The Madness Of Crowds And The Likelihood Of Bubbles*

Alex Chinco†

August 31, 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, $r < r_{\star}$, these social interactions will disperse any crowd of speculators that happens to get excited about the asset. But, as soon as the asset's past returns rise above this critical threshold, $r > r_{\star}$, the exact same social interactions will suddenly make the excited-speculator population boom. The resulting population explosion amplifies the effect of speculators' pre-existing psychological biases, causing arbitrageur constraints to bind and a speculative bubble to form.

The model predicts that speculative bubbles will be more common among assets whose speculators are more sensitive to fluctuations in past returns. More importantly, it also suggests how to estimate this key sensitivity parameter using data collected during normal times—i.e., when there's no speculative bubble currently taking place. And, I verify this prediction empirically: after similar price run-ups, industries in the top sensitivity quintile are more than twice as likely to experience a speculative bubble than those in the bottom quintile.

JEL Classification: G02, G11, G12

Keywords: Speculative Bubbles, Social Interactions, Displacement Events

*I would like to thank Zhi Da, Xavier Gabaix, Lawrence Jin, Neil Pearson, and Joel Peress for extremely helpful comments and suggestions. The paper has also benefited greatly from presentations at Insead and the Wabash River Conference.

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

†University of Illinois at Urbana-Champaign, Gies College of Business. alexchinco@gmail.com.

1 Introduction

Constrained arbitrageurs can't always correct the pricing errors caused by groups of biased speculators. This insight is known as the "limits of arbitrage" (Shleifer and Vishny, 1997). And, over the last 30 years, researchers have cataloged long lists of the many psychological biases afflicting speculators—e.g., overconfidence (Daniel et al., 1998; Scheinkman and Xiong, 2003), bounded rationality (Hong and Stein, 1999; Gabaix, 2014; Barberis et al., 2015), or sentiment (Baker and Wurgler, 2006)—and of the various trading constraints that arbitrageurs face—e.g., short-sale bans (Miller, 1977; Xiong, 2013), margin requirements (Gromb and Vayanos, 2002; Garleanu and Pedersen, 2011), equity limits (Shleifer and Vishny, 1997; Vayanos, 2004), or coordination frictions (Abreu and Brunnermeier, 2003).

By the logic of limits to arbitrage, any bias-constraint pair from these two lists might potentially combine at any moment to produce an equilibrium pricing error. There are myriad possibilities. But, here's the thing: large pricing errors such as speculative bubbles are rare. Kindleberger (1978) studied national price indexes and found only 34 bubble episodes dating back to the early 1600s, a period of nearly 400 years! Similarly, Greenwood et al. (2018) looked for industries that realized $\geq 100\%$ price appreciation over two years followed by a crash and found only 21 episodes since 1928 spread across 48 industries. There are 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 that some psychological bias will cause some trading constraint to bind?

Notice that these aren't questions we can answer within the existing limits-to-arbitrage framework. This framework explains how an equilibrium pricing error, such as a speculative bubble, can be sustained. But, we're not asking questions about *how* a speculative bubble can be sustained; we're asking questions about *how often* we should expect one 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 these questions, we 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 by recombining two common elements found in popular accounts of bubble formation in a new way: i) while speculators get overly excited following good news about fundamentals due to the "madness of crowds", ii) they only recover their senses "slowly and one by one" (Mackay, 1841).

Economic Model. There are three parts to my analysis. First, I develop an economic model to explain how displacement events work. Most of the model is taken directly from the standard limits-to-arbitrage playbook. The model focuses on a single risky asset and contains two kinds of traders: newswatchers and speculators. Newswatchers learn about the asset's fundamental value and incorporate this information into prices. Speculators only invest in the risky asset when excited, in which case they overvalue and have excess demand. Because the newswatchers are constrained, the excess demand coming from these excited speculators can push the 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 asset varies over time due to social interactions. These social interactions are governed by two key forces. First, speculators get overly excited about the asset following good news in a process resembling a social epidemic. Excited speculators are always trying to get the remaining apathetic speculators excited, 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 $r \in (0, \infty)$ is larger. This interaction between an asset's past return and its excited-speculator population is called "feedback trading" (Cutler et al., 1990).

However, feedback trading on its own only provides an *on switch* rather than an *on/off switch*. And, this is a problem because "only a relatively small proportion of large shocks leads to a speculative mania" (Kindleberger, 1978). So, to explain why feedback trading only sporadically amplifies speculator biases, I use 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 to explain why the "madness of crowds" (Mackay, 1841) might persist for a long time. But, in this paper, it explains why the madness of crowds only occasionally takes hold in the first place.

Combining these two forces allows small continuous changes in an asset's past returns to produce large discontinuous jumps in the number of excited speculators. I show that there exists a critical performance level, $r_\star$, such that, when $r < r_\star$, social interactions make the speculator population vanish. At sufficiently low values of r, the first speculator to get excited about the asset will tend to regain his senses before exciting any of his friends, which results is a market governed by fundamental values. However, when $r > r_\star$, the opposite intuition holds. The exact same social interactions now make the speculator population explode, causing arbitrageur constraints to bind and a speculative bubble to form. When this happens, we say that the market governed by "the madness of crowds" (Mackay, 1841), "speculative euphoria" (Minsky, 1970), "mob psychology" (Kindleberger, 1978), "irrational exuberance" (Shiller, 2000), etc... A positive shock to fundamentals that pushes the asset's return above $r_\star$ represents a "displacement event" (Minsky, 1992) in the model.

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 when it's initially still close to room temperature. It's just that, early on, 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. 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° 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 Implications. Next, I develop the testable implications of this model. The key observation 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 media coverage and 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. Thus, the model makes cross-sectional predictions about the likelihood of speculative bubbles as a function of the sensitivity of speculator persuasiveness to past returns. More sensitive assets should experience speculative bubbles more often.

More importantly, the model suggests how to estimate this sensitivity using data collected during normal times. This is completely new and different. If you want to learn about speculative bubbles using the existing limits-to-arbitrage framework, then you have to wait for arbitrageur constraints to bind. The relevant economic forces simply aren't operational during normal times. However, in a 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. In the model, it's possible to estimate each asset's speculator-persuasiveness sensitivity by studying the rate at which an asset's excited-speculator population decays during normal times in the same way that 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.

Econometric Analysis. Finally, I verify the model's main empirical prediction in monthly U.S. industry returns. I proxy for the each industry's excited-speculator population by counting the number of Wall Street Journal articles that reference an industry in their title per day. Then, to see how sensitive this population is to changes in past returns, I correlate the number of WSJ articles that reference an industry with that industry's monthly returns during the previous five years when no speculative bubble was taking place.

I define a speculative bubble as an episode where an industry realizes a 18-to-36 month boom where prices rise by more than 50% per year followed by a 18-to-36 month bust where prices fall by more than 25% per year. 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, the empirical results do not hinge on this exact definition. 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. I realize this might seem counter-intuitive: 'What hope is there of predicting the likelihood of speculative bubbles if we can't even agree about which past episodes should be classified as speculative bubbles?' But, this is a false choice. We know that suicide is the most common form of gun death in America[1] even though experts disagree about whether certain kinds of crime records should be classified as suicides or as accidents.[2]

Consistent with the prediction that assets with higher speculator-persuasiveness sensitivities will experience speculative bubbles more often, I find that industries in the top sensitivity quintile are more than twice as likely to experience a speculative bubble as those in the bottom quintile following good past performance. This result holds after controlling for changes in each industry's fundamental value. It also holds when controlling for the overall level of media coverage dedicated to the industry. Furthermore, I find that the estimated coefficients are stable over time. After fitting the model using the first half of the data sample, I show that the model's out-of-sample predictions for the second half of the data sample provide a measure of "market froth"[3] in an industry—i.e., the likelihood that the industry will experience a speculative bubble in the future. This bubble likelihood is something that gets talked about a lot, but it's something that previously had no analogue in economic models.

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 the existing literature.

Neo-Classical Models. The existing academic literature focuses on the question of existence because there's no such thing as a speculative bubble in neo-classical models. To illustrate, consider an asset with a $\$v$ per share payout tomorrow and three units available for purchase today. In a neo-classical model, speculators are fully rational rather than psychologically biased. In other words, they demand three units at a price of $p = \$v$ per share. In Figure

[1]"Gun Deaths In America." https://fivethirtyeight.com/features/gun-deaths/
[2]"What Counts As An Accident?" https://fivethirtyeight.com/features/gun-accidents/
[3]"Greenspan Is Concerned About 'Froth' In Housing." The New York Times. May 21, 2005.

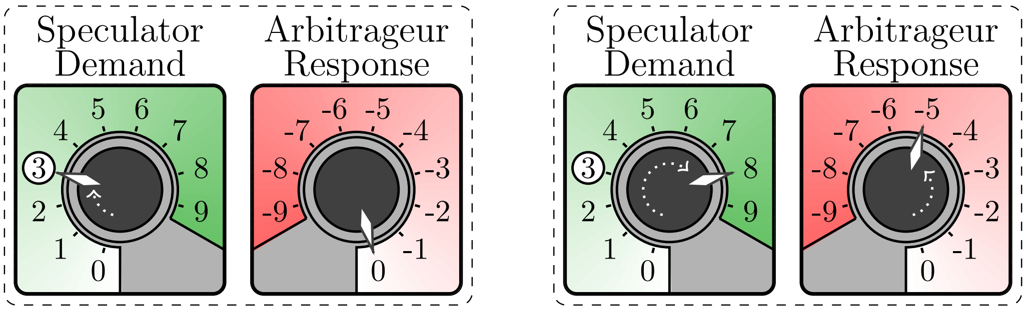

(a) Rational Demand **(b) Excess Demand**

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

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.

What's more, if there were some biased speculators who had excess demand, neo-classical models assume that there are rational arbitrageurs who would always be able to undo this error (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, $3 + 5 = 8$. Rational arbitrageurs could make money by selling five units of the asset at a price above $p = \$v$ per share to this group of biased speculators today and then buying back these units at the fundamental value tomorrow, 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 of 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. Instead of shorting the full five units as in Figure 1b, rational arbitrageurs might only be able to build a two-unit short position due to trading constraints, as shown by constraining the red arbitrageur-demand knob in Figure 2 to values of $\{0, -1, -2\}$. Thus, in the presence of both a psychological bias and a trading constraint, the price of an asset can rise above its fundamental value due to the resulting three extra units of aggregate demand coming from biased speculators, $(8 - 3) - 2 = 3$, after accounting for arbitrageurs' inadequate two-unit response. The limits-to-arbitrage framework provides necessary conditions for *how*

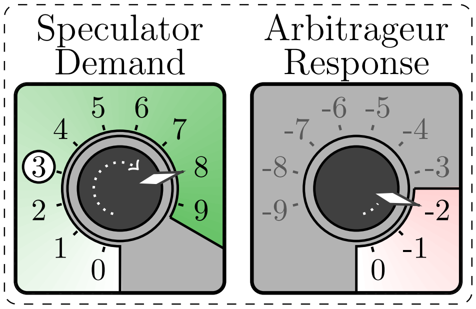

Figure 2. Limits to Arbitrage. Consider an asset with a $\$v$ payout tomorrow and three units available for purchase today (white circle). Green: speculator demand, $\{0, 1, \ldots, 9\}$, at price $p = \$v$. Red: arbitrageur response, $\{0, -1, -2\}$, when they can't short more than two units due to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

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 psychological bias. And, some trading constraint needs to prevent the red arbitrageur-response knob from undoing the error. But, it doesn't tell you *how often* you should expect to observe this configuration.

Displacement Events. By contrast, this paper is all about *how often* you should expect to observe this configuration. We know that 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 real-world 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 where they don't bind? And, how do displacement events switch us from a world like Figure 3a to a world like Figure 3b?

Predicting Crashes. There's been a lot of recent discussion about the importance of being able to predict the end of a speculative bubble. "They have to be predictable phenomena. I don't think any of this was particularly predictable." (Eugene Fama quoted in Cassidy, 2010) There's a sense that any complete theory of speculative bubbles should be able to 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. Any complete theory of speculative bubbles should describe both how speculative bubbles begin—the "Displacement"—as well as how they end—the "Crash". Both are important. 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 likely to form in the first place. Whether the resulting crash goes down in history books as "Black Monday" or "Black Tuesday" is utterly irrelevant.

Epidemic Models. The idea that social interactions might destabilize financial markets is not new. Several groundbreaking papers have explored this idea in great detail. However, this existing literature has focused on situations where biased speculator demand is always distorting prices. These papers ask: How might a social epidemic explain some regularly occurring pricing error? The outcome of interest in Shiller (1984) and Shive (2010) is return

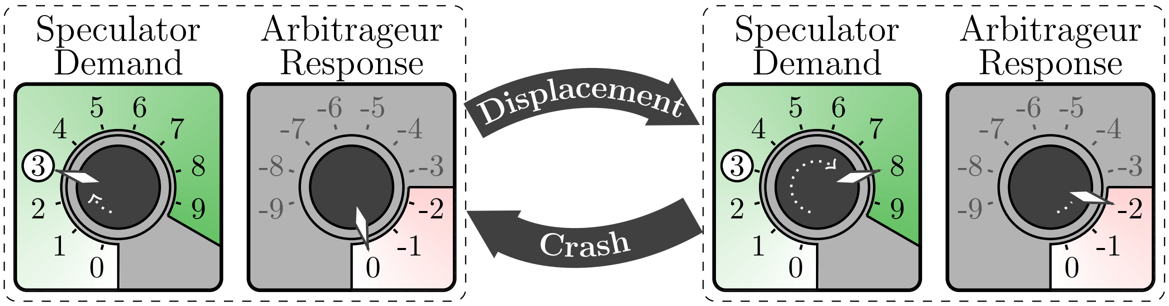

(a) Slack Constraint **(b) Binding Constraint**

Figure 3. How vs. How Often. Consider an asset with a $\$v$ payout tomorrow and three units available for purchase today (white circle). Green: speculator demand, $\{0, 1, \ldots, 9\}$, 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.

predictability in a *stationary* environment. Lux (1995) uses an epidemic model to formulate a "*cyclical* mechanism around fundamental values" in the stock market. Burnside et al. (2016) employs a similar model to study boom-bust *cycles* in the housing market. And, Han and Hirshleifer (2015) use social dynamics to explain the puzzling fact that active investors *persist* despite losing money. By contrast, this paper shows how social interactions can lead to the sudden emergence of a large pricing error, such as a speculative bubble. The goal is to explain when and where biased speculator demand is most likely to distort prices.

