result here that both individual and institutional investors generally do not get interested in individual stocks by reading about them alone. There is a strong interpersonal component to investing, as hypothesized in epidemic models." (Shiller and Pound, 1989) Traders go "mad in herds" (Mackay, 1841). Thus, since it's harder to excite additional speculators when there are fewer apathetic agents left to interact with, the crowd of excited speculators should grow most rapidly when it's small and then grow more and more slowly as it gets larger and larger. In short, the arrival rate should be convex in the current population size,  $\frac{\partial^2}{\partial n^2}[\Theta(n, \theta, r)] < 0$ .
3. Third, the probability that each excited speculator loses interest and departs the crowd should be independent of this current population size. That same Mackay (1841) epigram says that traders "recover their senses slowly and one by one". In other words, the per

capital departure rate,  $\Omega(n)/n = \omega$ , should be constant. Without loss of generality, I'm going to further assume that  $\omega = 1$ . If  $n_\tau(\omega)$  and  $\theta(\omega)$  represent the true values that depend on  $\omega$ , then this renormalization is equivalent to re-defining  $n_\tau \stackrel{\text{def}}{=} n_\tau(\omega)/\omega$  and  $\theta \stackrel{\text{def}}{=} \theta(\omega)/\omega$ .

4. Finally, [Shiller \(2000\)](#) describes how “whenever the market reaches a new high, public speakers, writers, and other prominent people suddenly appear, armed with explanations for the apparent optimism seen in the market”. He points out that “the new era thinking they promote is part of the process by which a boom may be sustained and amplified—part of the feedback mechanism that...can create speculative bubbles”. Thus, the first speculators who get excited and enter the market should find it easier to attract additional friends to join them when past returns are higher,  $\frac{\partial^2}{\partial n \partial r}[\Theta(n, \theta, r)]_{n=0} > 0$ . But, to make sure we aren't assuming the result, these price changes shouldn't have any higher-order effects,  $\frac{\partial^2}{\partial n^2}[\Theta(n, \theta, r)] = 0$ .

The definition below converts these four elements into properties of the growth, arrival, and departure rates for the excited-speculator population.

**Definition B.1** (*Feedback Trading*). We say that the population dynamics of excited speculators are governed by feedback trading if the following four conditions are satisfied:

1. For all  $r > 0$ , we have that  $G(0, \theta, r) = 0$ .
2. For all  $n \in (0, 1)$  and  $r > 0$ , we have that  $\frac{\partial^2}{\partial n^2}[\Theta(n, \theta, r)] < 0$ .
3.  $\Omega(n) = n$ .
4. For all  $r > 0$ , we have that  $\frac{\partial^2}{\partial n \partial r}[\Theta(n, \theta, r)]_{n=0} > 0$  and  $\frac{\partial^2}{\partial n^2}[\Theta(n, \theta, r)] = 0$ .

*Logistic Approximation.* An excited-speculator population that obeys the logistic-growth model in Equation (6) is clearly governed by feedback trading. But, a population governed by the arrival rate  $\tilde{\Theta}(n, \theta, r) = r \cdot (1 - e^{-\theta n})$ . These growth rates look superficially different. But, it turns out that they lead to identical behavior in the neighborhood of  $r_\star$ .

**Proposition B.1** (*Logistic Approximation*). Suppose that a population of excited speculators obeys the law of motion  $\tilde{G}(n, \theta, r)$ . If there exists some  $r_- > 0$  such that  $\frac{\partial}{\partial n}[\tilde{G}(n, \theta, r_-)]_{n=0} < 0$  and these excited speculators engage in feedback trading (Definition B.1), then the population will also display a sudden qualitative change in steady-state behavior at a critical return threshold,  $r_\star > r_-$ .

Here's the intuition behind this result. First, if the population of excited speculators engages in feedback trading, then we know that the initial arrival rate is increasing in the price level,  $\frac{\partial^2}{\partial n \partial r}[\Theta(n, \theta, r)]_{n=0} > 0$ , for all  $r > 0$ . Higher returns make it easier for the first excited speculator to recruit more of his friends. And, we also know that there exists a return level,  $r_- > 0$ , such that the initial per capita growth rate of the crowd of excited speculators is negative,  $\frac{\partial}{\partial n}[G(n, \theta, r_-)]_{n=0} < 0$ . So, via the implicit-value theorem, we know that there must exist some critical return threshold,  $r_\star > r_-$ , such that

$$\frac{\partial}{\partial n}[G(n, \theta, r)]_{n=0} \begin{cases} < 0 & \text{if } r < r_\star \\ > 0 & \text{if } r > r_\star \end{cases}$$

