# The Ex Ante Likelihood of Bubbles

# Alex Chinco<sup>a</sup>

<sup>a</sup>Zicklin School of Business, Baruch College, New York, New York 10010 Contact: alexchinco@gmail.com, bhttps://orcid.org/0000-002-0255-1688 (AC)

Received: August 28, 2020 Revised: August 9, 2021 Accepted: October 3, 2021 Published Online in Articles in Advance:

https://doi.org/10.1287/mnsc.2022.4351

Copyright: © 2022 INFORMS

**Abstract.** The limits of arbitrage explain how a speculative bubble is sustained; they do not explain how likely one is to occur. To do that, you need a theory about the thing that sporadically causes arbitrageur constraints to bind. I propose a first such theory, which is based on social interactions between speculators. The theory says that bubbles should be more likely in assets where increases in past returns make excited-speculators relatively more persuasive to their peers. I empirically verify this ex ante prediction about bubble likelihoods and show that it is robust to some ex post disagreement about bubble definitions.

History: Accepted by Victoria Ivashina, finance.
Supplemental Material: The data files and online appendix are available at https://doi.org/10.1287/mnsc. 2022.4351.

Keywords: limits to arbitrage • speculative bubbles • social interactions

# 1. Introduction

Suppose there has recently been some good news about an asset. As a result, the asset's price has increased a bit. Given everything else you know about the asset, how likely is it that this modest price increase will morph into a full-blown speculative bubble (i.e., a large additional price boom driven by excited-speculators who overvalue the asset followed by a crash when arbitrageur constraints are finally relaxed)?

Investors and policymakers really want to be able to answer this sort of question. When investors ask things like "by driving rates ever lower, is the Fed helping to fuel dangerous bubbles?," they are asking about the likelihood of a future bubble.<sup>1</sup> Likewise, when Alan Greenspan expressed his concerns in 2005 about "froth in the housing market," he was expressing concerns about the exact same thing.<sup>2</sup> Although researchers tend to focus on predicting when a bubble will pop, for Greenspan the timing of the crash was beside the point. The ensuing housing collapse would have been just as bad for U.S. investors had prices peaked in 2007 or 2008 rather than 2006.

Popular accounts of bubble formation are unable to answer questions about when and where a speculative bubble is most likely to occur next. In fact, most authors openly acknowledge that this is a soft spot in their narratives. For example, Minsky (1970) admits that "an event that is not of unusual size or duration can trigger a sharp financial reaction." And, Shiller (2000) calls this fact "unsatisfying to those of us seeking scientific certitude." Kindleberger (1978) does tell us that bubbles tend to occur on the back of good news,  $\Pr[GoodNews_t | Bubble_{t+1}] \approx 1$ . But, this is an observation about the probability of good news given a bubble, not the probability of a bubble given good news. Good news is not a reliable predictor of future bubbles,  $\Pr[Bubble_{t+1} | GoodNews_t] \ll 1$ . "Virtually every mania is associated with a robust economic expansion, but only a few economic expansions are associated with a mania" (Kindleberger 1978).

Unfortunately, the existing academic literature cannot speak to the ex ante likelihood of bubbles either. The standard approach to modeling speculative bubbles goes by the name of "limits to arbitrage" (Shleifer and Vishny 1997). This recipe calls for equal parts bias and constraint. The bias-for example, overconfidence (Daniel et al. 1998, Scheinkman and Xiong 2003) or extrapolation (Hong and Stein 1999, Barberis et al. 2015)-causes speculators to overvalue an asset. The constraint—for example, a short-sale ban (Miller 1977) or margin requirement (Gromb and Vayanos 2002)prevents arbitrageurs from fixing the problem. These two ingredients produce a speculative boom. The crash occurs when the constraint is finally lifted, allowing arbitrageurs to trade against speculators' excess demand.

Although this limits-to-arbitrage framework does give us conditions under which a bubble can exist, it says nothing about when and where these conditions are most likely to be satisfied in the future. In fact, most papers in this literature sidestep the issue of ex ante likelihoods entirely. They start by simply assuming that a bubble has already begun. "We assume that the price surpasses the fundamental value at a random point in time  $t_0''$  (Abreu and Brunnermeier 2003). Each paper then studies how the specific bias or constraint responsible for the bubble will impact the rest of its life span. How much will prices rise during the boom? When will arbitrageurs pop the bubble? How far will prices fall during the crash?

To fill this gap in our understanding, I develop a first simple theory explaining the ex ante likelihood of bubbles. This requires a story for why, "on occasion and under some circumstances (Barberis 2018)," a speculator bias will cause some arbitrageur constraint to bind and a bubble to form. And, to be empirically useful, this bubble-generating mechanism must have two key properties.

First, the limits-to-arbitrage literature has documented many different speculator biases and arbitrageurs constraints, and different bias-constraint pairs may be responsible for different bubble episodes. So the bubble-generating mechanism I propose must not be specific to one particular bias-constraint pair. It must work the same way no matter which pair happens to bind in any given bubble episode.<sup>3</sup>

Second, the key parameter governing the ex ante likelihood of a bubble must be relatively stable over time both during and between bubble episodes. That way, a researcher can estimate this parameter during normal times and use it to predict the likelihood of a future bubble. A theory in which bubbles just happen because the key parameter suddenly jumps would explain nothing. It would take us from being unable to predict future bubbles to being unable to predict future parameter changes.

The bubble-generating mechanism I propose is based on the observation that speculators "go mad in herds [but] only recover their senses slowly and one by one" (Mackay 1841). This is a statement about how social interactions govern the dynamics of an excited-speculator population. These social interactions usually make the population disappear. But, occasionally they spark exponential growth. When this happens, whatever bias the resulting mob of excited-speculators suffers from will cause arbitrageur constraints to bind and a bubble to form.

This bubble-generating mechanism satisfies both requirements. It is a theory about population dynamics not pricing errors. So it works the same way regardless of the bias-constraint pair involved. What is more, the same population dynamics operate both during and between bubble episodes. So a researcher can estimate the key parameters governing these population dynamics during normal times when social interactions are busy driving the excited-speculator population to zero. The existing behavioral finance literature on social interactions treats them as yet another source of bias (Shiller 1984, Shive 2010, Burnside et al. 2016, Hirshlei-fer 2020). By contrast, this paper uses social interactions as an on/off switch to regulate the size of the excited-speculator population regardless of which bias they suffer from. The key insight is that social interactions can produce a sharp qualitative change in the size of the excited-speculator population. And, it is possible to characterize what makes an asset more susceptible to such events in the future.

I begin in Section 2 with an economic model built out of two separate components. The first is an offthe-shelf limits-to-arbitrage model, which determines how a bubble will unfold once excited-speculators enter the market. The second is a time-varying population of excited-speculators. During normal times, social interactions cause this population to vanish, leaving prices unaffected by any biases they might suffer from. However, under certain conditions, the exact same social interactions can generate explosive population growth. When this happens, speculator biases (whatever those are) cause arbitrageur constraints (whatever those are) to bind and a bubble to form.

This economic model predicts that speculative bubbles should occur more often in assets where increases in past returns make excited-speculators relatively more persuasive to their peers. Section 3 empirically verifies this ex ante prediction about bubble likelihoods using data on industry-level stock returns via a case-control methodology. Section 4 compares these findings with Greenwood et al. (2019) (GSY) and evaluates the model's out-of-sample forecasting power.

There are three main takeaways from this constellation of empirical results. To start with, they provide an empirical confirmation of the first theory explaining when and where speculative bubbles are most likely to occur next. When investors and policymakers ask economists about the ex ante likelihood of a bubble, economists now have something to base their answer on. To be sure, the economic model is extremely stylized. Part of the goal is to highlight that questions about ex ante bubble likelihoods are outside the scope of the existing limits-to-arbitrage framework. The second main takeaway is that the theory is not too stylized to be empirically useful.

Last but not least, there is no general consensus about which historical market episodes are speculative bubbles. So you might be asking yourself: "what hope do we have of predicting future bubbles if we cannot even agree on what was a bubble after the fact?" The third main takeaway is that it is possible to test ex ante predictions about bubble likelihoods even in the presence of some ex post disagreement about bubble definitions. By analogy, suicide is recognized as the most likely form of gun death in America even though coroners sometimes disagree about whether a particular gun death was a suicide, homicide, or accident.<sup>4</sup>

Definitions matter. But, once we acknowledge that a general kind of phenomenon called a "speculative bubble" sometimes occurs, we do not need complete agreement about the right way to define a speculative bubble in order to have a scientifically rigorous conversation about what makes one more likely to occur in the future. A better definition would allow us to have a more precise conversation. But, a meaningful conversation can still be had using a working definition.

# 1.1. Related Literature

This paper builds on the limit-to-arbitrage literature (Shleifer and Summers 1990, Shleifer and Vishny 1997, Barberis and Thaler 2003, Gromb and Vayanos 2010), which explains how a pricing error can be sustained in equilibrium. Empirical papers, such as Ofek and Richardson (2001), Brunnermeier and Nagel (2004), and Xiong and Yu (2011), were particularly influential in showing that the theoretical limits-to-arbitrage recipe produces speculative bubbles in the real world.

Cutler et al. (1990) and De Long et al. (1990) suggest that bubble episodes are associated with feedback trading just like in my economic model. And, the peer effects literature offers a microfoundation for this feedback trading (Shiller and Pound 1989, Hong et al. 2004, Kaustia and Knüpfer 2012, Bursztyn et al. 2014, Bailey et al. 2018).

The current paper is closely related to GSY. Although both papers predict something about bubbles, they do not predict the same thing about bubbles. GSY predicts whether an already booming industry will crash immediately. I predict whether a not-yet-booming industry will bubble over at all.

# 2. Economic Model

This section presents an economic model that predicts the ex ante likelihood of an asset experiencing a speculative bubble. I start in Section 2.1 with a standard limits-to-arbitrage setup based on Hong and Stein (1999). This component of the model specifies the equilibrium price of an asset as a function of the number of excited-speculators. These excited-speculators suffer from some psychological bias, which causes them to overvalue the asset. There is also a group of arbitrageurs who face some constraint, which prevents them from correcting the resulting pricing error.

Then, in Section 2.2, I describe the bubble-generating mechanism that determines whether there are any excited-speculators in the market. This mechanism is based on the idea that speculators "go mad in herds [but] only recover their senses slowly and one by one"

(Mackay 1841). Most of the time, social interactions between speculators cause the population of excitedspeculators to vanish, leaving prices unaffected by any bias they might have. However, "on occasion and under some conditions" (Barberis 2018), the very same social interactions cause the excited-speculator population to explode. When that happens, speculators' bias (whatever this is) will cause arbitrageurs' constraint (whatever that is) to bind and a bubble to form.

Finally, in Section 2.3, I bring these two separate components together to show that a novel parameter,  $\theta$ , predicts the likelihood of a future bubble episode. This parameter captures how much more persuasive speculators are to their peers when an asset has experienced high past returns. When house prices rise by 10%, it is national news. Everyone starts talking about the housing market,  $\theta \gg 0$ . By contrast, although a 10% increase in the price of textile stocks is a big deal for market participants, it is not going to allow these participants to recruit many friends,  $\theta \approx 0$ . Thus, following a small initial run-up in prices, housing should be more likely to experience a speculative bubble than textiles, regardless of the specific bias-constraint pair involved.

# 2.1. Limits to Arbitrage

Here is how the limits of arbitrage determine the equilibrium price of an asset conditional on the size of the excited-speculator population in the market.

Time proceeds in discrete steps indexed by t = 1, 2, ... There is a single risky asset with a per period payout of  $v_t$  dollars per share:

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

 $\mu_v \gg 0$  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 its per period payout, and  $\varepsilon_{v,t} \sim$  Normal(0, 1) is an independent identically drawn (IID) shock.  $\psi \ge 0$  is the supply of risky shares.

The market contains two kinds of agents: newswatchers and speculators. Newswatchers play two separate roles. First, they incorporate news about an asset's payouts into the price. Every period there is a new unit mass of newswatchers indexed by  $j \in [0, 1]$ .  $\gamma > 0$  denotes newswatchers' risk-aversion coefficient. The *j*th newswatcher in period *t* chooses his demand,  $x_{j,t}$ , to maximize his expected end-of-period utility from consuming the asset's time *t* payout given the price,  $p_t$ :

$$x_{j,t} = \arg\max_{x} \mathbb{E}_{j} \left[ -e^{-\gamma \cdot (v_{t}-p_{t}) \cdot x} \right].$$
(2)

The newswatchers within each cohort have heterogeneous beliefs. Prior to the start of period t, the jth newswatcher observes a private signal,  $s_{j,t}$ , about the asset's time t payout, where

$$v_t = s_{j,t} + \varepsilon_{j,t}.\tag{3}$$

 $s_{j,t}$  denotes the *j*th newswatcher's beliefs after observing his private signal, and  $\varepsilon_{j,t} \sim \text{Normal}(0, 1)$  is noise.