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 is not specifically about this particular psychological bias. It's about the likelihood that some psychological bias will cause some arbitrageur constraint to bind. The whole point is to distinguish the reason that excited speculators distort prices (in this case extrapolative beliefs) from the mechanism causing excited speculators to enter the market in the first place (the madness of crowds).

I choose to study extrapolative beliefs because the price spiral implied by this psychological bias explains key stylized facts about the life cycle of speculative bubbles 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 unfold once a crowd of excited speculators enters the market. As emphasized in Corollary 3.1, it doesn't affect the ex ante likelihood that a crowd of excited speculators will enter the market.

Furthermore, because existing papers operate entirely within the standard limits-to-arbitrage framework—i.e., because there's no on/off switch in these models—their proposed mechanism for how speculative bubbles start is totally different. 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 speculative bubbles occur, not when a new population of excited speculators decides to enter the market, but when the existing population fundamental traders decides to exit. In other words, displacement events in these papers are the result of rational egress and not irrational exuberance. This paper pairs the price spiral implied by extrapolative beliefs once a speculative bubble starts 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, each speculator doesn't always behave rationally. Sometimes a speculator will get overly excited about the asset and enter the market. The number of excited speculators in the market will ebb and flow over time due. Let $N_\tau \geq 0$ denote the number of speculators currently excited about the risky asset. Likewise, let $n_\tau \overset{\text{def}}{=} N_\tau / K$ denote the fraction of all speculators who are currently excited about the risky asset.

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 exact same kinds of arguments. However, the persuasiveness of these arguments is an increasing function of the asset's recent returns, $r \in (0, \infty)$. Currently excited speculators find it easier to persuade their friends following good 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. 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). Brunnermeier and Nagel (2004), Greenwood and Nagel (2009), and Griffin et al. (2011) all give evidence that sophisticated traders exploited inexperienced investors during the .com bubble. And, 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; Duflo 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; Bailey et al., 2016; Ahern, 2017). Traders are more likely to enter a market if they already know someone who’s done so.

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

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

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

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

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 parts. The trailing n term 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. Whereas, the leading expression, $\theta \cdot r \cdot (1 - n)$, represents this per capita excitation rate.

There are three factors to consider: $(1 - n)$, r, and θ. The first factor, $(1 - n)$, represents the fraction of the speculator population that’s currently apathetic. 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 to interact with. The second factor, $r \in (0, \infty)$, 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."[4] You can also see evidence of this behavior documented in Bailey et al. (2016). Finally, the third factor $\theta \in (0, 1)$ represents how much an asset's past returns affect the persuasiveness of its excited speculators—i.e., the speculator-persuasiveness sensitivity for this particular risky asset. It is an asset-specific constant. 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 media coverage and 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.

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 someone who belongs to the crowd of excited speculators at time τ happens to calm down and come to his senses during the time interval $(\tau, \tau + \Delta\tau]$:

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

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

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

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

[4] "Instagram, Facebook photos spur Millennials to become homeowners." USA Today. Apr 11, 2018.

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 $\Omega(n) = \omega \times n$ for some positive constant $\omega > 0$ is without loss of generality (see Appendix B.1). All that matters is that the rate at which each excited speculator calms down and recovers his senses is independent of both the size of the excited-speculator crowd, n, and the risky asset's past returns, r.

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:

$$\tfrac{\mathrm{d}n}{\mathrm{d}\tau} = \mathrm{G}(n, \theta, r) \stackrel{\mathrm{def}}{=} \Theta(n, \theta, r) - \Omega(n) \qquad \text{for } n \in [0, 1) \tag{5}$$

In other words, $\mathrm{G}(n, \theta, r)$ 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):

$$\mathrm{G}(n, \theta, r) = \theta \cdot r \cdot (1 - n) \times n - n \tag{6}$$

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

$$\mathrm{g}(n, \theta, r) \stackrel{\mathrm{def}}{=} \tfrac{1}{n} \cdot \mathrm{G}(n, \theta, r) = \theta \cdot r \cdot (1 - n) - 1 \qquad \text{for } n \in (0, 1) \tag{7}$$

The quantity $\mathrm{g}(n, \theta, r)$ 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 τ.

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):

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

Standard texts (see Arnol'd, 2012) show that, if $\mathrm{G}(n, \theta, r)$ 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 $\tau \geq 0$:

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

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 the governing equation in 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 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 this sudden qualitative change in a population of excited speculators governed by the logistic-growth model around the critical return threshold, $r_\star$.

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

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. 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 speculator currently excited about the risky asset, $N_\tau = n_\tau \cdot K = 1$. 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:

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

The expressions underneath the braces characterize these probabilities when $N_\tau = 1$ given

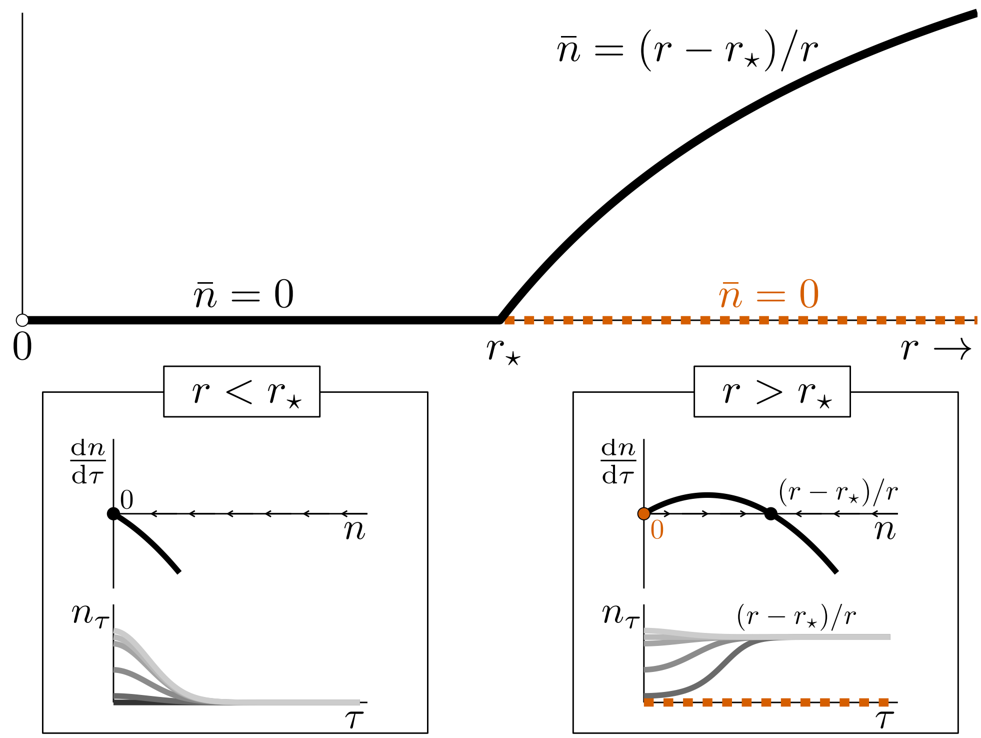

Figure 4. Sudden Qualitative Change. Top Panel. Steady-state solutions. x-axis: past return of the risky asset, $r \in (0, \infty)$. y-axis: steady-state solutions, $\bar{n} \in \mathcal{SS}(\theta, r)$, for a population of excited speculators governed by the law of motion in Equation (6). Solid black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

the functional forms in Equations (2) and (4). And, 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, $r < r_\star = 1/\theta$. 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. What's more, the exact same logic also will apply to new partner in crime, which means that the excited-speculator population will remain above zero in steady state. And, at first, the excited-speculator population will grow exponentially fast. 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.

The Calming Effect. It's important to emphasize that, 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 as an explanation for the persistence of mass hysteria, this paper uses it as an explanation why mass hysteria only occasionally takes hold. This paper uses the idea as an off switch rather than a repeater. 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 whenever their effects are weaker than this constant decay rate.

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 $\widetilde{\mathrm{G}}(n,\theta,r) \overset{\text{def}}{=} \Theta(n,\theta,r) = \theta \cdot r \cdot (1-n) \times n$ rather than $\mathrm{G}(n,\theta,r) = \Theta(n,\theta,r) - \Omega(n) = \theta \cdot r \cdot (1-n) \times n - n$ as in Equation (6). In this alternative model, the only steady-state solutions for the excited-speculator population would be $\widetilde{\mathcal{SS}}(\theta,r) = \big\{0,1\big\}$. 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. 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).

However, we can now use the model above to define the notion of 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 $r \mapsto r_\epsilon \overset{\text{def}}{=} r \cdot (1+\epsilon) > r_\star$.*

And, this definition naturally resolves all of the apparent irregularities mentioned above. 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 θ—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 $t = 1, 2, 3, \ldots$ 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:

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

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 $\varepsilon_{v,t} \sim \mathrm{N}(0,1)$ is a shock that's independently and identically distributed across time. Assume there are $\psi \geq 0$ shares of the risky asset.

Newswatchers. 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 jth 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:

$$x_{j,t} = \arg\max_x \mathrm{E}_j\big[-e^{-\gamma \cdot (v_t - p_t) \cdot x} \big] \tag{12}$$

Newswatcher utility displays constant relative risk-aversion. Above, $\mathrm{E}_j[\cdot]$ denotes the expectation of the jth 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 jth newswatcher observes a private signal about the risky asset's time-t payout. After observing this signal, the jth newswatcher believes the risky asset's payout can be decomposed as follows

$$v_t = s_{j,t} + \varepsilon_{j,t} \tag{13}$$

where $s_{j,t}$ denotes the jth newswatcher's beliefs after observing his private signal and $\varepsilon_{j,t} \sim \mathrm{N}(0,1)$ 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:

$$v_t = \mathrm{E}[s_{j,t}] = \textstyle\int_0^1 s_{j,t} \cdot \mathrm{d}j \tag{14}$$

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 $r_{t-1} \overset{\text{def}}{=} p_{t-1}/p_{t-2}$:

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

Again, 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 θ. 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. It also explains why there's no time subscript on r in Equation (8). 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? If there are $n_t > 0$ excited speculators in the market, then I model the aggregate demand coming from this crowd, z_t, as proportional to the asset's past returns:

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

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:

$$\int_0^1 x_{j,t} \cdot \mathrm{d}j + z_t = \psi \tag{17}$$

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 jth newswatcher in period t chooses his demand according to $x_{j,t} = \arg\max_x \mathrm{E}[-e^{-\gamma \cdot (v_t - p_t) \cdot x} | s_{j,t}]$.

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. (2015) 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 us 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. Again, the goal is not to add a new constraint to the already voluminous limits-to-arbitrage literature. The goal is to provide a theory of when and where these limits 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:

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

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:*

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

This price is increasing in the fundamental value, v_t; it's decreasing in the number of shares, ψ; 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

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

$$= 1[n_t > 0] \tag{20b}$$

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 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 $r_\star = 1/\theta$.

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 $\mathrm{B}(\theta, r) = 1$. Each of these instances

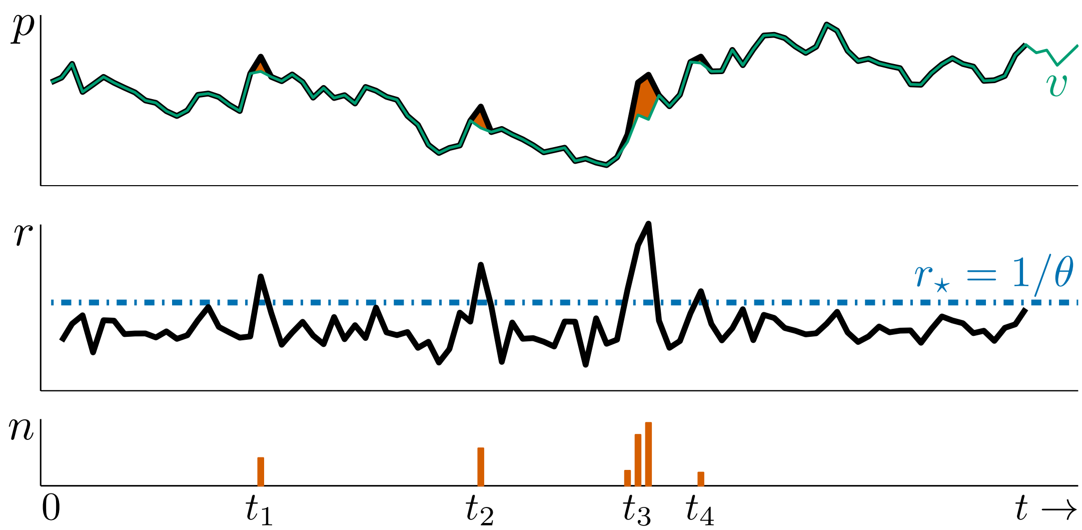

Figure 5. Sample Realization. Data simulated using parameters $\psi = 0$, $\mu_v = 1.0$, $\kappa_v = 0.1$, $\sigma_v = 0.1$, $\theta = 0.4$, and $\lambda = 0.5$. x-axis in all panels represents time, $t = 1, 2, \ldots, 100$. **Top Panel.** Black line is risky asset's price, p. Thin green line is its 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, $r_\star = 1/\theta$. When $r < r_\star$, there are no excited speculators in the market, $n = 0$. When $r > r_\star$, 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 $\mathrm{B}(\theta, r) = 1$ are labeled on the x-axis, $\{t_1, t_2, t_3, t_4\}$.

is denoted by a label, $\{t_1, t_2, t_3, t_4\}$, on the x-axis in the bottom panel. Recall that $\mathrm{B}(\theta, r) = 1$ implies that $r > r_\star = 1/\theta$, 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. So, 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. 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 $\widetilde{\mathrm{G}}(n,\theta,r) \overset{\text{def}}{=} r \cdot (1 - e^{-\theta \cdot n}) - n$ would deliver qualitatively similar population dynamics as using $\mathrm{G}(n,\theta,r)$ even though $\widetilde{\mathrm{G}}(n,\theta,r)$ and $\mathrm{G}(n,\theta,r)$ look superficially very different. In Appendix B.1 I give conditions under which an alternative law of motion, $\widetilde{\mathrm{G}}(n,\theta,r)$, will exhibit feedback trading just like $\mathrm{G}(n,\theta,r)$.