Put differently, if we Taylor expand a law of motion that leads to feedback trading around the point,  $(0, \theta, r_\star)$ , then we see that the remaining criteria in the definition of feedback

trading imply that, to a second-order approximation, this growth rate must behave just like the logistic growth model. i.e., the restrictions in Definition B.1 imply that there exist positive constants,  $\varphi, \chi > 0$ , such that when  $n \in [0, \epsilon)$  and  $r \in (r_\star - \delta, r_\star + \delta)$  for sufficiently small values of  $\epsilon, \delta > 0$ :

$$G(n, \theta, r) = \varphi \cdot (r - r_\star) \cdot n - \chi \cdot n^2 + O[n^3]$$

Clearly, for a dynamical system with this functional form,  $n = 0$  is a steady-state solution for all  $r > 0$  since  $G(0, \theta, r) = \varphi \cdot (r - r_\star) \cdot 0 - \chi \cdot 0^2 = 0$ . What's more, given the derivative at zero,  $\varphi \cdot (r - r_\star)$ , we can see that  $n = 0$  will only be a stable steady-state solution when  $r < r_\star$ . As soon as  $r > r_\star$ , this steady-state solution will switch from stable to unstable as in Figure 4. And, notice that, if  $\varphi = \chi = \theta$  and  $r_\star = 1/\theta$ , then the growth rate in the equation above is identical to the logistic growth model. So, although this model is an extremely stylized model of social interactions, it's nevertheless emblematic of a more general phenomenon.

**Proof** (Proposition B.1). The definition of feedback trading implies the following sign restrictions for the derivatives of  $G(n, \theta, r)$ :

1.  $n = 0$  is a steady-state solution for all  $r > 0$  implies that  $\frac{\partial}{\partial r}[G(n, \theta, r)]_{n=0} = 0$ .
2.  $\frac{\partial^2}{\partial n^2}[\Theta(n, \theta, r)] < 0$  and  $\Omega(n) = n$  imply that  $\frac{\partial^2}{\partial n^2}[G(n, \theta, r)] < 0$ .
3.  $\frac{\partial^2}{\partial n \partial r}[\Theta(n, \theta, r)]_{n=0} > 0$  and  $\Omega(n) = n$  imply that  $\frac{\partial^2}{\partial n \partial r}[G(n, \theta, r)]_{n=0} > 0$ .
4.  $\frac{\partial^2}{\partial r^2}[\Theta(n, \theta, r)] = 0$  and  $\Omega(n) = n$  imply that  $\frac{\partial^2}{\partial r^2}[G(n, \theta, r)] = 0$ .

If i) there exists some  $r_- > 0$  such that  $\frac{\partial}{\partial n}[G(n, \theta, r_-)]_{n=0} < 0$  and ii) for all  $r > 0$  we have that both  $\frac{\partial^2}{\partial n \partial r}[G(n, \theta, r)]_{n=0} > 0$  and  $\frac{\partial^2}{\partial r^2}[G(n, \theta, r)]_{n=0} = 0$ , then via the implicit-value theorem there must be some critical return level,  $r_\star > r_-$ , such that

$$\frac{\partial}{\partial r}[G(n, \theta, r)]_{n=0} \begin{cases} < 0 & \text{if } r < r_\star \\ > 0 & \text{if } r > r_\star \end{cases}$$

Now, consider a Taylor expansion of  $G(n, \theta, r)$  around the point  $(0, \theta, r_\star)$  where  $n \in [0, \epsilon)$ ,  $\theta$  is constant, and  $r \in (r_\star - \delta, r_\star + \delta)$  for sufficiently small  $\epsilon, \delta > 0$ :

$$\begin{aligned} G(n, \theta, r) \approx & \overbrace{G(0, \theta, r_\star)}^{=0} + \overbrace{\frac{\partial}{\partial n}[G(0, \theta, r_\star)] \cdot n}^{=0} + \overbrace{\frac{\partial}{\partial r}[G(0, \theta, r_\star)] \cdot (r - r_\star)}^{=0} \\ & + \frac{1}{2} \cdot \frac{\partial^2}{\partial n^2}[G(0, \theta, r_\star)] \cdot n^2 + \frac{\partial^2}{\partial n \partial r}[G(0, \theta, r_\star)] \cdot n \cdot (r - r_\star) + \frac{1}{2} \cdot \underbrace{\frac{\partial^2}{\partial r^2}[G(0, \theta, r_\star)] \cdot (r - r_\star)^2}_{=0} \end{aligned}$$