There are two additional assumptions mirroring Admati (1985). First, the heterogeneous beliefs of each newswatcher cohort are correct on average:

$$v_t = \mathbb{E}[s_{j,t}] = \int_0^1 s_{j,t} \cdot \mathrm{d}j.$$
(4)

Second, although different newswatchers get different private signals, every newswatcher's private signal has unit precision,  $Var[\varepsilon_{j,t}] = 1$ .

Newswatchers also play the role of constrained arbitrageur. Their constraint takes the form of bounded rationality.<sup>5</sup> "At every time *t*, newswatchers formulate their asset demands based on the static optimization notion that they buy and hold … More critically … [newswatchers] do not condition on current or past prices. In other words, the equilibrium concept is a Walrasian equilibrium with private valuations, as opposed to a fully revealing rational expectations equilibrium" (Hong and Stein 1999).

This choice of constraint is convenient because it implies that newswatcher beliefs about future realizations of  $n_t$  do not affect an asset's current price level. This independence completely separates the limits of arbitrage (how a bubble is sustained) from the bubble-generating mechanism (how likely they are to occur). In essence, it is like using log utility in a macrofinance model when things like hedging demand are not germane to the model's conclusions.

In addition to newswatchers, the market also contains  $n_t \in [0, 1)$  excited-speculators. For now, take the number of excited-speculators in the market at time *t* as given. Excited-speculators cause an asset's price to move for nonfundamental reasons. They do so whenever  $n_t > 0$ . Excited-speculators' demand at time *t* is proportional to recent performance:

$$z_t(n) \stackrel{\text{def}}{=} (\lambda \cdot r_{t-1}) \times n.$$
(5)

 $\lambda > 0$  is the strength of excited-speculators' bias. A larger  $\lambda$  means excited speculators are more biased.  $r_{t-1} \stackrel{\text{def}}{=} (p_{t-1} + v_{t-1}) / p_{t-2}$  is the past realized return. In Section 3, I will think of  $r_{t-1}$  as an industry's return over the past two years.

The functional form of Equation (5) is consistent with excited-speculators having extrapolative beliefs. And, there is much evidence that, during bubble episodes, investors extrapolate returns (Cutler et al. 1990, De Long et al. 1990, Hong and Stein 1999, Case et al. 2012, Glaeser and Nathanson 2017, Barberis 2018, Barberis et al. 2018). However, this paper is not proposing a new model of extrapolative beliefs during bubbles. The predictions about ex ante bubble likelihoods do not depend on the details of this functional form as I show in Corollary 1. The aggregate demand of newswatchers and excited-speculators must clear markets:

$$\int_0^1 x_{j,t} \cdot \mathrm{d}j + z_t(n) = \psi.$$
(6)

Newswatchers ignore the information content of prices, so they have demand

$$x_{j,t} = (s_{j,t} - p_t) / \gamma.$$
 (7)

If the *j*th newswatcher's private signal results in beliefs that are higher than the price,  $(s_{j,t} - p_t) > 0$ , then he will buy; otherwise, he will sell.

**Proposition 1** (Limits to Arbitrage). When there are  $n_t = n$  excited-speculators in the market, equilibrium asset prices are given by

$$p_t(n) = v_t - \gamma \times \psi + \gamma \cdot (\lambda \cdot r_{t-1}) \times n.$$
(8)

This price is increasing in the size of an asset's payout,  $v_t$ , decreasing in the number of shares,  $\psi$ ; and increasing in the number of excited-speculators,  $n \in [0, 1)$ .

# 2.2. Bubble-Generating Mechanism

Here is the bubble-generating mechanism I propose. There are  $K \gg 1$  speculators. Let  $N_{\tau} \ge 0$  denote the number of speculators excited about a particular asset, and let  $n_{\tau} \stackrel{\text{def}}{=} N_{\tau}/K$  denote the corresponding excited fraction. This mechanism controls whether there are any excited-speculators in the market,  $n_{\tau} > 0$ .

In this paper, social interactions between speculators are not a separate bias. Instead, they are an on/ off switch. They regulate whether some speculator bias will cause some arbitrageur constraint to bind. To make this distinction as clear as possible, I set the model up so that the limits of arbitrage and this on/ off switch operate on completely different timescales. The fraction of excited-speculators at time t,  $n_t \in [0, 1)$ , will be the steady-state outcome of social interactions that play out in continuous time,  $\tau \ge 0$ . Because  $r_{t-1}$ will represent an industry's return over the past two years in Section 3, it is helpful to think about t as a clock that measures time in two-year increments and  $\tau$  as a clock that measures time in days,  $\Delta \tau \approx \Delta t/500$ .

Popular accounts describe bubbles as "social epidemics" where "news of price increases spurs investor enthusiasm which spreads by psychological contagion from person to person, amplifying stories that might justify the price increases and bringing in a larger and larger class of investors, who, despite doubts about the real value of an investment, are drawn to it partly through envy of others' successes and partly through a gambler's excitement" (Shiller 2000).

I use two rules to capture the feedback between an asset's past returns and the current number of speculators excited about it. The first rule is that apathetic speculators get excited via interactions with their n

excited friends:

$$\lim_{\Delta \tau \searrow 0} \mathbb{P}\mathbf{r}[n_{\tau+\Delta \tau} - n_{\tau} = +1/K \mid n_{\tau} = n, r] = \theta \cdot r \cdot (1-n) \times n.$$
(9)

 $\theta \cdot r \cdot (1 - n)$  is the per capita excitation rate. (1 - n) is the size of the apathetic speculator population.  $r \in$  $(0, \infty)$  is an asset's past performance.  $\$1 \cdot r$  is the amount of money you would have today if you had invested one dollar in the asset two years ago. Whenever  $r \gg 1$ , apathetic speculators missed out on a profitable investment opportunity two years ago.  $\theta \in (0, 1)$ reflects how much more persuasive an asset's excitedspeculators are following an increase in past returns. In the late 1990s, tech returns generated a lot of wordof-mouth buzz,  $\theta \approx 1$ , whereas textile stocks would have to have very strong returns for excited-speculators to have any sway on their friends,  $\theta \approx 0$ .

The second rule is that speculators "recover their senses slowly and one by one" (Mackay 1841):

$$\lim_{\Delta \tau \searrow 0} \mathbb{P}\mathbf{r}[n_{\tau + \Delta \tau} - n_{\tau} = -1/K \mid n_{\tau} = n, \ r] = 1 \times n.$$
 (10)

Multiplying *n* by one equates the phrase "slowly and one by one" with a constant per capita departure rate. 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. Moreover, because Equation (10) does not contain *r*, this rate is also independent of the asset's past return.<sup>6</sup>

Both of these rules are smooth and continuous functions, so you might expect that a slight increase in an asset's past return would always result in a slight increase in its excited-speculator population. But this is not what happens. The two rules actually produce a sudden change in steady-state behavior as an asset's past returns cross a critical threshold called a *"bifurcation"* (Hirsch et al. 2012, Guckenheimer and Holmes 2013, Kuznetsov 2013, Strogatz 2014).

Let  $n_{\tau}(n_0, \theta, r)$  be the fraction of speculators excited about an asset at time  $\tau \ge 0$  if  $n_0$  were excited<sup>7</sup> at time  $\tau = 0$ :

$$n_{\tau}(n_0,\theta,r) \stackrel{\text{def}}{=} \left\{ n \in [0,1) : n = \int_0^{\tau} [\theta \cdot r \cdot (1-n_u) - 1] \cdot n_u \cdot \mathrm{d}u \right\}.$$
(11)

A steady-state population,  $\bar{n} \in [0, 1)$ , is a population such that

1 0

$$\mathcal{SS}(\theta, r) \stackrel{\text{def}}{=} \{ \bar{n} \in [0, 1) : n_{\tau}(\bar{n}, \theta, r) = \bar{n} \quad \forall \tau \ge 0 \}.$$
(12)

We say that a particular steady-state value,  $\bar{n} \in SS(\theta, r)$ , is stable if small perturbations away from  $\bar{n}$  die out over time.<sup>8</sup> The following proposition characterizes a critical return threshold,  $r_{\star}$ , such that there will be a stable nonzero steady-state population of excited-speculators in the market whenever  $r > r_{\star}$ .

# **Proposition 2** (Bubble-Generating Mechanism). *Define* $r_{\star} \stackrel{\text{def}}{=} 1/\theta$ .

1. If  $r < r_{\star}$ , there is only one steady-state value for the excited-speculator population,  $SS(\theta, r) = \{0\}$ . This lone steady state,  $\bar{n} = 0$ , is stable.

2. If  $r > r_*$ , there are two steady-state values,  $SS(\theta, r) = \{0, (r - r_*)/r > 0\}$ . Only the strictly positive steady state,  $\bar{n} = (r - r_*)/r > 0$ , is stable.

When  $r < r_*$ , any initial population of excited-speculators,  $n_0 > 0$ , quickly loses interest. However, as soon as an asset's return crosses the critical threshold,  $r > r_*$ , that same initial population gives rise to a persistent crowd of excited-speculators.

To see why social interactions can serve as an on/off switch, consider what happens when there is only one excited-speculator,  $N_{\tau} = n_{\tau} \cdot K = 1$ . In this situation, the entire population of excited-speculators will go extinct if its lone member cannot excite at least one of his apathetic friends before he himself comes to his senses:

$$\underbrace{\mathbb{P}\mathbf{r}[\Delta N_{\tau} = +1 \mid N_{\tau} = 1, r]}_{=\theta \cdot r \cdot (1-1/K) \cdot 1 \cdot \Delta \tau \approx \theta \cdot r \cdot \Delta \tau} < \underbrace{\mathbb{P}\mathbf{r}[\Delta N_{\tau} = -1 \mid N_{\tau} = 1]}_{=1 \cdot \Delta \tau}.$$
(13)

This simplifies down to the following inequality:  $\theta \cdot r < 1$ . And, by rearranging terms, it is immediately clear there will be no excited-speculators left whenever past returns are sufficiently low,  $r < r_* = 1/\theta$ , as shown in the lower left panel of Figure 1.

However, although the same economic forces are at work when  $r > r_{\star}$ , these forces produce the exact opposite result. If a single speculator happens to get excited when  $r > r_{\star}$ , this lone agent will likely be able to excite at least one friend before he himself comes to his senses. And, the same will also be true for this newly excited friend, which leads to exponential growth in the excited-speculator population. Thus, when  $r > r_{\star}$ , the excited-speculator population will remain stably above zero in steady state as shown in the lower right panel of Figure 1.

# 2.3. Ex Ante Likelihood

I now fold the limits-of-arbitrage setup from Section 2.1 and the bubble-generating mechanism from Section 2.2 into a single asset-pricing model, which can be used to predict the likelihood of a future bubble in a way that does not depend on the nitty gritty details of the particular bias-constraint pair involved.

I connect the previous two subsections by assuming that the number of excited-speculators in the limits-to-arbitrage model from Section 2.1 is given by the steady-state solution in Proposition 2 when  $r = r_{t-1}$ :

$$n_{t} = \begin{cases} (r_{t-1} - r_{\star})/r_{t-1} & \text{if } r_{t-1} > r_{\star} = 1/\theta \\ 0 & \text{otherwise.} \end{cases}$$
(14)

#### Figure 1. Bubble-Generating Mechanism



*Notes.* (Top panel) The *x* axis: an asset's past return,  $r \in (0, \infty)$ . The *y* axis: steady-state solutions,  $\bar{n} \in SS(\theta, r)$ , for a population of excited-speculators. The solid black line reports stable steady states; the dashed red line reports unstable ones. Population displays a bifurcation at  $r_* = 1/\theta$ . (Lower panels) Transition to steady state when  $r < r_*$  vs. when  $r > r_*$ . The *x* axis (top): fraction of speculators who are currently excited about an asset, *n*. The *y* axis (top): growth rate of excited-speculator population,  $\frac{dn}{d\tau}$ . When  $r < r_*$ , this growth rate is always negative for all n > 0 as indicated by the solid line remaining below the *x* axis. By contrast, when  $r > r_*$ , this growth rate is positive for some population values, n > 0, as indicated by the solid line arching above the *x* axis. The *x* axis (bottom): time since initial group of  $n_0 \ge 0$  speculators got excited about an asset at time,  $\tau = 0$ . The *y* axis, (bottom): number of speculators excited about an asset at time  $\tau > 0$ ,  $n_{\tau} = n_{\tau}(n_0, \theta, r)$ . Different shades of gray denote different initial population sizes,  $n_0 \in [0, 1)$ . When  $r < r_*$ , any initial population of speculators  $n_0 > 0$  that happens to get excited will quickly lose interest and disperse so  $n_{\tau}(n_0, \theta, r) \rightarrow \bar{n} = 0$ . But, when  $r > r_*$ , the excited-speculator population will converge to  $\bar{n} = (r - r_*)/r > 0$  whenever a single speculator happens to get excited.