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:

$$\mathrm{G}(n,\theta,r) \mapsto \widetilde{\mathrm{G}}(n,\theta,r) \overset{\text{def}}{=} \mathrm{G}(n,\theta,r) + \sigma_n \cdot n \cdot \tfrac{\mathrm{d}\varepsilon_g}{\mathrm{d}\tau} \tag{21}$$

In the equation above, $\sigma_n > 0$ is a positive constant reflecting the instantaneous volatility of the excited-speculator population growth rate, and $\varepsilon_{g,\tau} \sim \mathrm{N}(0,1)$ 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 $\epsilon \in [0, \frac{1}{\theta \cdot r})$:

$$\widetilde{\mathrm{G}}(n,\theta,r) \overset{\text{def}}{=} \theta \cdot r \cdot (1 + \epsilon \cdot n) \cdot (1 - n) \times n - n \tag{22}$$

The $(1 + \epsilon \cdot n)$ term in the equation above captures the idea that an inflow of excited speculators at time τ 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 ϵ, 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 return level, $r_\star = 1/\theta$, at which a non-zero population of excited speculators suddenly enters the market. And, this paper is entirely about where this threshold return level 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.

3 Empirical Implications

Let's now explore the empirical implications of this model. Assets where speculator persuasiveness is more sensitive to past returns should experience speculative bubbles more often. And, more importantly, it's possible to estimate this key sensitivity parameter 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 the key sensitivity parameter for each asset by studying how quickly an asset's excited-speculator population decays during normal times.

3.1 Cross-Sectional Prediction

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

Bubble Likelihood. Which asset's are most likely to experience a speculative bubble following 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, $r > r_\star$. When this happens, a crowd of excited speculators enters the market, $n > 0$. And, the extrapolative demand coming from these excited speculators causes arbitrageur 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, $r_\star$, determined for each asset? ii) How do newswatchers incorporate news 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, $r_\star = 1/\theta$. 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):

$$\mathrm{E}_{t-1}[\Delta v_t] = \kappa_v \cdot (\mu_v - v_{t-1}) \tag{23a}$$

$$\mathrm{Var}_{t-1}[\Delta v_t] = \sigma_v^2 \tag{23b}$$

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

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

$$\tfrac{\partial}{\partial\theta}\mathrm{E}_{t-1}[\, \mathrm{B}(\theta, r_t) \,|\, r_{t-1} < r_\star \,] > 0 \tag{24}$$

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

Economic Intuition. The economic intuition behind this result is straightforward. When house prices rise by 10%, it's national news. Everyone starts talking about housing. And, it becomes a lot easier for second-home buyers to convince their friends to speculate in the housing market, too (Bailey et al., 2016). 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 in this figure display 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 by $r_\star = 1/0.4 = 2.5$ 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, $r_\star = 1/0.2 = 5.0$. 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, $2.5 < r_{t_1-1} < 5.0$ 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 *'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 λ:*

$$\tfrac{\partial}{\partial\lambda}\mathrm{E}_{t-1}[\,\mathrm{B}(\theta, r_t)\,|\,r_{t-1} < r_\star\,] = 0 \tag{25}$$

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."[5] 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?[6] In U.S. stocks?[7] Or, maybe in Chinese real estate?[8] 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

Proposition 3.1 characterizes how the likelihood of a speculative bubble should vary across assets as a function of θ—i.e., the sensitivity of each asset's speculators to its past returns. More sensitive assets should experience speculative bubbles more often. To test this prediction, we have to be able to estimate this key sensitivity parameter for each asset. And, we have to be able to compute this estimate using data collected during normal times prior to any speculative bubble taking place. And, the economic model above suggests how to do this.

[5] *China Blowing Major Bubbles In 2017.* Forbes. Dec 19, 2016.
[6] "Bitcoin is heading to \$10,000" CNBC. Oct 20, 2017.
[7] "Goldman's Blankfein: Things Have Been Going Up for Too Long" Wall Street Journal. Sep 6, 2017.
[8] "Chinese Efforts to Stem Housing Bubble Show Promise" Bloomberg. Jun 11, 2017.

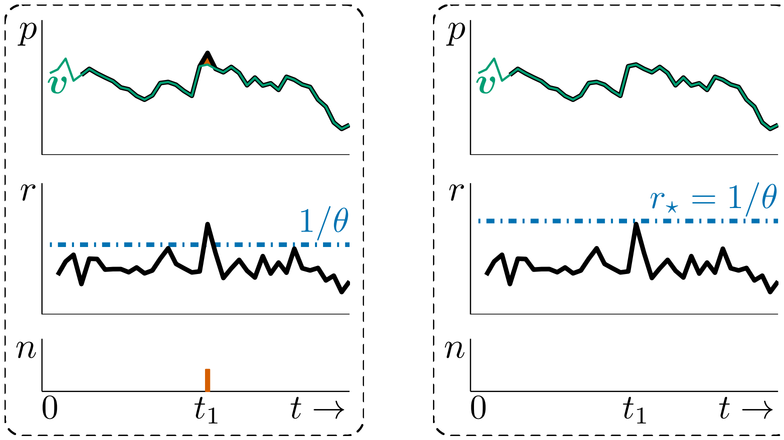

(a) $\boldsymbol{\theta = 0.4}$ **(b) $\boldsymbol{\theta = 0.2}$**

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 $t = 1, 2, \ldots, 25$. 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, $r_\star = 1/\theta$. Red bars, bottom: number of excited speculators in the market, n. $\boldsymbol{\theta = 0.4}$. 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 $r_{t_1 - 1} > 1/0.4 = 2.5$. $\boldsymbol{\theta = 0.2}$. 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 $r_{t_1 - 1} < 1/0.2 = 5.0$.

Economic Intuition. A theory of displacement events must explain not only why speculative bubbles sometimes occur but also why they typically do not. The exact same economic forces must be at work both during and between bubble episodes because these forces. Saying that individual investors act according to completely different rules during speculative bubbles is no explanation at all. It's like saying: "The Gods did it." (Deutsch, 2011) Thus, it should be possible to learn about the social interactions that sometimes cause the excited-speculator population to explode by studying how these same social interactions usually cause the excited speculator population to vanish. The economic model above suggests doing this by estimating the rate at which an asset's excited-speculator population decays during normal times. In the bottom panels of Figure 4, the excited-speculator population converges towards its steady state $\bar{n}$ at the same rate regardless of whether this steady state is $\bar{n} = 0$ (left, $r < r_\star$) or $\bar{n} > 0$ (right, $r > r_\star$).

Normal-Times Estimate. Here's how I statistically implement this economic intuition. 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:

$$\tau_{1/2} \overset{\text{def}}{=} \min \left\{ \tau > 0 \, : \, \tfrac{n_\tau}{n_0} \leq \tfrac{1}{2} \right\} \tag{26}$$

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

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

The idea is to learn about the sensitivity of each asset's speculators to past returns, θ, by studying how small changes in past returns affect this half-life. Let ϵ denote a small fluctuation in past returns, $r \mapsto r_\epsilon \overset{\text{def}}{=} r \cdot (1+\epsilon)$. Think about ϵ 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 five years while ϵ will correspond to the asset's return in the most recent month. And, 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:

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

If increasing r 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:

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

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

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

$$\tfrac{\partial}{\partial\theta}\mathrm{E}[\mathrm{C}(\theta, r; \epsilon)] > 0 \tag{30}$$

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

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

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 θ*.*

Figure 7 depicts the central intuition underpinning this result. 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, $r < r_\star$. 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 $r_\epsilon < r_\star$. 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 7a to Panel 7b

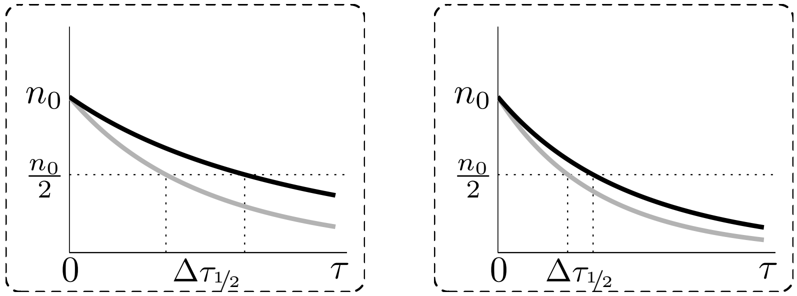

(a) $\boldsymbol{\theta = 0.4}$ **(b) $\boldsymbol{\theta = 0.2}$**

Figure 7. Normal-Times Estimate. Transition of an initial population of $n_0 > 0$ excited speculators to the steady-state value of $\bar{n} = 0$ when $r < r_\star$. x-axis: time $\tau \geq 0$. Grey line: size of excited-speculator population at time τ given past return r, $n_\tau(\theta, r)$. Black line: size of of excited-speculator population at time τ 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 ϵ-fluctuation in past returns on this half-life and corresponds to the horizontal distance between the two vertical dotted lines. $\boldsymbol{\theta = 0.4}$**.** Transition in a market where the risky asset has a parameter value of $\theta = 0.4$ as in Figure 5. $\boldsymbol{\theta = 0.2}$**.** Transition in an otherwise identical market where $\theta = 0.2$.

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 θ? Proposition 3.2 provides a way of estimating the sensitivity of speculator persuasiveness to past returns using data collected during normal times. And, it does so without specifying why some assets have a higher value of θ than others. Put another way, this paper provides a model of how cross-sectional differences in θ will affect asset prices, but it does not explain where these cross-sectional differences come from. This omission might seem problematic at first. But, it is a feature not a bug.

While it would be interesting to understand the sources of variation in θ, 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 .com stocks an exciting proposition to day traders in the late 1990s? Were those mid-2000s housing 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 easy-to-estimate parameters. To illustrate this fundamental principle, 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 λ, will be proportional to the ratio of an asset's fundamental volatility to its demand-noise volatility. Thus, asset's with more volatile fundamentals will be less liquid because market makers are at a larger informational disadvantage. By contrast, asset's 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 λ (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 and others have more volatile 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 verifies the above economic model's main empirical prediction using data on the cross-section of monthly U.S. industry returns. Consistent with the prediction that assets with higher speculator-persuasiveness sensitivities will experience speculative bubbles more often, I find that industries in the top sensitivity quintile are more than twice as likely to experience a speculative bubble as those in the bottom quintile following good past performance. Moreover, after fitting the model using only first half of the data sample, I show that the model makes accurate out-of-sample predictions about the likelihood of a speculative bubble in during the second half of the data sample.

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 econometric variables observed in the data. For example, $\texttt{PastReturn}_{i,t}$ will denote the statistical estimate for the ith industry's past return in month t, which corresponds to the model parameter r.

Industry Returns. Following Greenwood et al. (2018), I use the Fama and French (1997) classification scheme to assign each stock in the CRSP database to one of 49 industries.[9] I use

[9] 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

this database to compute the value-weighted return of the stocks in each industry, $\texttt{Return}_{i,t}$, during the period from January 1950 to December 2017. This yields a sample that contains $2017 - 1949 = 73$ years of monthly observations, which amounts to a total of 816 monthly observations for each industry.

When computing each industry's monthly return, I restrict the universe of stocks to include only stocks that are actively traded on U.S. exchanges with a share code of 10, 11, and 12. I also do not consider the returns of stocks in the "other" industry because this does not represent a cohesive collection of stocks that are connected to one another in any meaningful way. There's no way for speculators to have an exciting conversation about a miscellaneous group of unrelated stocks. Rather than use the time series reported by Ken French on his website, I compute industry returns manually so as to include recently listed stocks, which have historically been an important part of speculative bubbles in the stock market (see Greenwood et al., 2018).

To provide a sense of what these returns look like, the solid black lines in Figures 8a and 8b depict the outcome of continuously re-investing \$1 in each industry starting in January 1950. In other words, the black lines depict the quantity $\$1 \times \prod_{h=1}^{816}(1 + \texttt{Return}_{i,\text{Dec1949}+h})$. If you had invested \$1 in the aerospace industry in January 1950, then you would have had an additional \$5,271.75 at the end of December 2017. Whereas, if you had invested that same \$1 in the steel industry, you would have only ended up with an extra \$483.59.

I estimate the past return of each industry, $\texttt{PastReturn}_{i,t}$, by computing the cumulative return of the industry over the previous five years:

$$\texttt{PastReturn}_{i,t} \overset{\text{def}}{=} \prod_{h=59}^{0} (1 + \texttt{Return}_{i,t-h}) \tag{32}$$

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 ith industry five years ago. If $\texttt{PastReturn}_{i,t} = 0$, then they would have lost their entire \$1 investment by month t; whereas, if $\texttt{PastReturn}_{i,t} = 2$, then their initial \$1 investment would have doubled in value by month t. Note that this definition of $\texttt{PastReturn}_{i,t}$ also reflects what happens when speculators study price indexes that have been normalized to one at a previous date (Case and Shiller, 1987).

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) whlsl: wholesale.