Thus, in order for  $n > 0$  to be a steady-state solution, we must have that

$$0 = \frac{1}{2} \cdot \underbrace{\frac{\partial^2}{\partial n^2}[G(0, \theta, r_\star)] \cdot n^2}_{<0} + \underbrace{\frac{\partial^2}{\partial n \partial r}[G(0, \theta, r_\star)] \cdot n \cdot (r - r_\star)}_{>0}$$

This is only possible for  $n > 0$  if  $(r - r_\star) > 0$ , yielding a solution:

$$n = -2 \cdot \frac{\frac{\partial^2}{\partial n \partial r}[G(0, \theta, r_\star)]}{\frac{\partial^2}{\partial n^2}[G(0, \theta, r_\star)]} \cdot (r - r_\star) > 0$$

Furthermore, this positive solution will only be stable if

$$0 > \partial_n G(n, \theta, r) = \frac{\partial^2}{\partial n^2}[G(0, \theta, r_\star)] \cdot n + \frac{\partial^2}{\partial r \partial n}[G(0, \theta, r_\star)] \cdot (r - r_\star)$$

Plugging in the functional form for  $n$  yields:

$$\frac{\partial^2}{\partial n^2}[G(0, \theta, r_\star)] \cdot n + \frac{\partial^2}{\partial r \partial n}[G(0, \theta, r_\star)] \cdot (r - r_\star) = -\frac{\partial^2}{\partial r \partial n}[G(0, \theta, r_\star)] \cdot (r - r_\star)$$

Thus, we can conclude that the solution is stable for  $r > r_\star$  since  $\frac{\partial^2}{\partial r \partial n}[G(0, \theta, r_\star)] > 0$ . We've

just shown that any population growth rate that displays feedback trading will also display a sudden qualitative change in steady-state solutions at some critical value  $r_*$  that is identical in number and stability to the logistic growth model.  $\square$

## B.2 Random Fluctuations

What would happen if the excited-speculator population followed a stochastic law of motion? In the presence of random fluctuations, it's possible that the sudden change in the steady-state excited-speculator population as the risky asset's past return crosses  $r_* = 1/\theta$  disappears. This subsection shows that adding noise does not eliminate the sharp change in population dynamics around  $r_*$ .

*Stochastic Process.* Suppose that we redefine the law of motion in Equation (6) as follows:

$$\tilde{G}(n, \theta, r) \stackrel{\text{def}}{=} G(n, \theta, r) + \sigma_g \cdot n \cdot \frac{d\varepsilon_g}{d\tau}$$

In the equation above,  $\sigma_g > 0$  is a positive constant reflecting the instantaneous volatility of the excited-speculator population growth rate, and  $\varepsilon_g \sim N(0, 1)$  is a white-noise process. Introducing random fluctuations in this fashion implies that the excited-speculator population will adhere to the following stochastic law of motion:

$$dn_\tau = \theta \cdot (r - 1/\theta) \cdot n_\tau \cdot d\tau - \theta \cdot r \cdot n_\tau^2 \cdot d\tau + \sigma_g \cdot n_\tau \cdot d\varepsilon_{g,\tau} \quad \text{for } n_\tau \in [0, \infty) \quad (\text{B.2})$$

I've included time subscripts in the above equations,  $n_\tau$  and  $d\varepsilon_{g,\tau}$  rather than  $n$  and  $d\varepsilon_g$ , to emphasize which elements in this equation are time-varying and which are constants.

Equation (B.2) is just a noisy version of the law of motion described in Equation (6). But, there is one noteworthy difference:  $n_\tau \in [0, \infty)$  rather than  $n_\tau \in [0, 1)$ . Because the diffusion term  $\sigma_g \cdot n_\tau \cdot d\varepsilon_{g,\tau}$  contains  $n_\tau$ , the noise dies away as the excited-speculator population shrinks towards zero. And, as a result, the population will never go negative, which would be a physically meaningless outcome. But, it is possible for the population to exceed unity,  $n_\tau > 1$ . One way to make sense of this outcome is to think about the quantity  $U$  as the typical number of apathetic speculators in the market rather than the total number. Thus, whenever  $n_\tau > 1$ , there are more excited speculators in the market than the usual number of total speculators in the market. While it's possible to use  $N_\tau/U_\tau$  as the key state variable in the model (Safuan et al., 2013), doing so complicates the exposition without adding any new economic insight. So, I stick with the simpler setup in my analysis.