It is as if on the first day of each discrete time period t lasting roughly two years, a single speculator gets excited about an asset. After he enters, the madness of crowds either takes over or does not depending on both the asset's past return,  $r = r_{t-1}$ , and its  $\theta$  parameter. The resulting steady-state excited-speculator population (if one exists at all) is responsible for any nonfundamental demand shock realized in period t.

A speculative bubble occurs at time t when newswatchers push an asset's time (t - 1) return above  $r_{\star}$ , causing a crowd of excited-speculators to flood the market at time t. Let  $b_t$  be an indicator variable for whether there is a nonzero population of excited-speculators in period t:

$$b_t = \mathbf{B}(\theta, r_{t-1}) \stackrel{\text{def}}{=} \mathbb{1}[r_{t-1} > 1/\theta] = \mathbb{1}[n_t > 0].$$
(15)

When  $b_t = 1$ , Proposition 1 says there will be an equilibrium pricing error.

It is easiest to understand the asset-pricing implications of this model by studying a sample price path. Figure 2 shows a single realization simulated using parameters  $\psi = 0$ ,  $\mu_v = 1.0$ ,  $\kappa_v = 0.1$ ,  $\sigma_v = 0.1$ ,  $\theta = 0.4$ , and  $\lambda = 0.5$ . The figure is meant to illustrate economic intuition, not match empirical facts. That being said, if each period *t* corresponds to roughly two years, then these parameters generate speculative booms that last 30 months and speculative busts that last 22 months on average. This closely matches the averages of 27.4 and 23.0 months reported in Table 1a.

The black line in the top panel depicts the equilibrium price each period,  $p_t(n)$ , whereas the thin green line depicts the size of an asset's payout,  $v_t$ . Because the risky asset is in zero net supply for this simulation,  $\psi = 0$ , these two lines fall right on top of one another when there are no excited-speculators in the market,  $n_t = 0$ . Additionally, this is exactly what happens most of the time. In the middle panel, the black line represents the asset's realized return in the previous period,  $r_{t-1} = (p_{t-1} + v_{t-1})/p_{t-2}$ .

The asset's past returns are typically below the dashed blue line representing the critical value of  $r_* = 1/\theta$ . However, there are four different points during this simulation where newswatchers pushed the asset's past return above the critical threshold level (i.e., where  $b_t = B(\theta, r_{t-1}) = 1$ ). These instances are denoted by  $t_1$ ,  $t_2$ ,  $t_3$ , and  $t_4$  on the *x* axis in the bottom panel. The height of the red bars in the bottom panel depicts the size of the excited-speculator population during each period,  $n_t$ .

The way in which the limits-to-arbitrage setup and the bubble-generating mechanism interact with one another introduces a delay. An asset's return in period

#### Figure 2. Sample Price Path



Notes. Simulated outcomes using  $\psi = 0$ ,  $\mu_v = 1.0$ ,  $\kappa_v = 0.1$ ,  $\sigma_v = 0.1$ ,  $\theta = 0.4$ , and  $\lambda = 0.5$ . The *x* axis represents time, t = 1, 2, ..., 100. (Top panel) The black line is price level,  $p_t(n)$ . The thin green line is per period payout,  $v_t$ . Red shaded regions denote times when excited-speculators caused arbitrageur constraints to bind and a speculative bubble to form. (Middle panel) The black line is the realized return in the previous period,  $r_{t-1} = (p_{t-1} + v_{t-1})/p_{t-2}$ . The dashed blue line is the threshold return level,  $r_* = 1/\theta$ . When  $r_{t-1} < r_*$ , there are no excited-speculators,  $n_t = 0$ . (Bottom panel) Red vertical bars report the number of excited-speculators,  $n_t$ . Four bubbles,  $b_t = 1$ , are labeled  $t_1$ ,  $t_2$ ,  $t_3$ , and  $t_4$ .

*t* determines the size of the excited-speculator population in period (t + 1). This delay is essential to how the model works. Morally speaking, it is not possible to predict the likelihood of a future bubble without some sort of delay.

This delay is absent from previous work on social interactions between speculators (Shiller 1984, Shive 2010, Burnside et al. 2016, Hirshleifer 2020). Because these papers were looking for a way to amplify existing speculator biases, they studied contemporaneous feedback between asset returns and speculator demand. Both contemporaneous feedback and delayed feedback are likely present in real-world markets. The existence of contemporaneous feedback will clearly affect how speculative bubbles unfold once they begin. However, by definition, contemporaneous feedback cannot affect the likelihood of a future bubble, so I exclude this force from my main analysis. Nevertheless, I show in Section A.3 of the online appendix that the model's predictions about ex ante bubble likelihoods carry over to a setting that includes contemporaneous feedback.

We know from Equation (1) how an asset's payout will fluctuate over time. Proposition 1 tells us how newswatchers determine an asset's equilibrium price during normal times as a function of its current payout,  $p_t = v_t - \gamma \times \psi$ . Proposition 2 then tells us how the critical boundary between normal times and bubble episodes is set,  $r_* = 1/\theta$ . From these three pieces of information, we can infer how likely it is that an asset's past returns will cross this critical boundary. Thus, we have all the ingredients in place to predict the likelihood of a future bubble.

**Proposition 3** (Ex Ante Likelihood). Assume  $r_{t-1} < r_*$ . Controlling for an asset's fundamentals,  $(\mu_v, \kappa_v, \sigma_v)$ , the probability of a speculative bubble at time (t + 1) is increasing in  $\theta$ :

$$\partial_{\theta} \mathbb{E}_{t-1}[b_{t+1} \mid r_{t-1} < r_{\star}] > 0.$$
(16)

In other words, suppose two assets have identical payout parameters,  $(\mu_v, \kappa_v, \sigma_v)$ , and time (t - 1) returns,  $r_{t-1} < r_{\star}$ . Then, the asset with the higher speculator-persuasiveness sensitivity,  $\theta$ , will be more likely to experience a speculative bubble at time (t + 1).

In this paper, I will treat  $\theta$  as an exogenous assetspecific constant encapsulating all of the things that make one asset's speculators more or less persuasive to their peers following high returns. I will not offer an explanation for why  $\theta$  differs across assets. This is an interesting topic for future research, but there are two things worth remembering here. First, this is the very first theory that explains the ex ante likelihood of bubbles. There is no other theory with strong microfoundations to fall back on. Second, you do not need to know why  $\theta$  varies across assets when using it to make predictions. By analogy, the Capital Asset Pricing Model (Sharpe 1964) is not a bad model because William Sharpe never explained why some stocks have higher market betas. You do not need this information when using beta to price assets.<sup>9</sup>

A key fact about Proposition 3 is that the ex ante predictions do not depend on the severity of speculator bias,  $\lambda > 0$ . This captures the sense in which the bubble-generating mechanism in Section 2.2 operates independently from the specific bias-constraint pair that happens to bind in Section 2.1. "How?" and "How likely?" are two fundamentally different kinds of questions.

**Corollary 1** (How vs. How Likely). Assume  $r_{t-1} < r_*$ . Controlling for an asset's fundamentals,  $(\mu_v, \kappa_v, \sigma_v)$ , changes in  $\lambda$  do not affect the probability that it will experience a speculative bubble at time (t + 1):

$$\partial_{\lambda} \mathbb{E}_{t-1}[b_{t+1} \mid r_{t-1} < r_{\star}] = 0.$$
(17)

Figure 3 illustrates the logic behind Corollary 1. The left panel depicts periods t = 40, ..., 80 of the sample

	Industry	Start date	retPast2Yr	netPast2Yr	retPast5Yr	bookToMkt	volatility	Peak date	Trough date	Boom length	Bust length	retBoom	retBust
1	Softw	06/78	44.8	35.3	29.0	0.75	44.5	06/83	11/84	60	17	90.2	-54.1
5	Cnstr	06/79	46.6	25.8	121.8	1.16	39.8	11/80	07/82	17	20	254.7	-64.2
С	Gold	07/79	39.6	15.3	10.1	0.43	45.0	08/60	06/82	14	21	306.1	-69.6
4	Oil	07/79	34.4	10.1	129.7	0.98	25.7	11/80	07/82	16	20	210.3	-49.2
ß	Mines	08/79	45.3	11.4	121.7	0.87	22.7	03/81	07/82	19	16	124.6	-58.8
9	Softw	03/95	43.3	26.9	203.9	0.21	35.7	02/00	09/02	59	31	204.7	-78.6
	Hardw	03/95	43.3	27.0	38.0	0.43	30.8	08/00	09/02	65	25	273.3	-84.6
8	Soda	05/96	42.4	-13.0	82.4	0.63	22.6	06/98	05/00	25	23	161.3	-44.0
6	Chips	86/60	33.7	-9.4	145.3	0.23	65.7	03/00	09/02	18	30	234.9	-84.6
10	LabEq	11/99	29.6	-17.5	161.5	0.27	43.7	02/00	09/02	С	31	145.8	-71.4
11	Steel	05/00	36.5	5.2	82.1	0.34	56.3	08/00	09/02	С	25	161.2	-76.0
12	Fin	06/00	38.7	5.5	342.9	0.28	58.0	08/00	09/02	2	25	148.1	-54.3
13	Coal	03/04	36.4	30.7	452.7	0.43	35.8	06/08	02/09	51	8	118.2	-79.6
14	Cnstr	06/04	48.2	25.3	111.8	0.42	32.8	01/06	02/09	19	37	92.5	-67.4
15	Mines	06/04	43.4	20.4	25.9	0.39	30.5	10/07	02/09	40	16	210.3	-66.5
Avg			40.4	13.3	137.3	0.52	39.3			27.40	23.00	182.4	-66.9
Sd			5.4	16.5	120.8	0.29	13.0			21.84	7.40	66.0	12.7
<i>Notes</i> over the n book the n date.	. There are the past two nost recent 1 ToMkt, and umber of m Avg, averag	15 bubble ep years (raw a month prior volatilit onths from p e; Sd, standar	isodes. Each is a and net) and 50% to the peak in y (percentage pe eak date to troug rd deviation.	five-year local pr 6+ raw returns o which returns o which are all me r year) are all me gh date. retBooi	tice maximum fo ver the past five ver the past two assured using dat m (percentage) is	r an industry or years. Followin o years were le ta available as o the return dur	ccurring on the p g the peak, the ii g the peak, the ii ess than 50%. $r\epsilon$ f the start date. B ing two years pr	eak date. The ndustry's pric etPast2Yr ( com length is ior to the pea	ere must be at le ce level must the percentage), ne s the number of k date. ret Bus	ast one month c an fall by at leas t Past 2Yr (pe months from th t (percentage) i	during the run- t 40%. The star rcentage), ret e start date to <u>r</u> s the return fro	up with 100% It date of eacl Past5Yr (p peak date. Bu om peak date	6+ returns 1 bubble is arcentage), st length is t to trough

Table 1a. Bubble Episodes

price path in Figure 2, which was simulated using  $\lambda = 0.50$ . The right panel depicts the exact same time period for the exact same simulation in a world where speculator biases are 50% more extreme,  $\lambda = 0.75$  rather than 0.50. Increasing  $\lambda$  increases the length of the second and third bubble episodes. When  $\lambda = 0.50$ , the second bubble episode only lasts one period, whereas when  $\lambda = 0.75$ , it lasts two. Likewise, the third bubble episode goes from three to four periods long when  $\lambda$  increases from 0.50 to 0.75.

Yet the 50% increase in the severity of speculators' bias does not affect the number of bubble episodes (i.e., it does not affect ex ante likelihood of a bubble). This is an important distinction because policymakers are often interested in the ex ante likelihood of a bubble and not the specific bias-constraint pair involved. It is common to see articles discussing whether "China's stimulus program is prone to blow more bubbles in the economy next year."<sup>10</sup> The main concern in these news articles is the likelihood a future speculative bubble will occur, not the particular ingredients used in the limits-to-arbitrage recipe to create it.

# 3. Empirical Analysis

Predicting the likelihood of a future bubble means finding some parameter related to bubble formation that is relatively stable over time. That way, researchers can estimate it during normal times and use this estimate to make out-of-sample predictions about the likelihood of a future bubble episode. If we were to model bubbles as the result of a sudden inexplicable parameter change, then the model would not help us predict where these sudden changes were most likely to occur next.