Bubble Episodes. I say that industry i experienced a speculative bubble with a peak in month t_p if it realized annualized returns of $\geq 50\%$ during the 18-to-36 month period prior to the peak and annualized returns of $\leq -25\%$ during the 18-to-36 month period after the peak. In other words, for there to be a speculative bubble with a peak at month t_p, there must exist $\texttt{BoomLen}_{i,t_p}, \texttt{BustLen}_{i,t_p} \in \big\{18, 19, \ldots, 36\big\}$ such that

$$\left[\prod_{h=\texttt{BoomLen}_{i,t_p}-1}^{0} (1 + \texttt{Return}_{i,t_p-h}) \right]^{12/\texttt{BoomLen}_{i,t_p}} - 1 > +50\% \tag{33a}$$

$$\left[\prod_{h=1}^{\texttt{BustLen}_{i,t_p}} (1 + \texttt{Return}_{i,t_p+h}) \right]^{12/\texttt{BustLen}_{i,t_p}} - 1 < -25\% \tag{33b}$$

Note that this paper uses a slightly different definition of a speculative bubble than the one used in Greenwood et al. (2018). That paper considers episodes where an industry realizes $\geq 100\%$ returns over a fixed two year period, both raw and net of the market. The authors fix the length of the boom to two years because they want to test whether it's possible to predict whether an industry-level boom will be followed by a subsequent bust. However, this paper is interested in predicting the start of the entire boom-bust episode rather than the peak. So, there's no reason to fix the length of the boom component. And, the sample realization plotted in Figure 5 even suggests that the length of the boom will vary across bubble episodes.

Let $\texttt{BubbleStart}_{i,t}$ be an indicator variable denoting the start of one of these boom-bust episodes. Specifically, $\texttt{BubbleStart}_{i,t} = 1$ if and only if there exists some t_p such that $\texttt{BoomLen}_{i,t_p} = (t_p - t) + 1$ and $\texttt{BubbleStart}_{i,t-1} = 0$. Returning to Figures 8a and 8b, we see that there are 21 such bubble episodes spread across all 48 industries from January 1950 to December 2017. In the figure, the vertical grey bars denote the boom component of each episode while the vertical red bars demand the subsequent bust. The numbers above these bars count the total number of speculative bubbles since January 1950. Table 1 then provides summary stats for these bubble episodes. On average, industries earn annualized returns of 64.13% during booms that typically last 1.83 years. But, they lose 49.84% annually during the subsequent bust, which tends to last 1.71 years on average. This means that investors lose 18.60% of their wealth during speculative bubbles on average. Note that investors lose money even though their gains during the boom are proportionally larger than their losses during the bust since these proportional losses get applied to a much larger portfolio. And, a quick back-of-the-envelope calculation confirms that $8/5 \cdot 1/2 = 4/5$.

Estimated Theta. If we want to use the model from Section 2 to predict which industries are most likely to experience a speculative bubble, then we need to estimate the key sensitivity

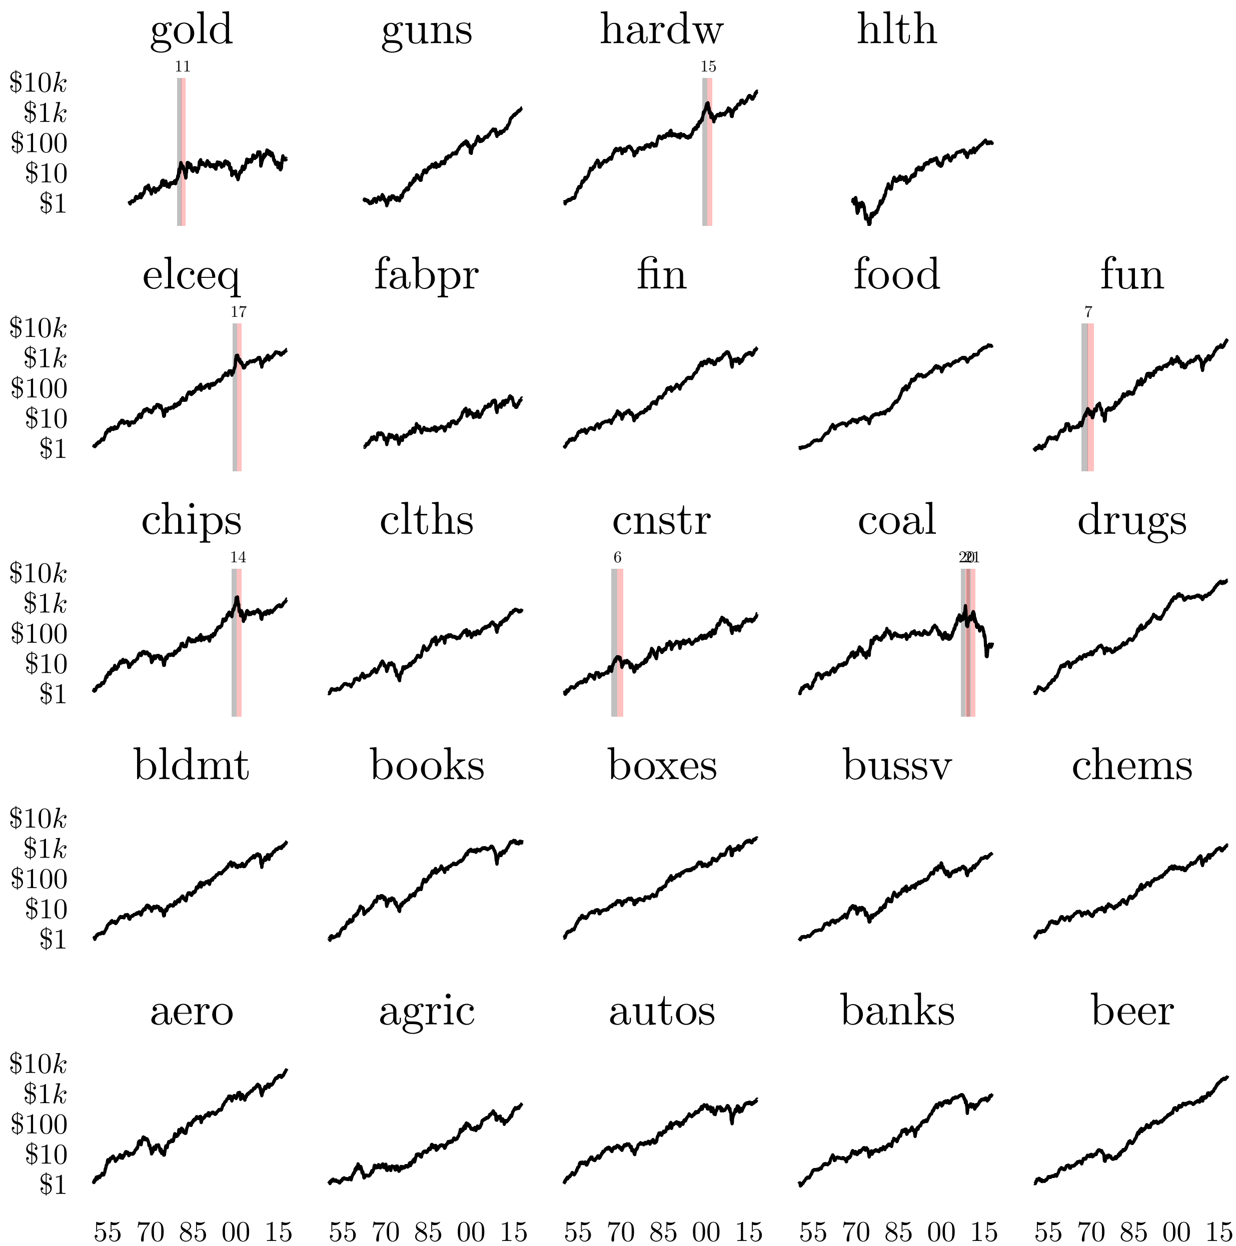

Figure 8a. Bubble Episodes. x-axis: time in months from January 1950 to December 2017. y-axis: dollar value of a continuously re-invested portfolio that started with $1 in January 1950 on a logarithmic scale. Each panel represents results for an industry-specific portfolio, which are labeled using the abbreviations provided in Fama and French (1997). Thick black line: portfolio size in month t. Vertical grey/red bars: boom-bust episodes. Number above each bar: total number of bubble episodes since January 1950.

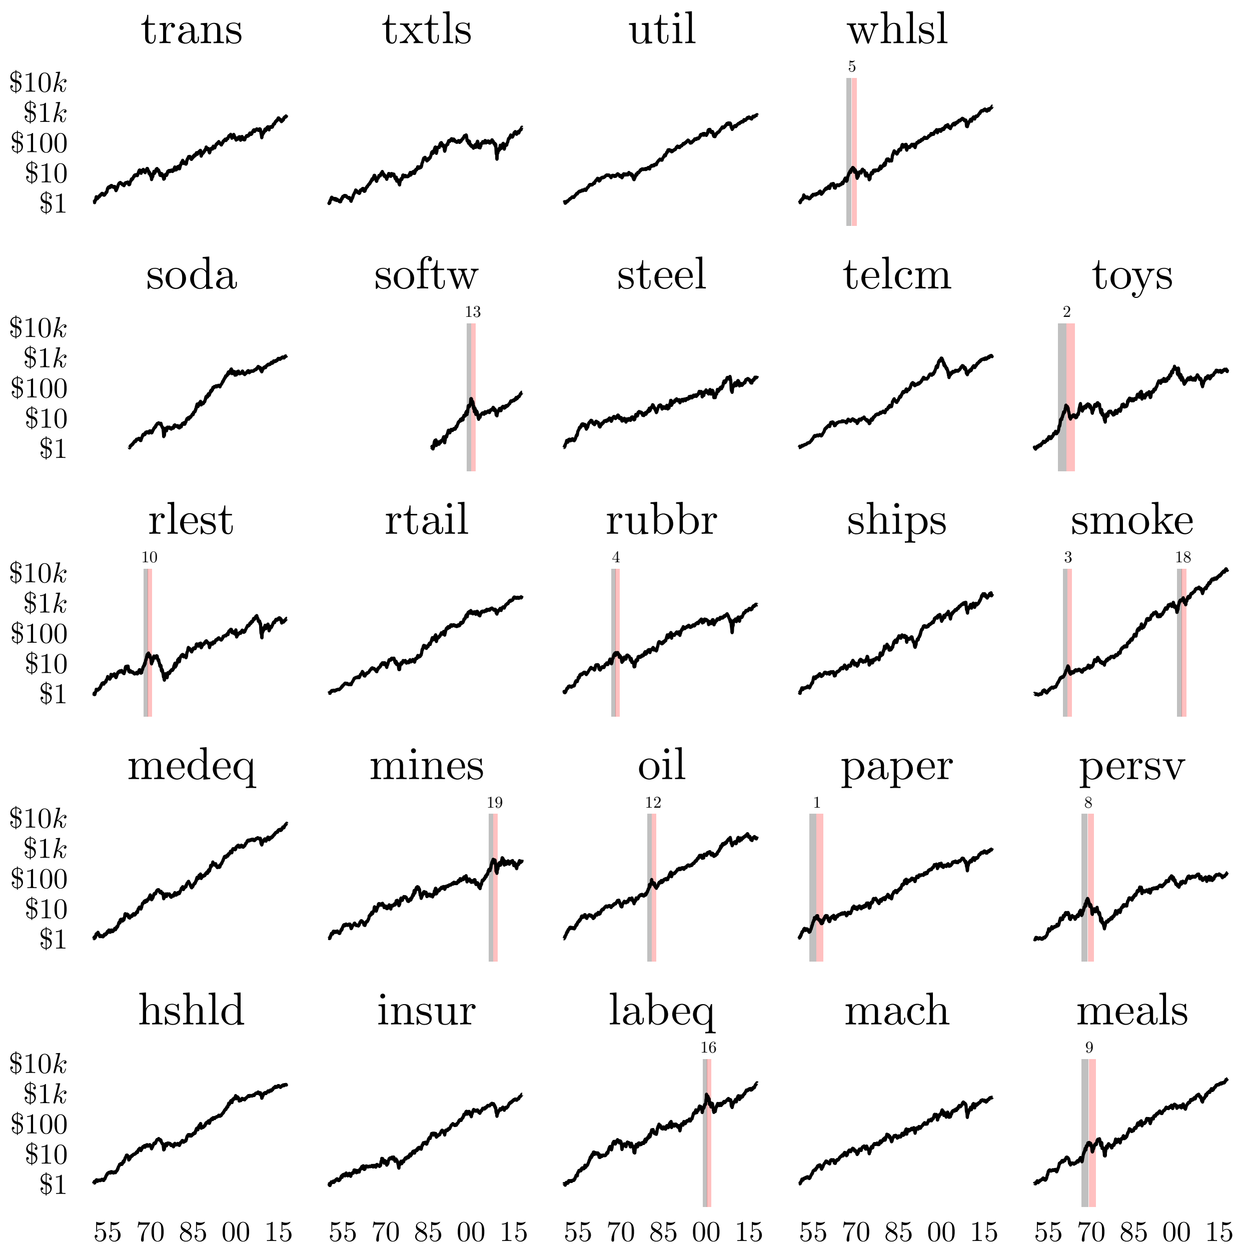

Figure 8b. Bubble Episodes, Ctd. x-axis: time in months from January 1950 to December 2017. y-axis: dollar value of a continuously re-invested portfolio that started with \$1 in January 1950 on a logarithmic scale. Each panel represents results for an industry-specific portfolio, which are labeled using the abbreviations provided in Fama and French (1997). Thick black line: portfolio size in month t. Vertical grey/red bars: boom-bust episodes. Number above each bar: total number of bubble episodes since January 1950.

parameter, θ. Proposition 3.2 suggests a way to do this by looking at the correlation between short-run fluctuations in an industry's returns and changes in the rate at which excited speculators lose interest in the industry. I proxy for changes in the half-life of the excited-speculator population 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 $r < r_{\star}$, then it's likely that the transient population of excited speculators is remaining in the market longer.