*Sudden Qualitative Change.* It turns out that, just like before, the stationary distribution for the excited-speculator population displays a sudden change in character as the asset's past return level crosses a critical return threshold,  $r_* = 1/\theta$ . When  $r < r_*$ , any initial population of excited speculators almost surely goes extinct; whereas, when  $r > r_*$ , this is no longer the case. Adding noise does not eliminate the sudden qualitative change as shown in Figure 12.


**Proposition B.2** (*Sudden Qualitative Change, Random Fluctuations*). *Suppose the excited-speculator population is governed by the law of motion in Equation (B.2) with  $\sigma_g = \sqrt{2}$ .*

1. *If  $r > r_* = 1/\theta$ , then the stationary distribution for  $\lim_{\tau \rightarrow \infty} n_\tau = n_\infty$  is characterized by the following probability-density function (PDF) given any initial  $n_0 \in (0, \infty)$*

$$n_\infty(\theta, r) \sim \text{Ga}(\theta \cdot r - 1, \theta \cdot r)$$

*where  $\text{Ga}(a, b) \stackrel{\text{def}}{=} \frac{b^a}{\Gamma(a)} \cdot \frac{x^{a-1}}{e^{b \cdot x}}$  is the PDF for the Gamma distribution.*

2. *If  $r < r_* = 1/\theta$ , then  $n_\infty(\theta, r) = 0$  almost surely.*



**Figure 12. Stationary Distribution.** Statistical properties of the stationary distribution for the excited-speculator population as characterized in Proposition B.2 where  $r_* = 1/\theta$ .

**Proof** (Proposition B.2). Let  $y_\tau \in (0, \infty)$  denote a stochastic process

$$dy_\tau = m(y_\tau) \cdot d\tau + \sigma_y \cdot s(y_\tau) \cdot d\varepsilon_{y,\tau}$$

where  $\tau \geq 0$ ,  $m(y)$  denotes the drift term,  $\sigma_y > 0$  is a positive constant,  $s(y) > 0$  is the diffusion term, and  $d\varepsilon_{y,\tau}$  is a standard Brownian-motion process. Assume that  $s(0) = 0$  and that  $m(\infty) = -\infty$ . Finally, let  $F_y(\tau, y_0)$  denote the probability-density function (PDF) for this stochastic process at time  $\tau \geq 0$  given the initial value  $y_0$ .

The Stratonovich interpretation of the Fokker-Plank equation dictates that:

$$\begin{aligned} \frac{\partial}{\partial \tau} [F_y(\tau, y_0)] = & - \frac{\partial}{\partial y} \left[ \left( m(y) + \frac{1}{2} \cdot \sigma_y^2 \cdot \frac{d}{dy} [s(y)] \cdot s(y) \right) \cdot F_y(\tau, y_0) \right] \\ & + \frac{1}{2} \cdot \sigma_y^2 \cdot \frac{\partial^2}{\partial y^2} \left[ s(y)^2 \cdot F_y(\tau, y_0) \right] \end{aligned}$$

And, a stationary distribution has the property that  $\frac{\partial}{\partial \tau} [F_y(\tau, y_0)] = 0$  for all  $y_0 \in [0, \infty)$ . Thus, the stationary distribution must satisfy the following condition:

$$\frac{\partial}{\partial y} \left[ \left( m(y) + \frac{1}{2} \cdot \sigma_y^2 \cdot \frac{d}{dy} [s(y)] \cdot s(y) \right) \cdot F_y(\tau, y_0) \right] = \frac{1}{2} \cdot \sigma_y^2 \cdot \frac{\partial^2}{\partial y^2} \left[ s(y)^2 \cdot F_y(\tau, y_0) \right]$$

This restriction, together with the boundary conditions that  $m(\infty) = -\infty$  and  $s(0) = 0$ , gives us the following functional form for the stationary distribution:

$$\begin{aligned} F_y(\cdot) = & \frac{1}{K \cdot s(y)} \cdot \exp \left( \frac{1}{\sigma_y^2/2} \cdot \int_0^y \frac{m(y')}{s(y')^2} \cdot dy' \right) \\ \text{given } K = & \int_0^\infty \frac{1}{s(y)} \cdot \exp \left( \frac{1}{\sigma_y^2/2} \cdot \int_0^y \frac{m(y')}{s(y')^2} \cdot dy' \right) \cdot dy < \infty \end{aligned}$$

If we substitute in the functional form for the excited-speculator population dynamics in Equation (B.2), then we have:

$$\begin{aligned} m(n) = & \theta \cdot (r - 1/\theta) \cdot n - \theta \cdot r \cdot n^2 \\ s(n) = & n \end{aligned}$$