*Notes.* Periods t = 40, ..., 80 from the sample price path simulation in Figure 2. All random number seeds and simulation parameters (except  $\lambda$ ) are kept the exact same:  $\psi = 0$ ,  $\mu_v = 1.0$ ,  $\kappa_v = 0.1$ ,  $\sigma_v = 0.1$ , and  $\theta = 0.4$ . (a)  $\lambda = 0.50$ . Severity of excited-speculators' bias is the same as in Figure 2,  $\lambda = \lambda_0 = 0.50$ . (b)  $\lambda = 0.75$ . Severity of excited-speculators' bias is 50% higher than in Figure 2,  $\lambda = 1.5 \times \lambda_0 = 0.75$ . Speculative bubbles at times  $t_2$  and  $t_3$  now last longer than in Figure 2. But this change in bias severity does not affect the number of bubbles.

The model developed in the previous section argues that  $\theta$  (i.e., the sensitivity of speculator persuasiveness to changes in past returns) fits this criteria. The model predicts that speculative bubbles should occur more often in assets with high  $\theta$  values, where small increases in past returns make excited-speculators relatively more persuasive to their peers. This section empirically verifies this central prediction.

I begin in Section 3.1 by giving a working definition of an industry-level bubble based on the definitions of a "boom month" and "crash" given in Greenwood et al. (2019). This working definition captures the industrylevel episodes that often get called bubbles, such as the rise and fall of technology stocks during the late 1990s and early 2000s. However, in Section 4, I show that the model's ex ante predictions about the likelihood of a future bubble are robust to some ex post disagreement over how to define one. In layperson's terms, if you think that the bubble definition I am using wrongly includes some episode that was not actually a bubble, then feel free to leave it out. The results will be unchanged.

Next, in Section 3.2, I lay out the empirical approach I use. The previous literature on the econometrics of speculative bubbles has been interested in predicting the timing of the crash; however, the current paper is interested in the ex ante likelihood of a future speculative bubble regardless of whether the crash will go down in history books as "Black Monday" or "Black Tuesday." So, to avoid complications related to how each speculative bubble evolves, I use a case-control methodology. I match each bubble episode (cases of interest) to another otherwise identical industry-month observation where no subsequent bubble occurs (control observations) based on data observed prior to the start of each bubble episode. The end result is a matched data set containing pairs of ex ante identical industry-month observations, one that subsequently experiences a speculative bubble and another that does not.

In Section 3.3, I describe how I empirically estimate  $\theta$  for each industry using data prior to the start of any bubble. I do this by looking at the relationship between changes in daily media coverage and daily returns for a given industry. I refer to the empirical analog to  $\theta$  as theta. Finally, in Section 3.4, I show that differences in the value of theta across industrymonth observations in the matched data set predict which of each matched pair will subsequently realize a speculative bubble.

# **3.1. Bubble Episodes**

The empirical analysis in this paper will primarily focus on industry-level bubble episodes. I do this because many historical accounts of bubbles have a strong industry component. The DotCom bubble is a prototypical example. What is more, industry-level data offer several econometric advantages. Relative to

data offer several econometric advantages. Relative to using national-level price indexes à la Kindleberger (1978), analyzing industry returns results in higherpowered statistical tests. And statistical power is important given how rare bubbles are. Kindleberger (1978) documents only 34 national-level bubble episodes in over 400 years! Relative to using individual stock returns, analyzing industry returns also helps circumvent several difficult measurement problems, such as the problem posed by the entrance of new firms in bubble industries.

I use data from the Center for Research in Securities Pricing (CRSP) to compute the value-weighted monthly returns of the 49 Fama and French (1997) industries from January 1975 to December 2017.<sup>11</sup> I remove the "Other" industry because this industry designation does not represent a cohesive collection of firms, leaving industries i = 1, ..., 48. I only use U.S. firms (share codes: 10 or 11) listed on the New York Stock Exchange, American Stock Exchange, or NASDAQ (exchange codes: 1, 2, or 3). I exclude industry-month observations that contain fewer than five firms.

Let  $ret_{i,t}$  (percentage) denote the *i*th industry's value-weighted return in month *t*. Figure 4 shows the cumulative return of each industry. Each black ribbon depicts the value  $\$1 \cdot \prod_{s=Jan75}^{t} (1 + ret_{i,s})$  (i.e., the value in month *t* of a continuously reinvested portfolio that started by investing one dollar in the *i*th industry in January 1975).

Let retPast2Yr<sub>*i*,*t*</sub> (percentage) be the *i*th industry's raw return over the past two years:

$$retPast2Yr_{i,t} \stackrel{\text{def}}{=} \prod_{t'=t-23}^{t} (1 + ret_{i,t'}) - 1.$$
(18)

If retPast2Yr = -100%, you would have lost your entire one dollar investment from two years ago, whereas if retPast2Yr = 200%, your initial one dollar investment would have tripled in value. Let netPast2Yr<sub>*i*,*t*</sub> = def  $\prod_{t'=t-23}^{t} (1 + \text{ret}_{i,t'}) - \prod_{t'=t-23}^{t} (1 + \text{mkt}_{t'})$  (percentage) be an industry's return net of the market over the past two years, and let retPast5Yr<sub>*i*,*t*</sub> def  $\prod_{t'=t-59}^{t} (1 + \text{ret}_{i,t'}) - 1$  (percentage) be its past fiveyear raw return.

A speculative bubble involves a boom followed by a bust. "In classical accounts of financial market bubbles, the price of an asset rises dramatically over the course of a few months or even years, reaching levels that appear to far exceed reasonable valuations of the asset's future cash flows ... The bubble eventually ends with a crash, in which prices collapse even more quickly than they rose" (Barberis et al. 2018). So I start my definition of industry-level bubble episodes by looking for local maxima. Month *t* is a local maximum if an industry's cumulative returns are higher than at any point  $\pm 5$  years. Not all local maxima are bubbles. To be a bubble, the run-up must contain at least one boom month where the industry realizes 100%+ returns over the past two years (raw and net) and 50%+ raw returns over the past five years as in Greenwood et al. (2019). This rules out situations where an industry's price level gradually rose at a steady rate over several years. In addition, the local maximum must also be followed by a < -40% decline from peak to trough during the next five years. This rules out situations where an industry's returns merely plateau after the boom before beginning to rise once more.<sup>12</sup>

There are 15 industry-level bubble episodes in my sample period as shown in Table 1a. The start of each bubble episode is the last month prior to the peak when the industry had returns over the past two years less than 50%. This is the month just before the industry's returns really began to take off. The end of a bubble is the month of the trough in the industry's returns following the peak.

# 3.2. Case-Control Methodology

Suppose two otherwise identical assets both realize a modest price increase at time (t - 1). No bubble is currently taking place in either asset. One of these assets has a high value of  $\theta \gg 0$ ; the other has a low  $\theta \approx 0$ . Proposition 3 says that the asset with the higher  $\theta$  (i.e., the one whose speculators get a lot more persuasive following good returns) is more likely to experience a bubble in the future at time (t + 1).

To test this prediction using data on industry-level stock returns, I need to create a data set containing pairs of otherwise identical industry-month observations where one observation realizes a future bubble episode and the other does not. Then I need to show that differences in  $\theta$  predict which observation in each pair will experience the future bubble. This empirical goal is well suited to a case-control study.

Specifically, I look at the data available at the start date of each bubble episode in Table 1a. Then, for each prebubble observation, I find an otherwise similar industry-month observation with no subsequent bubble episode. Both the prebubble observations and the matched observations will have had modest positive returns at the time of matching. Both will have similar fundamentals and return volatility at the match date as well. Finally, I check whether  $\theta$  predicts which observation will realize a future bubble.

What is the right way to measure similarity between industry-month observations? One possibility would be to look only at past returns. However, Proposition 3 assumes that the two assets are the same both in terms of their past returns and in terms of their fundamental parameters. What is more, a key lesson from Greenwood et al. (2019) is that "there is much more to a bubble than a mere security price





*Notes.* The *x* axis: time in months from January 1975 to December 2017. The *y* axis (log scale): dollar value at time *t* of a continuously reinvested industry-specific portfolio that started with one dollar at the opening bell on the first trading day of January 1975. Green regions indicate the boom period during a bubble episode. Red regions indicate the subsequent bust period. Grayed-out regions indicate aftershock bubbles, which overlap with an earlier episode in the same industry. White hash marks indicate boom months.

increase" (Greenwood et al. 2019, p. 22). So I use more than just past returns when choosing a matched control for each bubble episode.

I match based on returns over the past two years (both raw and net of the market) and raw returns over the past five years. And I also match on book-to-market ratio and return volatility over the past two years. Let bookToMkt<sub>*i*,*t*</sub> denote the average book-to-market ratio of stocks in the *i*th industry as of time *t*, which I compute using data from Compustat. Let volatility<sub>*i*,*t*</sub> (percentage per year) denote the value-weighted average of annualized volatility for firms in the *i*th industry. Table 1b gives the matched industry-month observation associated with each bubble episode. I use the optmatch package in R to perform the match. I measure distance between industry-month observations using the Mahalanobis metric, which rescales Euclidean distance to account for the covariance between match variables. I want the matched sample to be drawn during normal times, so I do not look for matches during another bubble episode  $\pm 2$  years. A matched observation need not be contemporaneous with the start date of the corresponding bubble episode. For example, the machinery industry in October

	Industry	Match date	retPast2Yr	netPast2Yr	retPast5Yr	bookToMkt	volatility	Match distance
1	Gold	09/11	47.3	35.8	48.4	0.59	44.6	0.40
2	FabPr	06/79	38.8	18.0	106.9	1.04	39.7	0.42
3	Rubbr	10/11	41.9	14.4	14.9	0.48	47.1	0.26
4	BldMt	07/79	34.0	9.7	131.7	1.02	26.4	0.11
5	Smoke	07/80	52.6	14.3	128.1	0.94	21.5	0.28
6	Soda	12/87	42.3	24.0	198.8	0.29	39.4	0.38
7	Mach	10/94	51.0	30.1	43.7	0.43	30.6	0.24
8	Hshld	07/84	46.1	-7.1	89.9	0.64	23.2	0.22
9	BusSv	02/00	33.0	-9.3	155.8	0.23	68.7	0.29
10	MedEq	11/99	32.7	-14.3	182.5	0.19	43.7	0.31
11	Hardw	08/11	30.3	4.7	48.7	0.34	56.7	0.43
12	ElcEq	06/00	38.9	5.7	297.5	0.26	64.7	0.88
13	Hlth	06/82	44.6	35.9	480.5	0.48	35.7	0.38
14	Hlth	08/94	47.9	23.4	102.4	0.43	32.3	0.16
15	BldMt	11/12	39.2	15.8	17.8	0.47	31.5	0.23
Avg			41.4	13.4	136.5	0.52	40.4	0.33
Sd			6.9	15.5	121.9	0.28	14.2	0.18

Table 1b. Matched Controls

Notes. The industry-month observations are selected as matches for each bubble episode listed in Table 1a. retPast2Yr (percentage), netPast2Yr (percentage), retPast5Yr (percentage), bookToMkt, and volatility (percentage per year) are all measured using data available as of the match date for each industry. Match distance is the Mahalanobis distance between a matched observation and the bubble episode it was matched to. Avg, average; Sd, standard deviation.

1994 is matched to the speculative bubble in the computer hardware industry that begins in Match 1995.

Figure 5 plots the prebubble industry-month observations (cases; solid red lines) together with the matched controls (dotted red lines). Matches can be drawn from any observation taken during normal times, which are represented by the black regions in each panel. The gray regions represent time intervals where a bubble is taking place  $\pm 2$  years where no matches can be selected.

turnover<sub>*i*,*t*</sub> (percentage) is the value-weighted trading volume as a percentage of shares outstanding for firms in the *i*th industry in month *t*.  $age_{i,t}$  (years) is the value-weighted age of firms in the *i*th industry in month *t*. A firm's age is defined as the number of years since it first appeared in CRSP or Compustat. ageTilt<sub>*i*,*t*</sub> (percentage) is the difference between the *i*th industry's equal-weighted and age-weighted returns over the past two years. When  $ageTilt_{i,t} > 0$ , younger firms have outperformed older firms. newIssuance<sub>*i*,t</sub> (percentage) is the percentage of firms in the *i*th industry that issued equity during the past two years.  $\Delta$ sales<sub>*i*,*t*</sub> (percentage per year) is the value-weighted year-over-year sales growth of firms in the *i*th industry. CAPE is the market-wide cyclically adjusted price-toearnings (P/E) ratio from Robert Shiller's website.<sup>13</sup> Finally, retAccel<sub>*i*,*t*</sub>  $\stackrel{\text{def}}{=} \prod_{s=t-23}^{t} (1 + \text{ret}_{i,s}) - \prod_{s'=t-23}^{t-12} (1 + \text{ret}_{i,s'})$ (percentage) is the difference between the *i*th industry's value-weighted return over the past two years and its value-weighted return one year ago.