Here's how I estimate each industry's media coverage in a given month. First, since an article about agriculture might talk about things like *'crop planting'* or *'pesticides'* in the title rather than *'agriculture'*, I construct a set of industry-specific keywords using the SIC-subgroup descriptions found on the U.S. Department of Labor's website.[10] For example, the agriculture industry has keywords *'agriculture'*, *'pesticide'*, *'cotton gin'*, *'veterinary'*, *'crop harvesting'*, *'crop planting'*, *'landscaping'*, *'farm'*, *'crops'*, *'livestock'*, *'cattle'*, *'hogs'*, *'swine'*, *'farm animals'*, *'commercial fishing'*, and *'animal feed'*. Then, for each industry, I search the entire historical Wall Street Journal archives for articles that include that industry's keywords in its title. These archives are available via ProQuest, and they cover the time period from January 1950 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 $\texttt{\#Articles}_{i,t}$ denote the average number of articles per day during month t that include one of the ith industry's keyword in the title. I then define the percent of all media coverage dedicated to the ith industry in month t as follows:

$$\texttt{\%Articles}_{i,t} \overset{\text{def}}{=} 100 \times \left(\frac{\texttt{\#Articles}_{i,t}}{\sum_{i'} \texttt{\#Articles}_{i',t}} \right) \tag{34}$$

I proxy for θ by computing the partial correlation between each industry's return in month t and its media coverage in month t estimated during the previous five years:

$$\texttt{Theta}_{i,t} \overset{\text{def}}{=} \frac{\sum_{h=0}^{59} (\texttt{\%Articles}_{i,t-h+1} - \langle \texttt{\%Articles}_i \rangle) \cdot (\texttt{Return}_{i,t-h+1} - \langle \texttt{Return}_i \rangle)}{\sum_{h=0}^{59} (\texttt{Return}_{i,t-h+1} - \langle \texttt{Return}_i \rangle)^2} \tag{35}$$

This just means regressing the percent of all Wall Street Journal articles that mention the ith industry in month t' on the ith industry's returns in month t' using data on the 60 months $t' \in \{t-59, \ldots, t\}$. I calculate $\texttt{Theta}_{i,t}$ using a separate regression for each industry in each month. This estimation strategy is motivated by the limiting result for θ given in Proposition 3.2. Figures 9a and 9b report these estimates for each industry's key sensitivity parameter, θ. The summary statistics in Table 1 show that these estimated values typically range between

[10]See https://www.osha.gov/pls/imis/sic_manual.html

0.03 and 0.37. And, the average industry as a sensitivity parameter of 0.21.

To make sure that the estimate for each industry's speculator sensitivity only uses data from normal times, I do not compute $\texttt{Theta}_{i,t}$ in any month t where an industry was experiencing a speculative bubble at any point during the previous five years. These omitted values show up as gaps in the time-series plots found in Figures 9a and 9b. This restriction also means that I cannot include 2 of the 21 speculative bubbles in the regression analysis below, so there are a total of 19 speculative bubbles used in the main analysis.

Control Variables. I also use the CRSP and media-coverage data to compute several control variables. To account for variation in industry fundamentals, I also compute the average of an industry's value-weighted dividend-to-price ratio over the previous five years, $\texttt{Div2Price}_{i,t}$. To reflect differences in short-run return volatility across industries, I compute the standard deviation of monthly returns over the previous five years, $\texttt{ReturnVol}_{i,t}$. Finally, to distinguish between the effect of changes in speculator sensitivity and the effect of changes in overall media coverage, I compute the average percent of all articles dedicated to each industry over the previous five years:

$$\texttt{CoverageLevel}_{i,t} \overset{\text{def}}{=} \tfrac{1}{60} \cdot \textstyle\sum_{h=59}^{0} \texttt{\%Articles}_{i,t-h} \tag{36}$$

The model's predictions should operate through $\texttt{Theta}_{i,t}$ and not through $\texttt{CoverageLevel}_{i,t}$.

4.2 Empirical Evidence

Let's now use this data to test the model's main cross-sectional prediction.

Bubble Likelihood. Specifically, consider a probit regression that estimates the relationship between the start of a bubble episode in the ith industry in month $(t+1)$ and the interaction of that same industry's cumulative returns and speculator-persuasiveness sensitivity during the previous five years up to and including month t:

$$\begin{aligned} \mathrm{Pr}_t[\texttt{BubbleStart}_{i,t+1} = 1] = \Phi\Big[\hat{a} + \hat{b} \cdot \texttt{PastReturn}_{i,t} \\ + \hat{c} \cdot \texttt{Theta}_{i,t} \\ + \hat{d} \cdot (\texttt{PastReturn}_{i,t} \times \texttt{Theta}_{i,t}) \\ + \text{``}\ldots\text{''} \\ + \texttt{Resid}_{i,t+1} \Big] \end{aligned} \tag{37}$$

The industry-month data used in this regression includes industry-months where no speculative bubble is taking place as well as the first month of each bubble episode. In other words, each of the 19 speculative bubbles used in the analysis shows up in the regression data once. In the equation above, $\mathrm{Pr}_t[\,\cdot\,]$ denotes the probability conditional on month-t information, $\Phi[\,\cdot\,]$ is the cumulative distribution function (CDF) of the standard normal distribution, the

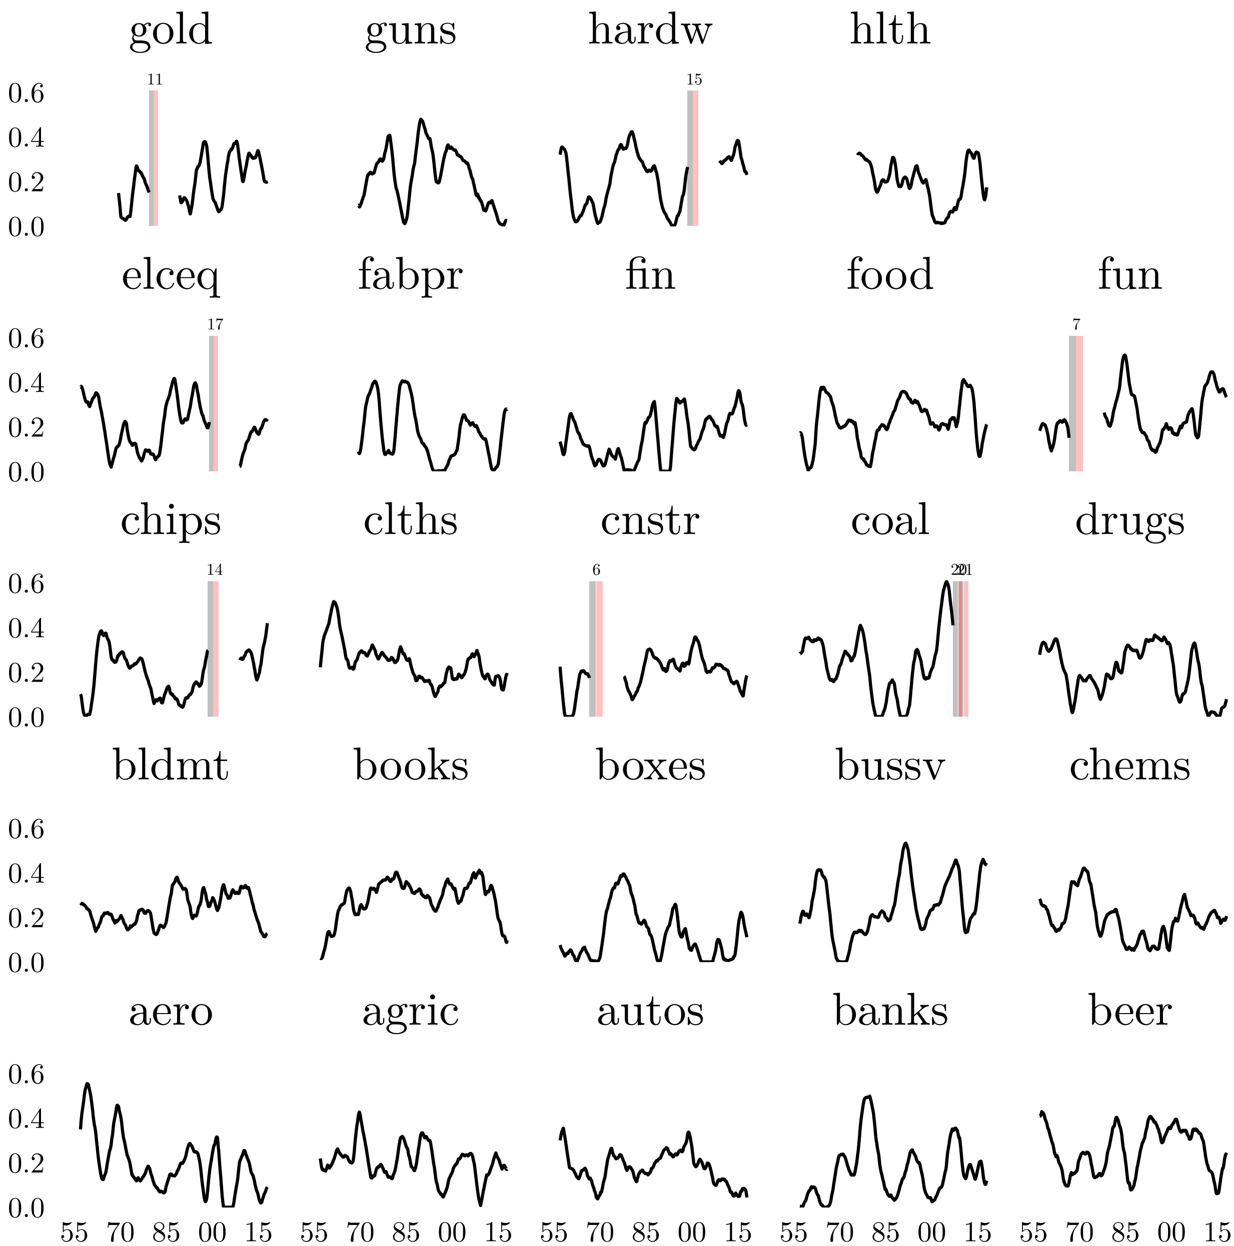

Figure 9a. Estimated Theta. x-axis: time in months from January 1955 to December 2017. y-axis: partial correlation between each industry's return in month t and its media coverage in month t computed over the previous five years. Each panel represents results for a single industry, which are labeled using the abbreviations provided in Fama and French (1997). Vertical grey/red bars: boom-bust episodes. Number above each bar: total number of bubble episodes since January 1950.

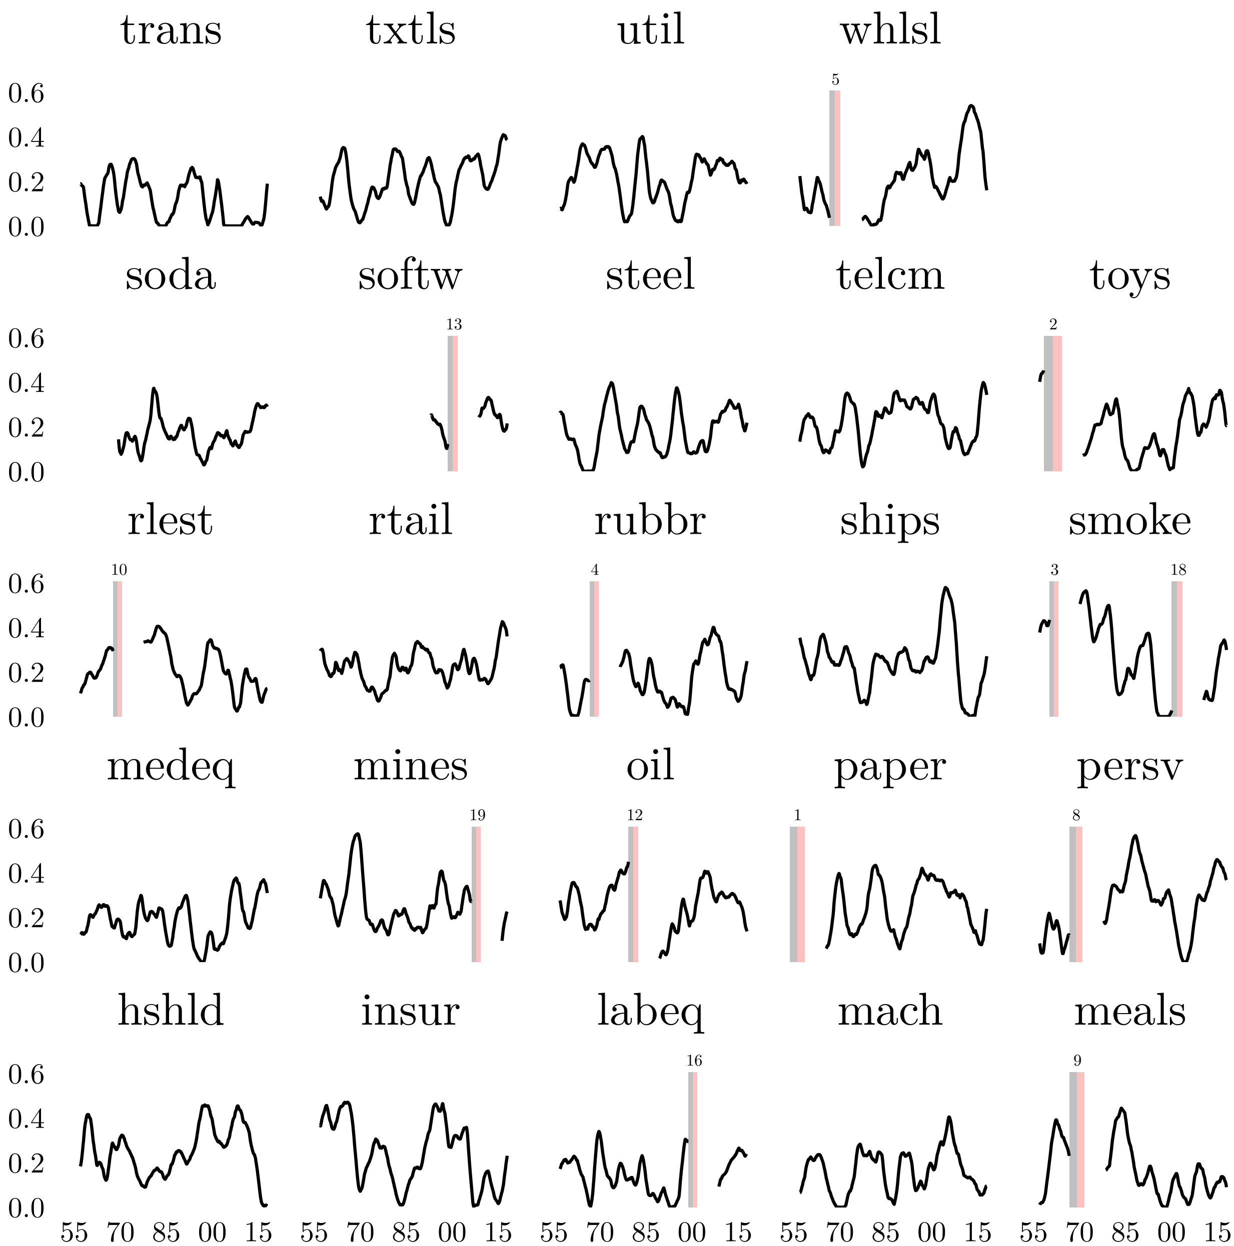

Figure 9b. Estimated Theta, Ctd. x-axis: time in months from January 1955 to December 2017. y-axis: partial correlation between each industry's return in month t and its media coverage in month t computed over the previous five years. Each panel represents results for a single industry, which are labeled using the abbreviations provided in Fama and French (1997). Vertical grey/red bars: boom-bust episodes. Number above each bar: total number of bubble episodes since January 1950.