So, when  $r > r_* = 1/\theta$ , the solution above dictates that:

$$F_n(r) = \frac{(\theta \cdot r)^{\theta \cdot r - 2}}{\Gamma(\theta \cdot r - 1)} \cdot \frac{n^{\theta \cdot r - 2}}{e^{\theta \cdot r \cdot n}}$$

And, this is the functional form of the PDF for the Gamma distribution,  $\text{Ga}(a, b) \stackrel{\text{def}}{=} \frac{b^a}{\Gamma(a)} \cdot \frac{x^{a-1}}{e^{b \cdot x}}$  with  $a = \theta \cdot r - 1$  and  $b = \theta \cdot r$ . This distribution is defined for all  $x \in (0, \infty)$ . When  $r < r_*$  this PDF is undefined. This corresponds to a solution where  $n_\tau = 0$  is an absorbing boundary. See Horsthemke and Lefever (2006, Ch. 6.4) for further details.  $\square$



### B.3 Continuous Feedback

Finally, in the economic model from Section 2, speculator interactions play out on a much faster timescale than assets are priced. First, speculators observe the risky asset's return in the previous period. Then, they interact with one another until a steady-state population has been reached. Finally, after this steady-state has been reached, any remaining excited speculators submit their demand to the market. What would happen if short-run changes in the excited-speculator population affected the risky asset's returns—i.e., what would change if there was continuous feedback between the excited-speculator population and the risky asset's return? I now show that, while modeling the continuous feedback between population dynamics and asset returns might seem more realistic, it turns out that this extension only affects the size of the steady-state excited-speculator population conditional on entering.

Consider an alternative law of motion where a 1% increase in the population of excited speculators increases the risky asset's return by a factor of  $\epsilon \in [0, \frac{1}{\theta \cdot r}]$ :

$$\tilde{G}(n, \theta, r) \stackrel{\text{def}}{=} \theta \cdot r \cdot (1 + \epsilon \cdot n) \cdot (1 - n) \times n - n \quad (\text{B.3})$$

The  $(1 + \epsilon \cdot n)$  term in the equation above captures the idea that an inflow of excited speculators at time  $\tau$  will increase the risky asset's return, which will then make it easier for future excited speculators to recruit their friends. If we set  $\epsilon = 0$ , then we get back the original law of motion in Equation (6). By increasing  $\epsilon$ , we allow transient fluctuations in the excited-speculator population to have a larger and larger effect on speculator persuasiveness via their effect on the asset's past returns.

**Proposition B.3** (*Sudden Qualitative Change, Continuous Feedback*). *Suppose the excited-speculator population is governed by the law of motion in Equation (B.3). Define  $r_\star \stackrel{\text{def}}{=} 1/\theta$ .*

1. *If  $r < r_\star$ , there's only one steady-state value for the excited-speculator population,  $\mathcal{SS}(\theta, r) = \{0\}$ . And, this lone steady state,  $\bar{n} = 0$ , is stable.*
2. *If  $r > r_\star$ , there are two steady-state values,  $\mathcal{SS}(\theta, r) = \{0, (1 - \epsilon)^{-1} \cdot (r - r_\star)/r > 0\}$ .*

*However, only the strictly positive steady state,  $\bar{n} = (1 - \epsilon)^{-1} \cdot (r - r_\star)/r > 0$ , is stable. In other words, continuous feedback doesn't affect the threshold return level,  $r_\star = 1/\theta$ , at which a non-zero population of excited speculators suddenly enters the market.*

Understanding the dynamics of the excited-speculator population interacts with an asset's past returns is very important on the intensive margin. It's very important if you want to understand how any particular bubble episode will unfold. But, it's not very important on the extensive margin. It's not essential if all you want to do is understand the likelihood that a crowd of excited speculators will enter the market in the first place.

**Proof** (Proposition B.3). The law of motion in Equation (B.3) can be re-written as follows

$$\tilde{G}(n, \theta, r) = (\theta \cdot r - 1) \times n - \theta \cdot r \cdot (1 - \epsilon) \times n^2 + O[n^3]$$

Thus, if we ignore third-order terms, then we can solve for the steady-state population of excited speculators using the same logic as in the proof of Proposition 2.1:

$$\bar{n} = \begin{cases} \frac{1}{1-\epsilon} \cdot \frac{r-r_\star}{r} & \text{if } r > r_\star \\ 0 & \text{otherwise} \end{cases}$$

The assumption that the strength of the continuous feedback is weak enough so that  $\epsilon < (\theta \cdot r)^{-1}$  ensures that  $\bar{n} < 1$ .  $\square$