To estimate  $\theta$  in Section 3.3, I collect the data on the number of times the *i*th industry is mentioned in a *Wall Street Journal* (WSJ) article each day. These data come from the WSJ Historical Archive available via ProQuest. Let WSJcoverage<sub>*i*,*t*</sub> (percentage) denote the

average percentage of WSJ articles in each of the past 24 months that reference the *i*th industry. And let  $\Delta$ WSJcoverage<sub>*i*,*t*</sub> (percentage per year) be the year-over-year change in the average percentage of WSJ articles referencing the *i*th industry.

Table 2 shows summary statistics for the 15 prebubble observations as well as the 15 matched control observations with no subsequent bubbles. The 15 prebubble industry-month observations look just like their matched counterparts in terms of past returns, book-to-market, and volatility in the first five rows. This is mechanical. The matched observations were selected because of their similarity along these dimensions. However, there is no reason why the prebubble observations and their matched counterparts need to look similar along other dimensions. Nevertheless, retAccel is the only variable with a systematic difference between cases and controls.

# **3.3. Estimating** $\theta$

The economic model predicts that bubbles will occur more often in assets where increases in past returns make excited-speculators much more persuasive to their peers. I use the parameter  $\theta$  to capture the sensitivity of speculator persuasiveness to changes in an asset's past returns. In this section, I explain how I estimate  $\theta$  for each observation in my case-control study. In Section 3.4, I show that prebubble cases tend to have higher  $\theta$  values than the matched control observations.

My estimation strategy is motivated by the fact that financial news outlets strategically choose which assets to write stories about with an eye toward maximizing their readership (Mullainathan and Shleifer 2005). Thus,



Figure 5. Bubble Episodes and Matched Controls

*Notes.* The *x* axis: time in months from January 1975 to December 2017. The *y* axis (log scale): dollar value at time *t* of a continuously reinvested industry-specific portfolio that started with one dollar at the opening bell on the first trading day of January 1975. Solid red lines denote the start date of a bubble episode in Table 1a. Gray regions denote bubble episodes  $\pm 2$  years. Dotted red lines denote the matched observations in Table 1b. Black regions indicate the normal times from which these otherwise similar control observations were selected.

when an industry receives lots of positive coverage in the WSJ, it suggests that many speculators are currently excited about the industry.<sup>14</sup>

WSJ news articles come from the WSJ Historical Archive available through ProQuest. I include stories about a particular industry on each trading day from January 1975 through December 2017. I classify a WSJ article about the *i*th industry as positive if its title contains more positive than negative words according to the dictionary in Loughran and McDonald (2011).<sup>15</sup> isPosArticle<sub>*i*,*d*</sub> is an indicator for the existence of a positive WSJ article about the *i*th industry on trading day *d*.

When  $\theta \gg 0$ , an increase in an industry's past returns will cause speculators to talk a lot more about the industry. When  $\theta \approx 0$ , the same increase in past returns will not generate much speculator attention. So to proxy for  $\theta$ , I calculate the probability that there is a positive WSJ article about the *i*th industry on the day following an increase in its returns:

theta<sub>*i*,*t*</sub> 
$$\stackrel{\text{def}}{=} \sum_{d \in (t-24, t]} \left( \frac{\text{retWentUp}_{i,d-1}}{\sum_{d \in (t-24, t]} \text{retWentUp}_{i,d-1}} \right)$$
  
× isPosArticle<sub>*i*,*d*</sub>.

(19)

Table 2. Summary Statistics for the Matched Data Set

	Ca	ses	Con	trols	
	Avg (1)	Sd (2)	Avg (3)	Sd (4)	Difference (5)
retPast2Yr	40.4	5.4	41.4	6.9	-0.97
					(0.43)
netPast2Yr	13.3	16.5	13.4	15.5	-0.15
					(0.03)
retPast5Yr	137.3	120.8	136.5	121.9	0.75
					(0.02)
bookToMkt	0.5	0.3	0.5	0.3	0.00
					(0.01)
volatility	39.3	13.0	40.4	14.3	-1.07
					(0.21)
turnover	146.6	101.6	140.8	129.2	5.82
					(0.14)
age	28.3	12.8	33.3	16.2	-4.98
					(0.93)
ageTilt	-0.3	7.5	2.9	12.3	-3.18
					(0.85)
newIssuance	17.7	7.5	22.2	12.3	-4.47
					(1.13)
$\Delta$ sales	19.9	11.6	36.8	47.7	-16.87
					(1.33)
CAPE	39.5	33.6	34.7	31.1	4.81
					(0.41)
retAccel	33.9	19.8	15.8	21.1	18.01**
					(2.41)
WSJcoverage	3.1	3.3	1.7	2.1	1.42
					(1.42)
$\Delta$ WSJcoverage	0.3	0.5	0.4	0.8	-0.13
					(0.52)
No. of observations	1	5	1	5	

Notes. Cases indicate mean and standard deviation for 15 prebubble observations. Controls indicate the same statistics for 15 matched control observations. Difference is the difference in means across samples. Numbers in parentheses are *t* statistics clustered by industry. retPast2 $Yr_{i,t}$  (percentage) is the value-weighted return over the past two years.  $\mathtt{netPast2Yr}_{i,t}$  (percentage) is the value-weighted return net of the market over the past two years.  $retPast5Yr_{i,t}$ (percentage) is the value-weighted return over the past five years.  $bookToMkt_{i,t}$  is the average book-to-market ratio in month t. volatility<sub>it</sub> (percentage per year) is the value-weighted daily volatility in month t. turnover<sub>*i*,*t*</sub> (percentage per month) is the valueweighted trading volume divided by shares outstanding in month *t*. age<sub>*i*,*t*</sub> (years) is the value-weighted firm age in month *t*. ageTilt<sub>*i*,*t*</sub> (percentage) is the difference between equal-weighted and ageweighted return over the past two years. newIssuance<sub>it</sub> (percentage) is the percentage of firms issuing equity in the past two years.  $\Delta sales_{i,t}$  (percentage per year) is the value-weighted yearover-year sales growth.  $CAPE_t$  is the market-wide cyclically adjusted P/E ratio in month *t*. retAccel<sub>*i*,*t*</sub> (percentage) is the value-weighted return in months [t - 23, t] minus value-weighted return in months [t-23, t-12]. WSJcoverage<sub>*i*,*t*</sub> (percentage) is the average percentage of WSJ articles in past two years about an industry. AWSJcoverage<sub>i,t</sub> (percentage per year) is the year-over-year change in the percentage of WSJ articles about an industry. Avg, average; Sd, standard deviation.

\*\*Statistically significant at the 5% level.

I compute this conditional probability using data on each trading day during the previous two years,  $d \in (t-24, t]$ , when an industry's returns went up on the previous day, retWentUp<sub>*i*,*d*-1</sub>  $\stackrel{\text{def}}{=} 1$ [ret<sub>*i*,*d*-1</sub>  $-\hat{\mu}_{i,d-1} > \hat{\sigma}_{i,d-1}$ ], where  $\hat{\mu}_{i,d-1}$  and  $\hat{\sigma}_{i,d-1}$  are the mean and

volatility of the industry's daily returns over the previous 252 days, respectively.

Table 3 describes the distribution of theta in the matched data set. The inset table shows that, for these industry-month observations (either a prebubble case or a matched control observation), a jump in returns on day (d-1) was immediately followed by a positive WSJ news article on day d roughly 6% of the time on average during the previous two years. However, theta was 6.4% points higher for the 15 prebubble cases, which saw positive returns followed by positive coverage 9.2% of the time, than for the 15 matched controls, which only saw positive returns followed by positive coverage 2.8% of the time. The number in parentheses is the t statistic for the 6.4%-point difference in means when clustering standard errors by industry.

The histogram accompanying Table 3 depicts the entire distribution of theta values (i.e., both prebubble cases (solid black circles) and matched controls (dotted white circles)). Although there is overlap in the two distributions, theta tends to be higher for the prebubble cases than for the matched controls. Only two control observations have theta > 5%, whereas 10 of the prebubble cases do. The 1996 speculative bubble in the Soda industry is the only case with theta = 0%.

There are four things about my estimation strategy worth emphasizing. First, when creating my list of otherwise similar control observations in Table 1b, I did not consider any industry-month observations during the two years following a bubble. So the estimated theta values in Table 3 only use data observed during normal times. These estimates do not incorporate any information taken during bubble episodes.

Second, my strategy for estimating  $\theta$  is not tailored to a specific bias-constraint pair. I am using the existence of a *Wall Street Journal* article with positive tone as a proxy for whether speculators get excited about an industry the day after it experiences positive returns. I am not looking at which bias is responsible, nor am I looking at which constraint might handicap arbitrageurs. This is important because different biases and constraints were likely at work during different bubble episodes in Table 1a.

Third, the economic model says that each industry is endowed with a speculator sensitivity parameter,  $\theta$ , which is relatively stable over time. A speculative bubble occurs when good news pushes an industry's recent performance above the critical threshold,  $r_{\star} = 1/\theta$ . In other words,  $\theta$  does not increase prior to a bubble; instead, returns increase, and this results in an inflow of excited-speculators when  $\theta \gg 0$ .

Figure 6 shows that this is exactly what is going on in the matched data set for the estimated theta values. theta is 6.4% points higher for the 15 prebubble

#### Table 3. Theta Distribution in the Matched Data Set



*Notes.* Full sample indicates mean and standard deviation of theta across all 30 observations in the matched data set. Prebubble cases and matched controls indicate the same statistics for just this subset of observations. Difference is the spread between average theta in prebubble cases and average theta in matched controls. The number in parentheses is the *t* statistic clustered by industry. In the histogram, each circle represents an observation in the matched data set. The *x* axis is the estimated value of theta. The *y* axis is the number of observations in the bin where each bin has a width of 1%. Solid black circles are bubble episodes listed in Table 1a. Dotted white circles are matched controls listed in Table 1b.

\*\*\*Statistically significant at the 1% level.

cases than for the 15 matched controls on the date of the match. But this is not due to theta increasing for the industries with future bubble in the preceding months. The difference in theta between the prebubble cases and the matched controls is stable over the preceding three years, with an average value of 7.1% during this time period.

Because the sensitivity of speculator persuasiveness to past returns is a novel economic quantity, I also report the full time series of theta estimates for each industry in Table 4. The gray regions denote time periods when a speculative bubble was taking place in an industry. In addition to the sparkline plot, I also report the mean and standard deviation of theta for each industry during normal times.

It is clear from Table 4 that theta is not perfectly constant over time for each industry. However, it is also clear from the table that the main source of variation in theta is across industries not over time. Some industries tend to have theta values close to zero (e.g., Boxes, Guns, Medical Equipment); others tend to

#### Figure 6. Evidence of Parameter Stability



*Notes.* Average difference between theta in the 15 prebubble cases (red) and in the 15 matched observations (black) during three years (36 months) prior to the match date. The *x* axis: months prior to the start of the speculative bubble. The *y* axis: average estimated theta value in a given month. Shaded regions show the average value  $\pm$  standard error clustered by industry.

have much larger values of theta (e.g., Autos, Finance, Oil).

Fourth and finally, in the empirical analysis, a single trading period corresponds to 2 years,  $\Delta t \approx 2$  years. In my model, I assume that the excited-speculator population dynamics play out on a much faster timescale,  $\Delta \tau \ll \Delta t$ , allowing it to reach a steady state every trading period. And consistent with this intuition, my approach to estimating theta uses daily data, a horizon  $\frac{2 \text{ years}}{1 \text{ day}} \times \frac{252 \text{ trading days}}{1 \text{ year}} \approx 500$  times shorter than a trading period.