(a) **Predictive Variables**	Avg	Sd	10%	50%	90%
`PastReturn`	10.12	9.75	−2.26	10.92	21.85
`Div2Price`	22.61	12.64	8.02	20.65	39.59
`ReturnVol`	20.14	5.97	13.43	19.40	27.62
`Theta`	0.21	0.12	0.03	0.21	0.37
`CoverageLevel`	2.13	2.54	0.17	1.31	5.12
		#Observations: 32,915			

(b) **Speculative Bubbles**	Avg	Sd	10%	50%	90%
`BoomLen`	1.83	0.44	1.50	1.58	2.50
`BustLen`	1.71	0.34	1.50	1.50	2.17
`TotalLen`	3.53	0.52	3.00	3.25	4.17
`BoomReturn`	64.13	13.64	50.86	61.41	77.69
`BustReturn`	−49.84	18.72	−28.64	−45.01	−72.13
`TotalReturn`	−18.60	29.60	−52.88	−14.95	15.97
		#Observations: 21			

Table 1. Summary Statistics. Summary statistics for variables used in the econometric analysis. Sample period covers all months from January 1955 to December 2017. **Predictive Variables.** Summary statistics computed across all 32,915 industry-month observations, (i, t), using data from previous five years, $\{t-59, t-58, \ldots, t\}$. `PastReturn`: industry's cumulative return in percent per year. `Div2Price`: industry's average dividend-to-price ratio in percent. `ReturnVol`: standard deviation of an industry's monthly returns in percent per month. `Theta`: partial correlation between each industry's return in month t and its media coverage in month t. `CoverageLevel`: average percent of WSJ articles that mention an industry. **Speculative Bubbles.** Summary statistics computed for the 21 bubble episodes. `BoomLen`: number of years from the start of the bubble to its peak during which the industry earned $\geq 50\%$ per year returns. `BustLen`: number of years after this peak during which the industry lost $\leq -25\%$ per year. `TotalLen`: sum of boom and bust lengths in years. `BoomReturn`: annualized percentage return during period prior to the peak. `BustReturn`: annualized percentage return during period after the peak. `TotalReturn`: annualized percentage return on a $1 investment over the course of an entire bubble episode.

terms with a hat represent estimated coefficient values, the "$\cdots$" term is a stand-in term for various control variables, and $\texttt{Resid}_{i,t+1}$ denotes the regression residual.

Table 2 reports the results of this baseline regression specification. Columns (1)-(3) show estimated coefficient values without accounting for the sensitivity of each industry's speculators to past returns, `Theta`—i.e., ignoring the coefficients $\hat{c}$ and $\hat{d}$. These first three columns that speculative bubbles tend to follow on the heels of good performance. "Nearly all bubbles from tulips to South Sea to the 1929 U.S. stock market to the late 1990s internet occur on the back of good fundamental news." (Barberis et al., 2018) The coefficient estimate $\hat{b}$ is positive and statistically significant. However, Column (1) implies that an industry in the top `PastReturn` decile in month t only has a 0.10% chance of entering into a speculative bubble in month $(t+1)$. In other words, "only a relatively small proportion of large shocks lead to a speculative mania" (Kindleberger, 1978).

Columns (4)-(6) then re-estimate the specification in Equation (37) after including both `Theta` and `Theta` $\times$ `PastReturn`. These second three columns show that the positive relationship between `PastReturn` and `BubbleStart` was actually operating entirely through industries with high speculator sensitivities—i.e., through industries with high values of `Theta`. The coefficient estimate $\hat{d}$ is positive and highly significant, both statistically and economically. Take that same industry, which was in the top `PastReturn` decile. If that industry were in the bottom `Theta` decile in month t, then Column (4) implies that the industry would only have a 0.03% chance of entering into a speculative bubble in month $(t+1)$. By contrast, if the industry were in the top `Theta` decile, then it would have a 0.18% chance of entering into a speculative bubble in month $(t+1)$. What's more, after controlling for the effect of the sensitivity of speculators to fluctuations in past returns, the level of an industry's past returns is actually negatively related to the likelihood of a speculative bubble.

In short, the results in Table 2 support the cross-sectional predictions developed in Section 3. They suggest that an industry is more likely to experience a speculative bubble when it has both high past returns and speculators who are very sensitive to changes in past returns. Either one of these ingredients is not enough on its own. And, the fact that the coefficients on the dividend-to-price ratio and return volatility are insignificant suggests that this cross-sectional pattern is not the result of a confounding difference in the level or volatility of fundamental values across industries.

Sensitivity vs. Level. But, how do we know that this cross-sectional pattern is really being driven by the sensitivity of speculators to fluctuations in past returns? I'm using a measure of media coverage to estimate this sensitivity. So, perhaps the level of media coverage is really driving the result? It seems plausible that an industry with a higher absolute level of media coverage—i.e., an industry mentioned in a larger fraction of all WSJ articles—might be more

Dependent Variable: `BubbleStart`	(1)	(2)	(3)	(4)	(5)	(6)
`Intercept`	$-3.47^{\star\star\star}$ (0.13)	$-3.34^{\star\star\star}$ (0.17)	$-3.71^{\star\star\star}$ (0.34)	$-3.06^{\star\star\star}$ (0.18)	$-2.94^{\star\star\star}$ (0.21)	$-3.20^{\star\star\star}$ (0.40)
`PastReturn`	$1.71^{\star\star\star}$ (0.76)	$1.75^{\star\star\star}$ (0.75)	$1.85^{\star\star\star}$ (0.71)	$-2.02^{\star}$ (1.22)	$-1.99^{\star}$ (1.21)	1.72 (1.23)
`Div2Price`		-0.58 (0.56)	-0.35 (0.58)		-0.64 (0.58)	-0.50 (0.61)
`ReturnVol`			5.21 (4.02)			3.30 (4.21)
`Theta`				$-2.03^{\star\star}$ (0.95)	$-2.00^{\star\star}$ (0.95)	$-1.86^{\star\star}$ (0.96)
`PastReturn` × `Theta`				$16.43^{\star\star\star}$ (4.74)	$16.57^{\star\star\star}$ (4.78)	$15.64^{\star\star\star}$ (4.86)
#Observations	32,915					

Table 2. Bubble Likelihood. Results from a probit regression using data at the industry-month level from January 1955 to December 2017. All months following the start of a speculative bubble are removed from the data. `BubbleStart`: indicator that is one if a speculative bubble starts in industry i in month $(t+1)$. Right-hand-side variables for industry-month (i,t) computed using data from previous five years, $\{t-59, t-58, \ldots, t\}$. `PastReturn`: industry's cumulative return. `Div2Price`: industry's average dividend-to-price ratio. `ReturnVol`: standard deviation of an industry's monthly returns. `Theta`: partial correlation between an industry's monthly returns and its media coverage. Numbers in parentheses are standard errors. Statistical significance: $^{\star} = 10\%$, $^{\star\star} = 5\%$, and $^{\star\star\star} = 1\%$.

likely to experience a speculative bubble following good past performance.

The results in Table 3 show that this alternative hypothesis is not supported by the data. Columns (1)-(3) replicate the analysis in Columns (4)-(6) of Table 2 using `CoverageLevel` rather than `Theta`. There is only a weak relationship between the average media coverage given to an industry in the previous five years and its likelihood of entering into a speculative bubble. And, if such a relationship does exist, then it's likely to go the wrong direction. Industries that have had more media coverage in previous years are less likely to boom. Columns (4)-(6) of Table 3 then show the results of re-estimating this regression specification after including both `CoverageLevel` and `Theta`. These second three columns show that `Theta`'s effect on the likelihood of a speculative bubble is unchanged. The results are driven by the sensitivity of speculators to fluctuations in past returns rather than the absolute level of attention given to an industry.

Matching Analysis. A second concern that you might have about the specification in Equation (37) is that it's comparing apples to oranges. It's comparing industry-month ob-

Dependent Variable: `BubbleStart`	(1)	(2)	(3)	(4)	(5)	(6)
`Intercept`	$-3.71^{\star\star\star}$ (0.34)	$-3.55^{\star\star\star}$ (0.35)	$-3.61^{\star\star\star}$ (0.35)	$-3.20^{\star\star\star}$ (0.40)	$-3.02^{\star\star\star}$ (0.41)	$-3.08^{\star\star\star}$ (0.41)
`PastReturn`	$1.85^{\star\star\star}$ (0.71)	$1.80^{\star\star\star}$ (0.71)	$2.37^{\star\star\star}$ (0.88)	-1.72 (1.72)	-1.82 (1.82)	-1.34 (1.34)
`Div2Price`	-0.35 (0.58)	-0.32 (0.57)	-0.32 (0.57)	-0.50 (0.61)	-0.43 (0.60)	-0.40 (0.60)
`ReturnVol`	5.21 (4.02)	4.70 (4.10)	4.51 (4.13)	3.30 (4.21)	2.74 (4.28)	2.64 (4.39)
`CoverageLevel`		$-8.73^{\star}$ (5.26)	-3.46 (5.69)		$-9.20^{\star}$ (5.47)	-4.66 (5.51)
`PastReturn` × `CoverageLevel`			-46.99 (39.21)			-41.09 (39.86)
`Theta`				$-1.86^{\star\star}$ (0.96)	$-1.94^{\star\star}$ (0.97)	$-1.94^{\star\star}$ (0.97)
`PastReturn` × `Theta`				$15.64^{\star\star\star}$ (4.86)	$16.01^{\star\star\star}$ (4.93)	$16.09^{\star\star\star}$ (4.98)
#Observations			32,915			

Table 3. Sensitivity vs. Level. Results from a probit regression using data at the industry-month level from January 1955 to December 2017. All months following the start of a speculative bubble are removed from the data. `BubbleStart`: indicator that is one if a speculative bubble starts in industry i in month $(t+1)$. Right-hand-side variables for industry-month (i,t) computed using data from previous five years, $\{t-59, t-58, \ldots, t\}$. `PastReturn`: industry's cumulative return. `Div2Price`: industry's average dividend-to-price ratio. `ReturnVol`: standard deviation of an industry's monthly returns. `Theta`: partial correlation between an industry's monthly returns and its media coverage. `CoverageLevel` is the average percent of WSJ articles that mention an industry. Numbers in parentheses are standard errors. Statistical significance: * = 10%, ** = 5%, and *** = 1%.

servations that are very likely to be followed by a speculative bubble with industry-month observations that are very unlikely to be followed by a speculative bubble. And, there are a lot more observations in this latter group. So, to address this concern, I re-estimate the regression specification in Equation (37) on a matched sample of industry-month observations. The first half of this dataset consists of the 19 industry-month observations that are followed by a speculative bubble. The second half consists of a matched sample of 19 industry-month observations that were not followed by a speculative bubble but had similar past returns, dividend-to-price ratios, and levels of return volatility.

Table 4 shows the results of this new estimation strategy. Columns (1)-(4) show that `PastReturn`, `Div2Price`, and `ReturnVol` have to ability to predict the start of speculative

Dependent Variable: `BubbleStart`	(1)	(2)	(3)	(4)	(5)	(6)
`Intercept`	-0.11 (0.39)	0.14 (0.40)	-0.45 (0.82)	-0.43 (1.24)	1.22 (0.78)	1.21 (1.55)
`PastReturn`	0.73 (2.25)			0.58 (2.30)	$-12.88^{\star\star}$ (6.66)	$-12.90^{\star\star}$ (6.68)
`Div2Price`		-0.66 (1.64)		-0.19 (1.91)		0.21 (2.18)
`ReturnVol`			7.43 (13.08)	6.35 (15.12)		-0.33 (16.54)
`Theta`					$-7.30^{\star\star}$ (3.79)	$-7.47^{\star}$ (4.23)
`PastReturn` $\times$ `Theta`					$60.46^{\star\star}$ (27.13)	$61.13^{\star\star}$ (28.29)
#Observations				38		

Table 4. Matching Analysis. Results from a probit regression using data on the 19 industry-month observations where a bubble started paired with a matched sample of otherwise identical industry-month observation where no speculative bubble began. Sample period covers January 1955 to December 2017. Observations are matched based on past returns, dividend-to-price ratio, and return volatility. `BubbleStart`: indicator that is one if a speculative bubble starts in industry i in month $(t+1)$—i.e., for exactly half of the observations. Right-hand-side variables for industry-month (i, t) computed using data from previous five years, $\{t-59, t-58, \ldots, t\}$. `PastReturn`: industry's cumulative return. `Div2Price`: industry's average dividend-to-price ratio. `ReturnVol`: standard deviation of an industry's monthly returns. `Theta`: partial correlation between an industry's monthly returns and its media coverage. Numbers in parentheses are standard errors. Statistical significance: * = 10%, ** = 5%, and *** = 1%.

bubbles in this new dataset. This is as expected. It confirms that the matching procedure really found industry-month observations that look similar along these three dimensions. Columns (5)-(6) then introduce `Theta` to the specification just like before. These columns show that, if you look specifically at the industry-month observations most likely to precede a speculative bubble, then the effects of speculators' sensitivity to fluctuations in past returns is even stronger. Column (5) implies that the marginal effect of moving an industry month from the lowest `Theta` decile to the highest `Theta` decile is an increase in the probability of a speculative bubble in the following month of 26.52%.