# 3.4. Main Results

I find that differences in theta predict which observations experience a future speculative bubble just as suggested by the economic model. Let willBeBubble indicate whether an observation in the matched data set is one of the 15 prebubble cases. Column (1) in Table 5a regresses this indicator on the estimated theta:

willBeBubble<sub>*i*,*t*</sub> = 
$$\hat{\alpha} + \hat{\beta} \cdot \text{theta}_{i,t} + \hat{\varepsilon}_{i,t}$$
. (20)

The slope coefficient of 4.55 implies that a 1%-point increase in theta is associated with a 4.55%-point increase in the likelihood of a future speculative bubble. We saw in Table 3 that theta is 6.4% points higher in the prebubble cases than in the matched controls on average. This increase in theta among the 15 prebubble cases predicts a  $4.55 \times 6.4\% \approx 29\%$ -point increase in their ex ante bubble likelihoods. This is an economically large change relative to the base rate of 50%. If you guessed at random whether an industrymonth observation in the matched data set would be followed by a bubble, you would guess correctly 50% of the time. Columns (2) and (3) in Table 5a confirm that the

		Avg	Sd			Avg	Sd		Avg	Sd	
Aero		6.3	4.0	FabPr		7.7	4.3	Paper	 1.4	1.9	
Agric		3.5	3.0	Fin		13.6	6.6	PerSv	 0.4	0.9	
Autos		19.4	8.2	Food		9.6	4.5	RlEst	11.4	6.1	
Banks		13.0	7.9	Fun		14.2	7.1	Rtail	12.4	4.2	
Beer		4.7	3.4	Gold		6.4	4.1	Rubbr	5.4	3.6	
BldMt		2.0	2.8	Guns	•••••	1.5	1.4	Ships	 1.2	1.3	
v Books		9.5	4.9	Hardw		6.1	4.6	Smoke	 4.6	4.2	
Boxes		1.3	1.5	Hlth		12.1	7.4	Soda	2.4	2.2	
BusSv		4.0	4.6	Hshld		1.6	2.0	Softw	14.0	5.9	
Chems		9.6	4.7	Insur		7.8	4.5	Steel	8.8	5.8	
Chips		10.0	8.6	LabEq	-	1.4	2.0	Telcm	 12.3	6.2	
Clths	<b>b</b>	3.9	3.1	Mach		10.5	4.8	Toys	4.0	3.1	
Cnstr		9.2	5.5	Meals		7.9	5.5	Trans	7.1	4.2	
Coal		4.4	3.6	MedEq		1.0	1.3	Txtls	3.9	3.8	
Drugs		7.4	6.6	Mines		4.5	4.0	Util	4.2	3.8	
ElcEq		2.1	3.0	Oil		15.1	5.7	Whlsl	 2.3	2.1	

Table 4. Theta over Time and Across Industries

Notes. Sparkline plots depict theta from January 1975 to December 2017. Gray regions indicate bubble episodes. Black regions correspond to normal times. All plots have the same *y*-axis scale. Avg and Sd are mean and standard deviation, respectively, of theta for each industry computed using data observed during normal times.

main result is a stable predictive relationship over the past few decades.

Because the matched control for each prebubble industry-month observation was chosen based on its similarity in terms of past return, book-to-market ratio, and volatility, these variables cannot explain why theta predicts the ex ante likelihood of bubbles. But, this is not the only way that the prebubble cases might be different.

In columns (2)–(10) of Table 5b, I investigate the predictive power of other variables. To make the coefficients comparable, I divide each predictor by the difference between its average in the 15 prebubble cases (Table 1a) and its average in the 15 matched controls (Table 1b). As a result, the coefficient in each column corresponds to the predicted change in ex ante bubble likelihoods when moving from a variable's average value among matched controls to its average among prebubble cases.

For example, Table 3 tells us theta is 6.4% points higher in the prebubble cases. So column (1) in Table 5b reports the  $\hat{\beta}$  from the following regression:

willBeBubble<sub>*i,t*</sub> =  $\hat{\alpha} + \hat{\beta} \cdot (\text{theta}_{i,t}/6.4) + \hat{\varepsilon}_{i,t}$ . Table 2 tells us that retAccel is 18.01% points higher in the prebubble cases. So column (8) in Table 5b reports the  $\hat{\beta}$  from the following specification: willBeBubble<sub>*i,t*</sub> =  $\hat{\alpha} + \hat{\beta} \cdot (\text{retAccel}_{i,t}/18.01) + \hat{\varepsilon}_{i,t}$ .

Table 5b reveals that variables other than theta also have some forecasting power. The 15 prebubble cases get 1.42% more WSJ coverage on average (Table 2, column (5)). Column (9) in Table 5b says this difference predicts a 6.71%-point increase in the probability of a future bubble. The 15 prebubble cases have 16.87% higher sales growth on average. Column (6) in Table 5b says this difference predicts a 5.96%-point drop in the probability of a future bubble cases have 18.01% higher return acceleration on average. Column (8) in Table 5b says this difference predicts a 17.20%-point increase in the probability of a future bubble.

That being said, column (1) in Table 5b reports that the 6.4% case-control difference in average theta predicts a 28.97%-point increase in their ex ante bubble likelihoods. This effect size dwarfs the other estimates

Dependent variable: willBeBubble	Full sample (1)	Before 2000 (2)	On/after 2000 (3)
intercept	22.58**	20.94*	21.60
	(2.35)	(1.81)	(1.29)
theta	4.55***	5.56***	4.09***
	(4.78)	(2.95)	(3.73)
Adjusted R <sup>2</sup>	0.26	0.25	0.25
No. of observations	30	16	14

Table 5a.	Ex Ante	Likelihood	of	Bubbles

Notes. Each column reports the results of a separate univariate regression. willBeBubble is the indicator for the 15 prebubble cases in the matched data set. A coefficient of +1 implies a 1%-point increase in the likelihood of a future bubble. Each prebubble case is matched to the most similar industry-month observation without a subsequent bubble based on retPast2Yr (percentage), netPast2Yr (percentage), retPast5Yr (percentage), bookToMkt, and volatility (percentage per year) as of the start of the bubble. Prebubble and matched observations mechanically look similar along these dimensions. retPast2Yr\_i, (percentage) is the value-weighted return over the past two years. netPast2Yr\_i, (percentage) is the value-weighted return over the past two years. retPast5Yr\_i, (percentage) is the value-weighted return over the past five years. bookToMkt\_i, is the average book-to-market ratio in month t. volatility\_i, (percentage per year) is the value-weighted daily volatility in month t. theta (percentage) is the empirical proxy for sensitivity of speculator persuasiveness to increases in past returns. Numbers in parentheses are t statistics clustered by industry.

\*Statistical significance at the 10% level; \*\*statistical significance at the 5% level; \*\*\*statistical significance at the 1% level.

Dependent variable: willBeBubble	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)
theta	28.97*** (4.78)										33.74**
turnover	(100)	0.07									-0.08 (0.27)
age		(0.11)	-3.02								-3.55
ageTilt			(0.07)	-2.53							-2.90
newIssuance				(1.10)	-4.35						-0.85
∆sales					(1.04)	$-5.96^{***}$					(0.20) $-15.33^{***}$ (3.80)
CAPE						(2.91)	0.59				2.14**
retAccel							(0.40)	17.20**			(2.13) 15.04** (2.41)
WSJcoverage								(2.01)	6.71*		(2.41) -2.76 (0.44)
$\Delta$ WSJcoverage									(1.60)	-0.95	(0.44) 2.58
Adjusted R <sup>2</sup>	0.26	-0.04	0.00	-0.01	0.01	0.03	-0.03	0.14	0.03	(0.58) -0.03	0.52

Table 5b. Ex Ante Likelihood of Bubbles

Notes. Each column reports the results of a separate regression involving the 30 observations in the matched data set. A coefficient of +1 implies a 1%-point increase in the likelihood of a future bubble. willBeBubble is the indicator for the 15 prebubble cases. Each prebubble case is matched to the most similar industry-month observation without a subsequent bubble based on retPast2Yr (percentage), netPast2Yr (percentage), retPast5Yr (percentage), bookToMkt, and volatility (percentage per year) as of the start date of the bubble. theta (percentage) is the sensitivity of speculator persuasiveness to increases in past returns. turnover<sub>it</sub> (percentage per month) is the value-weighted trading volume divided by shares outstanding in month t. age<sub>it</sub> (years) is the value-weighted firm age in month t. ageTilt<sub>it</sub> (percentage) is the difference between equal-weighted and age-weighted returns over the past two years. newIssuance<sub>it</sub> (percentage) is the percentage of firms issuing equity in the past two years.  $\Delta$ sales<sub>it</sub> (percentage per year) is the value-weighted year-over-year sales growth. CAPE<sub>t</sub> is the market-wide cyclically adjusted P/E ratio in month t. retAccel<sub>it</sub> (percentage of WSJ articles in the past two years about an industry.  $\Delta$ WSJcoverage<sub>it</sub> (percentage per year) is the value-weighted return in months [t - 23, t - 12]. WSJcoverage<sub>it</sub> (percentage) is the average percentage of WSJ articles in the past two years about an industry. Intercept estimates are not reported for clarity. I divide each predictor by the difference between its average in the prebubble cases and the matched controls (e.g., the estimated 17.20 in column (8) is the  $\hat{\beta}$  from the regression willBeBubble<sub>it</sub> =  $\hat{\alpha} + \hat{\beta} \cdot (retAccel_{it}/18.01) + \hat{\varepsilon}_{it}$ ). Column (5) of Table 2 reports the average case-control difference for each predictor. Numbers in parentheses are t statistics clustered by industry.

\*Statistical significance at the 10% level; \*\*statistical significance at the 5% level; \*\*\*statistical significance at the 1% level.

in columns (2)–(10). Column (11) shows that the predictive power of theta is unchanged when including these other variables.

In addition, column (11) in Table 5b also highlights how theta is not just a proxy for overall media coverage. Including WSJcoverage and ∆WSJcoverage does not erode the coefficient on theta. An industry must receive some media coverage to have theta >0. There is no way for  $\Pr[isPosArticle_{i,d}]$  $retWentUp_{id-1} > 0$  if there are no articles about the *i*th industry, positive or otherwise. However, lots of media coverage on its own do not mechanically imply a high theta. For example, much of the WSJ coverage of semiconductor stocks in recent years has focused on privacy concerns involving products with microchips in them. Articles often do not mention chip makers' recent stock market gains or losses. So, although the semiconductor industry has gotten a lot of coverage lately, its theta is still low.

# 4. Supporting Results

This section provides three sets of supporting results that complement my main empirical analysis. I start in Section 4.1 by comparing and contrasting this paper with GSY. My definition of a speculative bubble is based on GSY. Yet, many of the variables that predict crashes in GSY do not predict future bubbles in this paper. Each paper is answering a different question about bubbles, and this accounts for the difference in our empirical results.

In Section 4.2, I include all industry-month observations that are sufficiently similar to some past bubble episode in my regression analysis. This amounts to making rolling out-of-sample predictions about whether the next frothy industry-month observation will bubble over, which mirrors the real-world applications that econometricians and policymakers care about most.

Last but not least, in Section 4.3, I show that the results in Section 3 are robust to some ex post disagreement about how to define a bubble. If you strongly believe that 1 of the 15 bubble episodes used in the analysis is not a bubble, then feel free to leave it out. The findings in Table 5a will be qualitatively unchanged.

# 4.1. Bubbles for Fama

I define a bubble as a local price maximum with a speculative boom during the run-up (retPast2Yr > 100%, netPast2Yr > 100%, and retPast5Yr > 50%) and a crash following the peak (< -40% decline). Yet, although my boom and bust definitions come from GSY, many variables that predict crashes in GSY do not predict the likelihood of a future bubble in this paper.

The main reason for the difference in our results is the difference in our research questions. GSY asks if any variables other than past returns predict whether an already booming industry will suffer an immediate crash. By contrast, I ask if  $\theta$  predicts whether a not yet booming industry will suffer a future bubble no matter when the peak occurs. This difference in questions has three important consequences.

First, although we both use the same boom and crash definitions, each paper has a different bubble definition. Because GSY is interested in predicting the timing of the crash, the authors only consider a subset of bubble episodes—namely, those involving a single boom month followed by an immediate crash. Bubbles with multiple booms or a single boom followed by a delayed crash are not considered bubbles in GSY. The list of booms without a crash in Greenwood et al. (2019, table 1b) contains several such episodes, like the mid-2000s bubble in the mining industry.

Second, although GSY does not use this language, it is also using a case-control methodology. Greenwood et al. (2019, table 1a) presents cases of interest (i.e., bubbles), and Greenwood et al. (2019, table 1b) lists a set of matched control observations (i.e., nonbubbles). Yet, although both papers are case-control studies, we each select our control observations in different ways.

GSY wants to show that "there is much more to a bubble than a mere security price increase" (Greenwood et al. 2019, p. 3). So they populate their table 1b in Greenwood et al. (2019) with observations that have similar past returns to their bubble episodes in table 1a in Greenwood et al. (2019). Then they document that other variables, such as new issuance, predict which table an observation was drawn from.

I want to show that differences in  $\theta$  predict which of two otherwise identical industries will experience a future bubble. So, having learned from GSY that there is more to bubbles than high returns, I look for more than just high returns when choosing control observations for Table 1b. I also match on book-to-market and volatility.