4.3 Out-of-Sample Predictions

What makes the model of displacement events presented in Section 2 different is that it can be used to make out-of-sample predictions about the likelihood of a speculative bubble. This is not something that you can do within the existing limits-to-arbitrage framework. And, I conclude the econometric analysis by examining these out-of-sample predictions.

Subsample Analysis. For the specification in Equation (37) to make useful out-of-sample predictions, it has to be the case that the coefficient estimates for `Theta` and `Theta` $\times$ `PastReturn` are relatively constant over time. The model in Section 2 is written down under the assumption that the effect of changing θ is constant. The maintained hypothesis is that, while the underlying forces that determine the sensitivity of each asset's speculators to fluctuations in past returns might change over time, the effect of increasing θ by 0.01 on the probability of realizing a speculative bubble is the same.

Table 5 shows the results of estimating the regression specification in Equation (37) separately for the first and second halves of the dataset used in Table 2. Columns (1)-(3) report results using data from January 1955 to December 1987; whereas, Columns (4)-(6) report results using data from January 1988 to December 2017. Comparing the coefficient estimates on `Theta` and `Theta` $\times$ `PastReturn`, we see that the effect of increasing speculators' sensitivity to fluctuations in past returns is indeed relatively constant over time.

Out-of-Sample Predictions. Figures 10a and 10b then plot the model-implied probability of each industry entering into a speculative bubble during the second half of the sample based on the coefficient estimates from the first half of the sample. These figures show that, even though the telecommunications and personal-services industries did not technically experience a $\geq$ 50%-per-year speculative bubble in the late 1990s, they were very likely to have reached this threshold. In other words, a policymaker in January 1997 could have used the point estimated from Columns (1)-(3) in Table 5 to learn about the likelihood of a speculative bubble occurring in one these industries during the following year. This is an exercise that is of great practical value.

5 Conclusion

By the logic of limits to arbitrage, any bias-constraint pair from these two lists might potentially combine at any moment to produce an equilibrium pricing error. But, large pricing errors such as speculative bubbles are rare. Why is this? What pins down the likelihood of a speculative bubble? How often should we expect that some psychological bias will cause some trading constraint to bind?

You can't answer these questions within the existing limits-to-arbitrage framework. The limits of arbitrage explain how an equilibrium pricing error can be sustained in equilibrium. But, we're not asking questions about *how* a speculative bubble can be sustained; we're asking questions about *how often* we should expect one to occur. In short, we need to introduce another ingredient—an on/off switch—something that sporadically amplifies the effect of

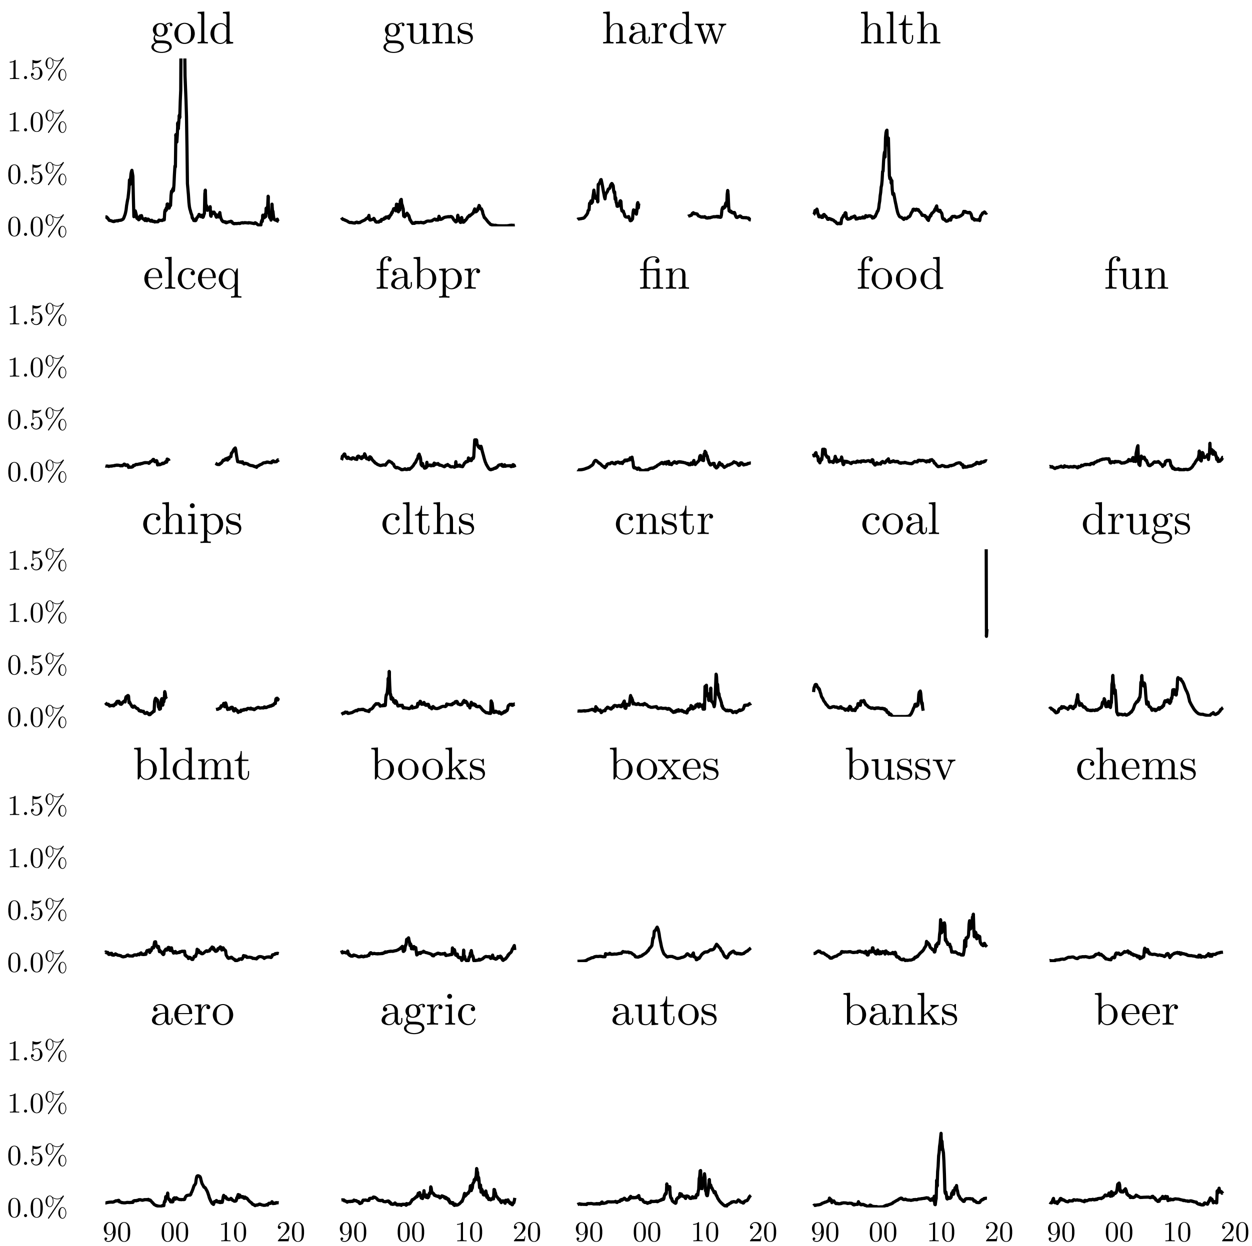

Figure 10a. Out-of-Sample Predictions. x-axis: time in months from January 1988 to December 2017. y-axis: predicted probability that an industry will begin a speculative bubble in month $(t+1)$ based on the regression specification in Equation (37) using parameter values estimated from January 1955 to December 1987. Each panel represents results for a single industry, which are labeled using the abbreviations provided in Fama and French (1997).

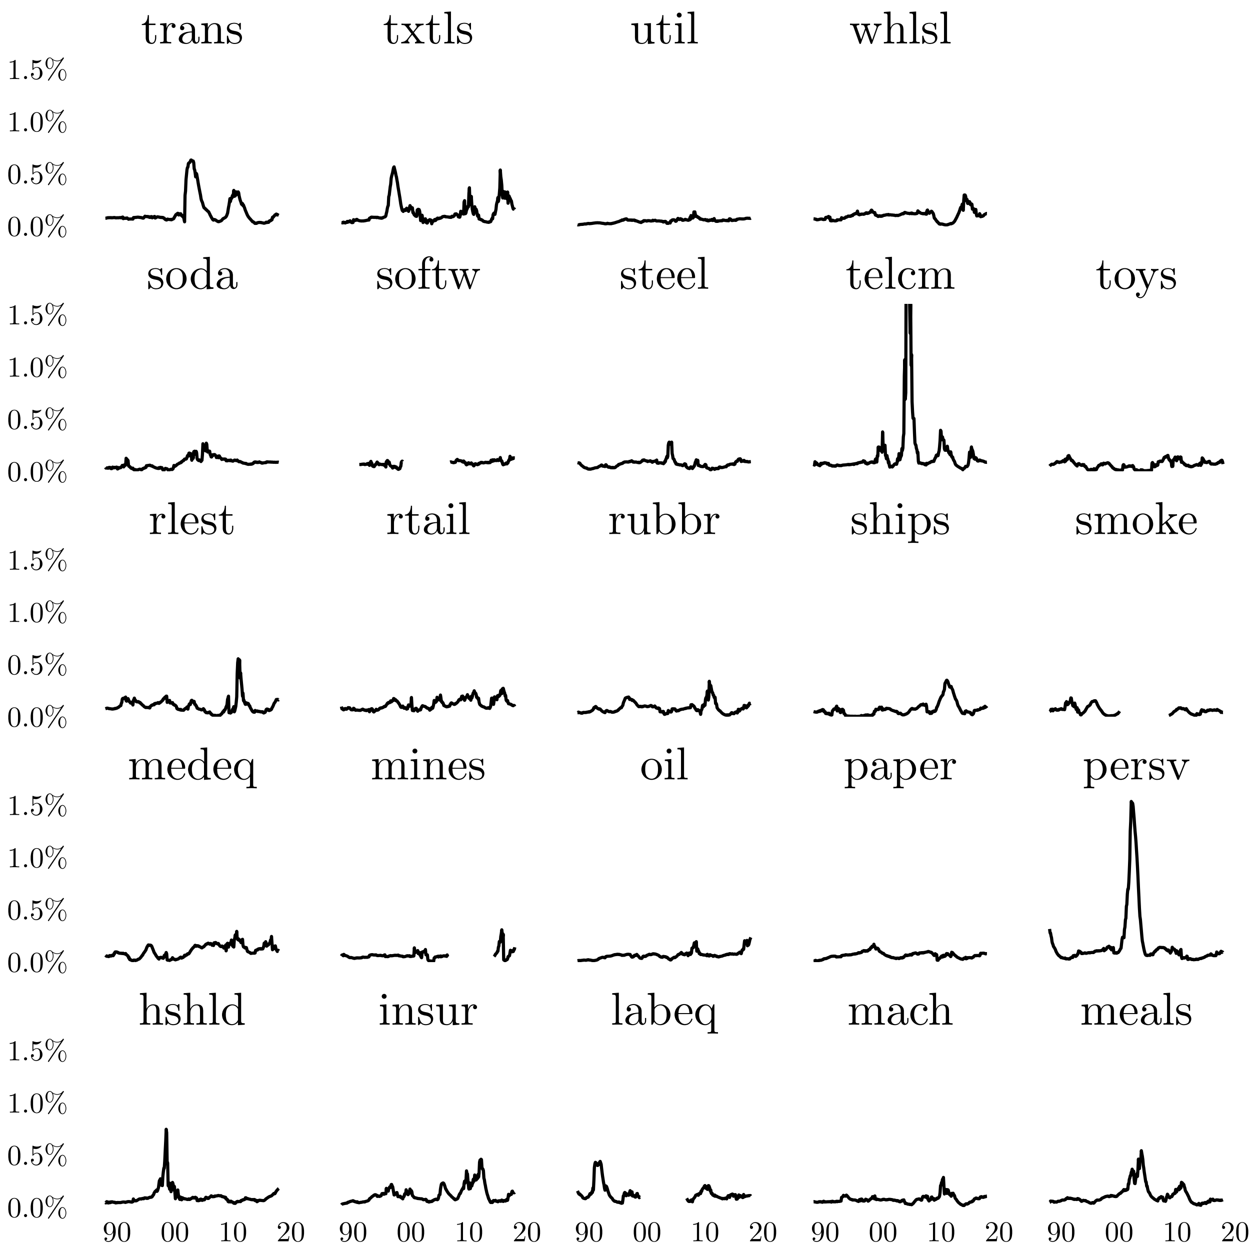

Figure 10b. Out-of-Sample Predictions, Ctd. x-axis: time in months from January 1988 to December 2017. y-axis: predicted probability that an industry will begin a speculative bubble in month $(t+1)$ based on the regression specification in Equation (37) using parameter values estimated from January 1955 to December 1987. Each panel represents results for a single industry, which are labeled using the abbreviations provided in Fama and French (1997).

Dependent Variable: `BubbleStart`						
		1955-1987			1988-2017	
	(1)	(2)	(3)	(4)	(5)	(6)
`Intercept`	$-2.83^{\star\star\star}$ (0.21)	$-2.69^{\star\star\star}$ (0.30)	$-2.36^{\star\star\star}$ (0.64)	$-3.56^{\star\star\star}$ (0.54)	$-3.33^{\star\star\star}$ (0.56)	$-3.94^{\star\star\star}$ (0.79)
`PastReturn`	$-4.25^{\star\star\star}$ (1.62)	$-4.06^{\star\star\star}$ (1.64)	$-4.37^{\star\star\star}$ (1.77)	0.87 (2.94)	0.94 (2.96)	1.58 (3.04)
`Div2Price`		-0.60 (0.86)	-0.85 (0.98)		$-1.67^{\star\star}$ (1.54)	$-1.33^{\star}$ (1.50)
`ReturnVol`			-4.12 (7.27)			7.56 (6.31)
`Theta`	-1.70 (1.07)	-1.64 (1.08)	-1.68 (1.09)	-3.36 (2.45)	-3.29 (2.46)	-2.62 (2.51)
`PastReturn` × `Theta`	$16.29^{\star\star\star}$ (5.95)	$16.03^{\star\star\star}$ (5.99)	$16.72^{\star\star\star}$ (6.21)	$22.22^{\star}$ (11.59)	$21.70^{\star}$ (11.62)	17.97 (11.99)
#Observations		16,476			16,439	

Table 5. Subsample Analysis. Results from a probit regression using data at the industry-month level from January 1955 to December 2017. All months following the start of a speculative bubble are removed from the data. Data has been split into two subsamples: January 1955 to December 1987 vs. January 1988 to December 2017. `BubbleStart`: indicator that is one if a speculative bubble starts in industry i in month $(t+1)$. Right-hand-side variables for industry-month (i,t) computed using data from previous five years, $\{t-59, t-58, \ldots, t\}$. `PastReturn`: industry's cumulative return. `Div2Price`: industry's average dividend-to-price ratio. `ReturnVol`: standard deviation of industry's monthly returns. `Theta`: partial correlation between an industry's monthly returns and its media coverage. Numbers in parentheses are standard errors. Statistical significance: $\star = 10\%$, $\star\star = 5\%$, and $\star\star\star = 1\%$.

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 by recombining two common elements found in popular accounts of bubble formation in a new way: i) while speculators get overly excited following good news about fundamentals due to the "madness of crowds", ii) they only recover their senses "slowly and one by one" (Mackay, 1841).

I verify the model's main empirical prediction in monthly U.S. industry returns. And, after fitting the model using the first half of the data sample, I show that the model's out-of-sample predictions for the second half of the data sample provide a measure of "market froth" in an industry—i.e., the likelihood that the industry will experience a speculative bubble in the future. This bubble likelihood is something that gets talked about a lot, but it's something that previously had no analogue in economic models.

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*.

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 (2016). Social networks and housing markets. NBER.

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*.

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 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*.

Deutsch, D. (2011). *The beginning of infinity: Explanations that transform the world*. Penguin UK.

Duflo, 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 (2015). Financial markets where traders neglect the informational content of prices. NBER.

Fama, E. (1965). The behavior of stock-market prices. *The 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 trades 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. and D. Hirshleifer (2015). Self-enhancing transmission bias and active investing. SSRN.

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*.

Lux, T. (1995). Herd behaviour, bubbles and crashes. *Economic Journal*.

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 event, 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 modelling*.

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 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. L. (1886). *The strange case of Dr. Jekyll and Mr. Hyde*. New York: Scribner.

Strogatz, S. (2014). *Nonlinear dynamics and chaos: with 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.-F. (1845). Recherches mathématiques sur la loi d'accroissement de la population. *Proceedings of the Royal Academy of Brussels*.

Xiong, W. (2013). Bubbles, crises, and heterogeneous beliefs. NBER.

Xiong, W. and Y. Yu, Jialin (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 $\mathrm{G}(n, \theta, r) = \theta \cdot r \cdot (1-n) \times n - n$. We can rewrite this law of motion as

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

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
$$\tfrac{\partial}{\partial n}[\mathrm{G}(n, \theta, r)]_{n=0,\, r<1/\theta} = \theta \cdot (r - 1/\theta) - 2 \cdot \theta \cdot r \cdot 0 < 0$$
2. If $r > 1/\theta$, then $(\theta \cdot r - 1) > 0$. So, there are now two ways for the right-hand size 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
$$\tfrac{\partial}{\partial n}[\mathrm{G}(n, \theta, r)]_{n=0,\, r>1/\theta} = \theta \cdot (r - 1/\theta) - 2 \cdot \theta \cdot r \times 0 > 0$$
$$\tfrac{\partial}{\partial n}[\mathrm{G}(n, \theta, r)]_{n=(r-1/\theta)/r,\, r>1/\theta} = \theta \cdot (r - 1/\theta) - 2 \cdot \theta \cdot r \times (r - 1/\theta)/r < 0$$

□

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

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

Market clearing then implies that:

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

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 $r_{t-2} < r_\star$ can be written as

$$\mathrm{E}_{t-2}[\,\mathrm{B}(\theta, r_{t-1}) \,|\, \mathrm{B}(\theta, r_{t-2}) = 0\,] = \Pr_{t-2}[\, r_{t-1} > r_\star \,|\, r_{t-2} < r_\star\,]$$

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

$$\mathrm{E}_{t-2}[\,\Delta v_{t-1}\,] = \kappa_v \cdot (\mu_v - v_{t-2})$$
$$\mathrm{Var}_{t-2}[\,\Delta v_{t-1}\,] = \sigma_v^2$$

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

$$\mathrm{F}_{t-2}(p) = \tfrac{1}{\sigma_v \cdot \sqrt{2 \cdot \pi}} \cdot e^{-\frac{1}{2 \cdot \sigma_v^2} \cdot (p - v_{t-2} - \mathrm{E}_{t-2}[\Delta v_{t-1}] + \gamma \cdot \psi)^2} \tag{A1}$$

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

$$\mathrm{E}_{t-2}[\,\mathrm{B}(\theta, r_{t-1}) \,|\, r_{t-2} < r_\star\,] = \int_{p_{t-2}/\theta}^{\infty} \mathrm{F}_{t-2}(p) \cdot \mathrm{d}p \tag{A2}$$

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

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, λ, 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:

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

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:

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

In the equation above, $\mathrm{O}[n_0]$ represents *'Big O'* notation. We say that $\mathrm{f}(x) = \mathrm{O}[x]$ as $x \to 0$ if there exists a positive constant $C > 0$ such that $|\mathrm{f}(x)| \leq C \cdot x$ 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 $\mathrm{G}(n_0, \theta, r)$:

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

$\square$

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 ϵ and the half-life of small excited-speculator populations when $r < r_\star$:

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

The functional form for $\tau_{1/2}(\theta, r)$ given by Equation (27) implies that for small fluctuations in returns we have that both $\partial_{\theta \cdot r}[\tau_{1/2}] = \log(2) \cdot (1 - \theta \cdot r)^{-2}$ and

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

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

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

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

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

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

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 ω, 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 $\mathrm{G}(0, \theta, r) = 0$.*
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 $\widetilde{\Theta}(n, \theta, r) = r \cdot (1 - e^{-\theta \cdot n})$. 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 $\widetilde{\mathrm{G}}(n, \theta, r)$. If there exists some $r_- > 0$ such that $\frac{\partial}{\partial n}[\widetilde{\mathrm{G}}(n, \theta, r_-)]_{n=0} < 0$ 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, $\frac{\partial}{\partial n}[\mathrm{G}(n, \theta, r_-)]_{n=0} < 0$. So, via the implicit-value theorem, we know that there must exist some critical return threshold, $r_\star > r_-$, such that

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

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 $\epsilon, \delta > 0$:

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

Clearly, for a dynamical system with this functional form, $n = 0$ is a steady-state solution for all $r > 0$ since $\mathrm{G}(0, \theta, r) = \varphi \cdot (r - r_\star) \cdot 0 - \chi \cdot 0^2 = 0$. 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 $\mathrm{G}(n, \theta, r)$:

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

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

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

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

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

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

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

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

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

Furthermore, this positive solution will only be stable if

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

Plugging in the functional form for n yields:

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

Thus, we can conclude that the solution is stable for $r > r_\star$ since $\frac{\partial^2}{\partial r \partial n}[\mathrm{G}(0, \theta, r_\star)] > 0$. 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 $r_\star$ 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 $r_\star = 1/\theta$ disappears. This subsection shows that adding noise does not eliminate the sharp change in population dynamics around $r_\star$.

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

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

In the equation above, $\sigma_g > 0$ is a positive constant reflecting the instantaneous volatility of the excited-speculator population growth rate, and $\varepsilon_g \sim \mathrm{N}(0, 1)$ 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:

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

I've included time subscripts in the above equations, n_τ and $\mathrm{d}\varepsilon_{g,\tau}$ rather than n and $\mathrm{d}\varepsilon_g$, 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 $\sigma_g \cdot n_\tau \cdot \mathrm{d}\xi_{g,\tau}$ contains n_τ, 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 in the market. While it's possible to use N_τ/U_τ 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, $r_\star = 1/\theta$. 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 sudden qualitative change as shown in Figure 11.

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* $r > r_\star = 1/\theta$, *then the stationary distribution for* $\lim_{\tau \to \infty} n_\tau = n_\infty$ *is characterized by the following probability-density function (PDF) given any initial* $n_0 \in (0, \infty)$
$$n_\infty(\theta, r) \sim \mathrm{Ga}\big(\theta \cdot r - 1, \theta \cdot r\big)$$
where $\mathrm{Ga}(a, b) \stackrel{\text{def}}{=} \frac{b^a}{\Gamma(a)} \cdot \frac{x^{a-1}}{e^{b \cdot x}}$ *is the PDF for the Gamma distribution.*
2. *If* $r < r_\star = 1/\theta$, *then* $n_\infty(\theta, r) = 0$ *almost surely.*

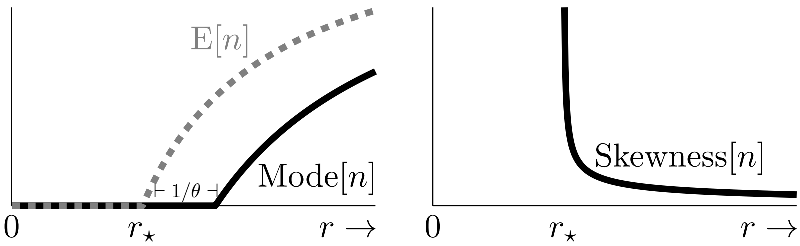

Figure 11. Stationary Distribution. Statistical properties of the stationary distribution for the excited-speculator population as characterized in Proposition B.2 where $r_\star = 1/\theta$.

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

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

where $\tau \geq 0$, $\mathrm{m}(y)$ denotes the drift term, $\sigma_y > 0$ is a positive constant, $\mathrm{s}(y) > 0$ is the diffusion term, and $\mathrm{d}\varepsilon_{y,\tau}$ is a standard Brownian-motion process. Assume that $\mathrm{s}(0) = 0$ and that $\mathrm{m}(\infty) = -\infty$. Finally, let $\mathrm{F}_y(\tau, y_0)$ 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:

$$\begin{aligned} \tfrac{\partial}{\partial \tau}[\mathrm{F}_y(\tau, y_0)] = &- \tfrac{\partial}{\partial y}\Big[\,(\mathrm{m}(y) + \tfrac{1}{2} \cdot \sigma_y^2 \cdot \tfrac{\mathrm{d}}{\mathrm{d}y}[\mathrm{s}(y)] \cdot \mathrm{s}(y)) \cdot \mathrm{F}_y(\tau, y_0)\,\Big] \\ &+ \tfrac{1}{2} \cdot \sigma_y^2 \cdot \tfrac{\partial^2}{\partial y^2}\Big[\,\mathrm{s}(y)^2 \cdot \mathrm{F}_y(\tau, y_0)\,\Big] \end{aligned}$$

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

$$\tfrac{\partial}{\partial y}\Big[\,(\mathrm{m}(y) + \tfrac{1}{2} \cdot \sigma_y^2 \cdot \tfrac{\mathrm{d}}{\mathrm{d}y}[\mathrm{s}(y)] \cdot \mathrm{s}(y)) \cdot \mathrm{F}_y(\tau, y_0)\,\Big] = \tfrac{1}{2} \cdot \sigma_y^2 \cdot \tfrac{\partial^2}{\partial y^2}\Big[\,\mathrm{s}(y)^2 \cdot \mathrm{F}_y(\tau, y_0)\,\Big]$$

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

$$\mathrm{F}_y(\cdot) = \tfrac{1}{K \cdot \mathrm{s}(y)} \cdot \exp\Big(\, \tfrac{1}{\sigma_y^2/2} \cdot \textstyle\int_0^y \frac{\mathrm{m}(y')}{\mathrm{s}(y')^2} \cdot \mathrm{d}y' \,\Big)$$

$$\text{given} \qquad K = \int_0^\infty \tfrac{1}{\mathrm{s}(y)} \cdot \exp\Big(\, \tfrac{1}{\sigma_y^2/2} \cdot \textstyle\int_0^y \frac{\mathrm{m}(y')}{\mathrm{s}(y')^2} \cdot \mathrm{d}y' \,\Big) \cdot \mathrm{d}y < \infty$$

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

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

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

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

And, this is the functional form of the PDF for the Gamma distribution, $\mathrm{Ga}(a, b) \overset{\text{def}}{=} \frac{b^a}{\Gamma(a)} \cdot \frac{x^{a-1}}{e^{b \cdot x}}$ with $a = \theta \cdot r - 1$ and $b = \theta \cdot r$. This distribution is defined for all $x \in (0, \infty)$. When $r < r_\star$ 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 $\epsilon \in [0, \frac{1}{\theta \cdot r})$:

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

The $(1 + \epsilon \cdot n)$ term in the equation above captures the idea that an inflow of excited speculators at time τ 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 ϵ, 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 we re-written as follows

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

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 $\bar{n} < 1$. $\square$