Third, GSY wants to predict whether an already booming industry will immediately crash, whereas I want to predict whether a not-yet-booming industry will bubble over at all. So GSY studies each bubble after it has already started, whereas I study each bubble before it starts.

Section B.2 in the online appendix shows how to modify my empirical exercise to match GSY's approach. Table B2 in the online appendix shows that if I extend my sample period, exclude bubbles where the first boom month is not immediately followed by a crash, choose control observations based only on past returns, and study the first boom date rather than the start date, then I can match GSY's findings.

The control observations in Table 1b are different from the ones listed in Greenwood et al. (2019, table 1b). But in Figure B2 in the online appendix, I show that if I were to match bubble episodes, I would select the same control observations as in Greenwood et al. (2019, table 1b). In short, there is no fundamental conflict between our results. It is possible to bring GSY's results in line mine by making the three changes above. However, if I were to do this, I would no longer be testing whether differences in  $\theta$  predict the likelihood of a future bubble among otherwise identical industry-month observations.

# 4.2. Out-of-Sample Predictions

Table 5a shows that theta predicts the likelihood of a future bubble in a data set containing 15 prebubble cases and 15 otherwise similar controls. This evidence confirms the key prediction of the economic model in Section 2. Yet, there are two reasons why it does not directly speak to econometricians' and policymakers' core question. Given everything else we know about an industry, how likely is it that a modest price increase today will morph into a speculative bubble in the future?

The first is about feasibility. My main empirical analysis pairs each bubble episode with a single control observation. I do this using data observed as of the start of each bubble episode, which is defined as the most recent month where the bubble industry had retPast2Yr<sub>*i*,*t*</sub> < 50%. But how are econometricians and supposed to know in real time if they are looking at the start of one of the bubbles in Table 1a?

The second is more quantitative. There are many industry-month observations that look similar to each prebubble case. Table 5a shows that theta predicts the ex ante likelihood of bubbles in a data set containing one match per bubble. But maybe its forecasting power deteriorates when looking at many matches per bubble episode?

I address both these concerns in Table 6 by using theta to make rolling out-of-sample predictions. The data in this table include all industry-months that are similar to the start of some previous bubble. In each month t, I include any industry with a match distance less than 0.88 (i.e., the maximum match distance in Table 1b) to the start date of a bubble that crashed prior to month t.

The resulting data set contains 187 industry-month observations. These 187 observations include the start dates for 11 of the 15 bubble episodes listed in Table 1a. This finding demonstrates that it is possible to recognize start dates in real time.

Each column in Table 6 reports a separate logit specification. I switch to a logit model because 11/187 =5.9% is close enough to zero that willBeBubble should be treated as a binary response. Column (1) reveals that each 1%-point increase in theta is associated with a 7.86% larger bubble likelihood. theta averages 10.1% in the 11 prebubble cases, whereas it only averages 6.3% in the remaining 176 observations. This 10.1% - 6.3% = 3.8%-point difference in theta predicts a  $3.8\% \times 7.86 \approx 30\%$  larger ex ante bubble likelihood. theta is a quantitatively useful predictor even when the probability of a future bubble is much less than 50%.

Columns (2) and (3) in Table 6 confirm that theta's out-of-sample forecasting power is not driven by the previously examined covariates. Column (4) shows there is nothing different about theta's predictions

Table 6. Out-of-Sample Predictions

Dependent variable: willBeBubble	(1)	(2)	(3)	(4)
theta	7.86**		12.88**	7.35**
	(1.96)		(2.18)	(2.02)
turnover	. ,	0.36	0.23	. ,
		(1.27)	(0.71)	
aqe		-7.60**	-6.95**	
5		(2.20)	(2.21)	
ageTilt		1.03	-0.28	
5		(0.70)	(0.17)	
newIssuance		2.54	2.45	
		(0.92)	(0.87)	
Asales		0.57	0.41	
		(1.07)	(0.81)	
CAPE		-0.69	-0.82	
		(0.54)	(0.61)	
retAccel		1.30	1.29	
		(0.74)	(0.60)	
WSJcoverage		-4.73	-7.81	
		(0.41)	(1.27)	
∆WSJcoverage		5.32	4.66	
		(0.25)	(0.17)	
No. of observations	187	187	187	191
No. of bubbles	11	11	11	15

Notes. Each column reports the results of a separate logistic regression. A coefficient of +1 indicates a 1% proportional increase in the likelihood of a future bubble. Intercept estimates are not reported for clarity. Each industry-month observations in columns (1)-(3) had a match distance of <0.88 to some previous bubble episode based on retPast2Yr (percentage), netPast2Yr (percentage), retPast5Yr (percentage), bookToMkt, and volatility (percentage per year) as of the start date of the episode. Eleven of the 15 episodes in Table 1a were sufficiently similar to some past bubble episode; the remaining four were dissimilar to every previous episode. Column (4) reestimates column (1) including these four episodes. willBeBubble (true/false) is the indicator for whether an observation was followed by a bubble. theta (percentage) is the sensitivity of speculator persuasiveness to increases in past returns.  $turnover_{i,t}$  (percentage per month) is the value-weighted trading volume divided by shares outstanding in month *t*.  $age_{i,t}$  (years) is the value-weighted firm age in month *t*. ageTilt<sub>*i*,*t*</sub> (percentage) is the difference between equal-weighted and age-weighted returns over the past two years. newIssuance<sub>it</sub> (percentage) is the percentage of firms issuing equity in the past two years.  $\Delta sales_{i,t}$  (percentage per year) is the value-weighted year-over-year sales growth.  $CAPE_t$  is the market-wide cyclically adjusted P/E ratio in month t.  $retAccel_{i,t}$  (percentage) is the value-weighted return in months [t - 23, t] minus return in months [t-23, t-12]. WSJcoverage<sub>*i*,*t*</sub> (percentage) is the average percentage of WSJ articles about an industry each month during the past two years.  $\Delta WSJcoverage_{i,t}$  (percentage per year) is the year-over-year change in average percentage of WSJ articles about an industry each month. Numbers in parentheses are *t* statistics clustered by industry.

\*\*Statistical significance at the 5% level.

# Table 7. Ex Post Disagreement

	No. of observations	$\begin{pmatrix} 15\\k \end{pmatrix}$	Avg[coef]	Sd[coef]	Avg[se]	Min[coef]	<pre>Pr[signif]</pre>
k = 0	30	1	4.55	0.00	0.95	4.55	1.000
1	28	15	4.58	0.32	1.00	4.29	1.000
2	26	105	4.59	0.45	1.05	4.00	1.000
3	24	455	4.62	0.57	1.10	3.63	1.000
4	22	1,365	4.65	0.70	1.17	3.24	0.996
5	20	3,003	4.69	0.83	1.24	2.79	0.981

*Notes.* Each row show the results of reestimating column (1) of Table 5a after omitting  $k \in \{0, 1, 2, 3, 4, 5\}$  of the 15 speculative bubbles. The resulting matched data set contains 30 observations when k = 0 episodes are omitted, 28 observations when k = 1 episode is omitted, 26 observations when k = 2 episodes are omitted, and so on. The first row with k = 0 is the same specification as in column (1) of Table 5a. The number of observations is the number of industry × month observations in the matched data set when omitting k speculative bubbles.  $\binom{15}{k}$ 

indicates the number of ways to omit *k* episodes from among the 15 bubble episodes. Avg[coef] indicates the average slope coefficient in Equation (20) when omitting different combinations of *k* episodes. Sd[coef] indicates the standard deviation of these slope coefficients. Avg[se] indicates the average standard error of the slope coefficient when using standard errors clustered by industry. Min[coef] indicates the smallest slope coefficient in any regression when omitting some combination of *k* episodes. Pr[signif] indicates the fraction of the time that the *p*-value was <0.05 when using standard errors clustered by industry.

in the four episodes—Oil (July 1979), Computer Hardware (March 1995), Laboratory Equipment (November 1999), and Steel (May 2000)—with start dates not included in the original 187.

# 4.3. Ex Post Disagreement

Suppose you believe that 1 of the 15 bubble episodes in my data set is not a bubble. There are  $\binom{15}{1} = 15$  different ways to remove one episode. No matter which episode we disagree about, the omission will not affect the 4.55-point estimate for the slope coefficient on theta. The average slope coefficient across the 15 different ways to omit k = 1 episode is  $4.57 \pm (0.31)$  in Table 7.

Walking down the rows in Table 7, we see that we would have to disagree on at least four episodes before seeing a single insignificant result. Still, 1,360 of the  $\binom{15}{4} = 1,365$  ways to omit four bubble episodes still yield statistically significant results.

Even if we removed a full third of the 15 bubble episodes in my main analysis, we would only see insignificant results 1.9% of the time. Only 58 of the  $\binom{15}{5}$  = 3,003 ways to omit five episodes would result in a statistically insignificant coefficient on theta. Not a single one of the 3,003 combinations would flip the sign, Min[*coef*] = 2.79 > 0.

These results do not imply that the definition of a "*speculative bubble*" is irrelevant. Definitions matter. If everyone agreed on how to define a speculative bubble, researchers would be able to study this phenomenon much more precisely. But we do not have to wait for a consensus definition before studying this phenomenon at all. We can test predictions about what makes a bubble more or less likely in the future

even when there remains some disagreement about how to define one after the fact.

# 5. Conclusion

This paper aims to expand the set of questions that economists ask about speculative bubbles. Suppose there has recently been some good news about an asset. And as a result, its price has increased a bit. How likely is it that this modest price increase will morph into a full-blown speculative bubble? This is a question about the likelihood of a future bubble, not about how the bubble will be sustained in equilibrium.

To answer a question about the likelihood of a future bubble, you need a theory not of the limits to arbitrage but of the force that sporadically causes these limits to bind. This paper proposes the first such theory. The on/off switch is based on the idea that speculators "go mad in herds [but] only recover their senses slowly and one by one" (Mackay 1841). The theory predicts that speculative bubbles should be more likely to occur in assets where increases in past returns make excited-speculators much more persuasiveness to their peers.

The theory gives the first way to forecast the likelihood of a future bubble episode based on information observed during normal times when no bubble is currently taking place. I empirically validate this prediction with a case-control study using data on industry-level stock returns. In the process, I demonstrate that it is possible to test such ex ante predictions about bubble likelihoods even in the presence of some ex post disagreement about bubble definitions.

# Acknowledgments

The author thanks Nick Barberis, Pedro Bordalo, Zhi Da, Jesse Davis, Xavier Gabaix, Valentin Haddad, Sam Hartzmark, Lawrence Jin, Josh Pollet, and Jeff Wurgler for helpful comments. Marco Sammon provided research assistance. This paper also benefited from presentations at INSEAD; Illinois; National Bureau of Economic Research; Finance Research Association; Maastricht; Notre Dame; Imperial College; Washington University in St. Louis; the University of Massachusetts Amherst; the University of California, Los Angeles; Cavalcade; and Western Finance Association. This paper previously circulated under the title "The Madness of Crowds and the Likelihood of Bubbles."

# Appendix. Technical Appendix

**Proof of Proposition 1.** Newswatchers have demand given by  $x_{j,t} = (s_{j,t} - p_t)/\gamma$ . Market clearing implies that  $\psi = \int_0^1 x_{j,t} \cdot dj + (\lambda \cdot r_{t-1}) \times n_t$ . Because newswatcher signals are correct on average,  $\mathbb{E}[s_{j,t}] = v_t$ , we can conclude that

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

Rearranging to isolate  $p_t$  on the left-hand side gives the desired result.  $\Box$ 

**Proof of Proposition 2.** Suppose the excited-speculator population obeys the law of motion  $\frac{dn}{dt} = G(n, \theta, r) = \theta \cdot r \cdot (1-n) \times n - n$ . We can rewrite this as

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

There are then two possibilities to consider.

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

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

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

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

**Proof of Proposition 3.** The probability of realizing a speculative bubble at time (t + 1) given that  $r_{t-1} < r_*$  can be written as

$$\mathbb{E}_{t-1}[B(\theta, r_t) \mid B(\theta, r_{t-1}) = 0] = \mathbb{P}r_{t-1}[r_t > r_\star \mid r_{t-1} < r_\star]$$

Given the stochastic process governing asset payouts, we know that

$$\mathbb{E}_{t-1}[\Delta v_t] = \kappa_v \cdot (\mu_v - v_{t-1})$$
  
$$\mathbb{V}ar_{t-1}[\Delta v_t] = \sigma_v^2.$$

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

$$\mathrm{pdf}_{t-1}(p) = \frac{1}{\sigma_v \cdot \sqrt{2 \cdot \pi}} \cdot e^{-\frac{1}{2\sigma_v^2} \cdot (p - \upsilon_{t-1} - \mathbb{E}_{t-1}[\Delta \upsilon_t] + \gamma \cdot \psi)^2}.$$
 (A.1)

This PDF can be used to write down an integral expression for the probability of a speculative bubble at time (t + 1) because  $r_t = p_t/p_{t-1}$ :

$$\mathbb{E}_{t-1}[B(\theta, r_t) | r_{t-1} < r_{\star}] = \int_{p_{t-1}/\theta}^{\infty} p df_{t-1}(p) \cdot dp.$$
(A.2)

Notice two facts about this integral. Fact 1 is that  $pdf_{t-1}(p)$  is a strictly positive function. Fact 2 is that  $\theta$  plays no part in  $pdf_{t-1}(p)$  itself; it only enters into Equation (A.2) as a boundary condition. Thus, increasing  $\theta$  simply increases the size of the interval over which a strictly positive function is being integrated. So,  $\mathbb{E}_{t-1}[B(\theta, r_t) | r_{t-1} < r_*]$  must be strictly increasing in  $\theta$ .  $\Box$ 

**Proof of Corollary 1.** This corollary follows from the fact that excited-speculators' beliefs only affect equilibrium prices during a speculative bubble when  $n_t > 0$ . However, the likelihood of entering into a speculative bubble is based on considerations made prior to the bubble when  $n_t = 0$ . Thus,  $\lambda$  does not show up in either the PDF in Equation (A.1) or the boundary conditions in Equation (A.2).  $\Box$ 

#### Endnotes

<sup>1</sup> See "Fed Set to Cut Rates for First Time in Decade. Is It a Risk?" at https://apnews.com/article/economy-donald-trump-us-news-ap-top-news-jerome-powell-dad5b590357d4380a1eb28720fc20774 (*AP News*, July 29, 2019).

<sup>2</sup> See "Greenspan Is Concerned About 'Froth' in Housing" at https://www.nytimes.com/2005/05/21/business/greenspan-is-concerned-about-froth-in-housing.html#:~:text=Greenspan%20emphasized%20that %20he%20sees,local%20bubbles%22%20around%20the%20country (*The New York Times*, May 21, 2005).

<sup>3</sup> Of course, the limits-to-arbitrage literature also tells us that different bias-constraint pairs will result in speculative bubbles that unfold in different ways. They will have booms of different sizes, peaks at different times, and crashes of different severities. Thus, because my bubble-generating mechanism must be agnostic about the bias-constraint pair responsible for a bubble, it will only speak to the likelihood of a future bubble. It will remain silent about how the resulting bubbles will unfold.

<sup>4</sup> See "Gun Deaths in America" at https://fivethirtyeight.com/ features/gun-deaths/ (*FiveThirtyEight*, June 6, 2016).

<sup>5</sup> I could model newswatchers who were short-sale constrained rather than myopic. In such a model, the option to sell at an inflated price during future bubble episodes would push up an asset's price today as in Scheinkman and Xiong (2003). If our goal was to understand cross-sectional differences in average price levels, this would be exactly the right model. But this is not our goal. We want to predict cross-sectional differences in the likelihood of a future bubble.

<sup>6</sup> The choice of  $1 \times n$  rather than  $\omega \times n$  for  $\omega > 0$  is without loss of generality (see Section A.1 in the online appendix). The key assumption is that the rate at which excited-speculators calm down must be less sensitive to changes in the size of the excited-speculator crowd, n, and the asset's past returns, r, than the rate at which apathetic speculators get excited. The functional forms in Equations (9) and (10) are the simplest way to model this assumption. It is also possible to incorporate stochastic fluctuations in these population dynamics (see Section A.2 in the online appendix).

<sup>7</sup> Standard texts (e.g., Arnol'd 2012) show that  $n_{\tau}(n_0, \theta, r)$  is unique for all  $\tau \ge 0$  and  $n_0 \in [0, 1)$  because  $\theta \cdot r \cdot (1 - n) \cdot n - n$  is continuously differentiable on an open interval containing [0, 1).

<sup>8</sup> More formally, we say that  $\bar{n} \in SS(\theta, r)$  is stable if for every  $\delta > 0$ , there is some  $\epsilon > 0$  such that  $|n_{\tau}(n_0, \theta, r) - \bar{n}| < \delta$  for all  $\tau \ge 0$  given any initial population  $n_0 \in (\bar{n} - \epsilon, \bar{n} + \epsilon)$ .

<sup>9</sup> It is also not essential that  $\theta$  be strictly constant over time. All that matters is that  $\theta$  varies on a much slower timescale than returns. And I verify that this is true empirically in Section 3.3.

<sup>10</sup> See "China Blowing Major Bubbles in 2017" at https://www. forbes.com/sites/kenrapoza/2016/12/19/china-bubble-economy-2017/?sh=61bdec704745 (*Forbes*, December 19, 2016).

<sup>11</sup> Industries are labeled as in Fama and French (1997): (1) Aero: aircraft; (2) Agric: agriculture; (3) Autos: automobiles and trucks; (4) Banks: banking; (5) Beer: beer and liquor; (6) BldMt: construction materials; (7) Books: printing and publishing; (8) Boxes: shipping containers; (9) BusSv: business services; (10) Chems: chemicals; (11) Chips: electronic equipment; (12) Clths: apparel; (13) Cnstr: construction; (14) Coal: coal; (15) Drugs: pharmaceutical products; (16) ElcEq: electrical equipment; (17) FabPr: fabricated products; (18) Fin: trading; (19) Food: food products; (20) Fun: entertainment; (21) Gold: precious metals; (22) Guns: defense; (23) Hardw: computers; (24) Hlth: healthcare; (25) Hshld: consumer goods; (26) Insur: insurance; (27) LabEq: measuring and control equipment; (28) Mach: machinery; (29) Meals: restaurants, hotels, and motels; (30) MedEq: medical equipment; (31) Mines: nonmetallic 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.

 $^{12}$  Section B.1 in the online appendix shows that the paper's main results are robust to defining a speculative bubble using return thresholds other than 100% and -40%.

<sup>13</sup> See http://www.econ.yale.edu/~shiller/data.htm to download these data.

<sup>14</sup> Manela (2014) also uses media coverage to proxy for information diffusion after Food and Drug Administration drug approval.

<sup>15</sup> https://sraf.nd.edu/loughranmcdonald-master-dictionary/.

# References

- Abreu D, Brunnermeier MK (2003) Bubbles and crashes. Econometrica. 71(1):173–204.
- Admati AR (1985) A noisy rational-expectations equilibrium for multi-asset securities markets. *Econometrica: J. Econom. Soc.* 629–657.
- Arnol'd V (2012) Geometrical Methods in the Theory of Ordinary Differential Equations (Springer Science & Business Media, New York).
- Bailey M, Cao R, Kuchler T, Stroebel J (2018) The economic effects of social networks: Evidence from the housing market. J. Political Econom. 126(6):2224–2276.
- Barberis N (2018) Psychology-based models of asset prices and trading volume. Bernheim BD, DellaVigna S, Laibson D, eds. Handbook of Behavioral Economics: Applications and Foundations 1, vol. 1 (North-Holland, Amsterdam), 79–175.
- Barberis N, Thaler R (2003) A survey of behavioral finance. Constantinides GM, Harris M, Stulz RM, eds. *Handbook of the Economics of Finance 1* (Elsevier, Amsterdam), 1053–1128.
- Barberis N, Greenwood R, Jin L, Shleifer A (2015) X-CAPM: An extrapolative capital asset-pricing model. *J. Financial Econom.* 115(1):1–24.
- Barberis N, Greenwood R, Jin L, Shleifer A (2018) Extrapolation and bubbles. J. Financial Econom. 129(2):203–227.

- Brunnermeier MK, Nagel S (2004) Hedge funds and the technology bubble. J. Finance. 59(5):2013–2040.
- Burnside C, Eichenbaum M, Rebelo S (2016) Understanding booms and busts in housing markets. J. Political Econom. 124(4):1088–1147.
- Bursztyn L, Ederer F, Ferman B, Yuchtman N (2014) Understanding the mechanisms underlying peer effects: Evidence from a field experiment on financial decisions. *Econometrica*. 82(4):1273–1301.
- Case KE, Shiller RJ, Thompson AK (2012) What have they been thinking? Homebuyer behavior in hot and cold markets. *Brook*ings Papers Econom. Activity, 265.
- Cutler DM, Poterba JM, Summers LH (1990) Speculative dynamics and the role of feedback traders. *Amer. Econom. Rev.* 80(2):63–68.
- Daniel K, Hirshleifer D, Subrahmanyam A (1998) Investor psychology and security market under- and overreactions. J. Finance 53(6):1839–1885.
- De Long B, Shleifer A, Summers L, Waldmann R (1990) Positivefeedback investment strategies and destabilizing rational speculation. J. Finance. 45(2):379–395.
- Fama EF, French KR (1997) Industry costs of equity. J. Financial Econom. 43(2):153–193.
- Glaeser E, Nathanson C (2017) An extrapolative model of houseprice dynamics. J. Financial Econom. 126(1):147–170.
- Greenwood R, Shleifer A, You Y (2019) Bubbles for Fama. J. Financial Econom. 131(1):20–43.
- Gromb D, Vayanos D (2002) Equilibrium and welfare in markets with financially constrained arbitrageurs. J. Financial Econom. 66(2–3):361–407.
- Gromb D, Vayanos D (2010) Limits of arbitrage. Annual Rev. Financial Econom. 2(1):251–275.
- Guckenheimer J, Holmes P (2013) Nonlinear Oscillations, Dynamical Systems, and Bifurcations of Vector Fields (Springer Science & Business Media, New York).
- Hirsch M, Smale S, Devaney R (2012) *Differential Equations, Dynamical Systems, and an Introduction to Chaos* (Elsevier, Amsterdam).
- Hirshleifer D (2020) Presidential address: Social transmission bias in economics and finance. J. Finance. 75(4):1779–1831.
- Hong H, Stein J (1999) A unified theory of underreaction, momentum trading, and overreaction in asset markets. *J. Finance*. 54(6):2143–2184.
- Hong H, Kubik J, Stein J (2004) Social interaction and stock-market participation. J. Finance. 59(1):137–163.
- Kaustia M, Knüpfer S (2012) Peer performance and stock market entry. J. Financial Econom. 104(2):321–338.
- Kindleberger C (1978) Manias, Panics, and Crashes: A History of Financial Crashes (Basic Books, New York).
- Kuznetsov Y (2013) *Elements of Applied Bifurcation Theory* (Springer Science & Business Media, New York).
- Loughran T, McDonald B (2011) When is a liability not a liability? textual analysis, dictionaries, and 10-ks. J. Finance. 66(1):35–65.
- Mackay C (1841) Extraordinary Popular Delusions and the Madness of Crowds (Richard Bentley, London).
- Manela A (2014) The value of diffusing information. J. Financial Econom. 111(1):181–199.
- Miller E (1977) Risk, uncertainty, and divergence of opinion. J. Finance. 32(4):1151–1168.
- Minsky H (1970) Financial instability revisited: The economics of disaster. Board of Governors of the Federal Reserve System working paper, Federal Reserve System, Washington, DC.
- Mullainathan S, Shleifer A (2005) Market for news. Amer. Econom. Rev. 95(4):1031–1053.
- Ofek E, Richardson M (2001) Dotcom mania: The rise and fall of internet-stock prices. J. Finance. 58(3):1113–1137.
- Scheinkman J, Xiong W (2003) Overconfidence and speculative bubbles. J. Political Econom. 111(6):1183–1220.
- Sharpe WF (1964) Capital asset prices: A theory of market equilibrium under conditions of risk. J. Finance. 19(3):425–442.

- Shiller R (1984) Stock prices and social dynamics. *Brookings Papers Econom. Activity* 2:457–510.
- Shiller R (2000) Irrational Exuberance (Princeton University Press, Princeton, NJ).
- Shiller RJ, Pound J (1989) Survey evidence on diffusion of interest and information among investors. J. Econom. Behavior Organ. 12(1):47–66.
- Shive S (2010) An epidemic model of investor behavior. J. Financial Quant. Anal. 45(1):169–198.
- Shleifer A, Summers L (1990) The limits of arbitrage. J. Finance. 52(1):35–55.
- Shleifer A, Summers LH (1990) The noise trader approach to finance. J. Econom. Perspect. 4(2):19–33.
- Strogatz S (2014) Nonlinear Dynamics and Chaos: With Applications to Physics, Biology, Chemistry, and Engineering (CRC Press, Boca Raton, FL).
- Xiong W, Yu J (2011) The Chinese warrants bubble. Amer. Econom. Rev. 101(6):2723–2753.