Herding Cycles*

Edouard Schaal
CREI, UPF and BGSE

Mathieu Taschereau-Dumouchel
Cornell University

December 2020 – Preliminary

Abstract

This paper explores whether rational herding can generate endogenous business cycle fluctuations. We embed a tractable model of rational herding into a business-cycle framework. In the model, technological innovations arrive with unknown quality. New innovations are not immediately productive and agents have dispersed information about how productive the technology will be. Investors decide whether to invest in the technology or not based on their private information and the investment behavior of others. Herd-driven boom-bust cycles may arise endogenously in this environment out of a single impulse shock when the technology is unproductive but investors' initial information is optimistic and highly correlated. When the technology appears, investors mistakenly attribute the high observed investment rates to high fundamentals, leading to a pattern of increasing optimism and investment until the economy reaches a peak, followed by a crash as agents ultimately realize their mistake. As such, the theory can shed light on bubble-like episodes in which excessive optimism about uncertain technology fueled general macroeconomic expansions that were followed by sudden recessions. We calibrate the model to the U.S. economy and show that the theory can explain boom-and-bust cycles in line with historical episodes like the Dot-Com Bubble of the late 1990s. Leaning-against-the-wind policies can be beneficial in this environment as they improve the diffusion of information over the cycle.

JEL Classification: E32, D80

*E-mail: eschaal@crei.cat, mt763@cornell.edu. We thank Vladimir Asriyan, Andy Atkeson, Olivier Coibion, Pablo Fajgelbaum, Jordi Gali, Wouter den Haan, Nir Jaimovich, Francesco Lippi, Claudio Michelacci, Luigi Paciello, Ricardo Reis, Victor Rios-Rull, Nancy Stokey, Victoria Vanasco, Jaume Ventura, Jon Vogel, Pierre-Olivier Weill and Yi Wen for helpful comments. We thank seminar and conference participants at the University of Texas-Austin, Banco de Chile, Banco de Portugal, University of Chicago, CIGS Tokyo, Cornell University, CREI, University of Edinburgh, EIEF, Federal Reserve Banks of Philadelphia and St. Louis, HKUST, London School of Economics, Paris School of Economics, Society for Economic Dynamics 2018, Sveriges Riksbank, UCLA and the University of Zurich. Edouard Schaal acknowledges financial support from the Spanish Ministry of Economy, Industry and Competitiveness through the program Proyecto I+D Excelencia 2017 (grant ECO2017-87827-P-AEI/FEDER, UE) and through the Severo Ochoa Programme for Centres of Excellence in R&D (SEV-2015-0563).

1 Introduction

Several historical recessions were preceded by periods of massive investment in a new technology. One salient example is the boom in information technologies in the 1990s that culminated with the stock market crash of 2001 ("Dot-Com bubble"). While the internet had been invented years earlier to connect academic and military networks, its commercial potential only became clearer in the 1990s, which led to large investments in communication networks, softwares, and IT equipments. After a booming period with large capital inflows in the IT sector, the Dot-Com crash ensued as some of the expected returns failed to materialize. Other boom-bust episodes follow similar patterns. For instance, the Roaring Twenties, a period of massive economic growth fueled by technological innovations in many sectors such as car manufacturing, communication, aviation and the chemical industry, ended in the Great Depression.[1]

Standard practice in modern business cycle analysis often treats the booms and the busts as separate episodes, both driven by their own sequence of exogenous shocks. But the historical evidence suggests that some booms and busts might be, instead, intrinsically linked as two sides of the same coin. In that view, the "booms sow the seeds of the subsequent busts" (Beaudry et al., 2020), and recessions should be analyzed through the lens of the preceding expansion. Distinguishing what underlies these relationships is key to a deeper understanding of business cycle phenomena and the conduct of stabilization policy.

This paper proposes a theory that can generate an *endogenous* boom-and-bust cycle out of a single impulse shock, by which we mean a general economic expansion followed by a recession that falls below the trend.[2] To that purpose, we embed a model of herding into a standard business cycle framework. In the model, random technological innovations arrive over time and agents can decide whether to invest or not in the new technology. The payoff from investing is initially unknown and investors use all available information to update their beliefs about the fundamental value of the technology. Information comes from both public and private sources. Key to our mechanism is the assumption that private signals can feature common noise. This assumption captures the idea that investors collect information from similar sources (news media, market reports, etc), hence that beliefs across investors can be correlated for reasons unrelated to the fundamental value of the technology. Importantly, investors do not initially know the extent of that correlation but progressively learn about it over time.

Agents also receive public signals. First, they can learn by observing the exogenous return on their investment, which provides noisy information about the technology. They can also learn from

[1] Xiong (2013) documents several instances of boom-bust episodes that follow the introduction of new technologies.

[2] By "endogenous", we mean that the entire boom-and-bust pattern is produced by the forces in the model. Our theory still relies on shocks, however, but only one-time shocks and does not rely on a particular sequence of positive then negative shocks. This approach is different from other theories of endogenous business cycles that generate deterministic periodic or chaotic dynamics (see Boldrin and Woodford (1990); Benhabib (1992); Guesnerie and Woodford (1992) for surveys).

endogenous market outcomes such as aggregate quantities or prices. In the model, this amounts to observing, with some noise, the mass of agents who invest in the new technology. As the individual investment decisions reflect the private information of the agents, this public signal operates as a *social learning* channel by aggregating some of the dispersed information for everyone to see.

How agents interpret this public signal is key for the emergence of boom-bust cycles. Such cycles are caused in our model by what we refer to as “false-positives”: low realizations of the technology fundamental accompanied by unusually large and positive realizations of the common noise. When observing the large amount of investment induced by such false-positive shocks, agents infer that private signals are positive. These signals, in turn, can be positive either because the fundamental value of the technology is high, or because the correlated component of the private signals is high enough to fool agents into believing that the technology is good. Investors cannot tell these stories apart but, given that the false-positive shock is deemed unusual at first, update their beliefs by increasing the likelihood of the high-technology state. More optimistic beliefs lead to further aggregate investment next period, which, in turn, leads to even more positive beliefs about the fundamental and so on. Through this positive feedback loop, the arrival of a low-value technology can create a long-lasting boom as investors are fooled by the initial investment craze.

But agents are rational and understand the possibility that they can sometimes be mistaken. As a result, they keep track in the background of the probability of being in a false-positive state, which appears increasingly likely as time unfolds. At some point, the most pessimistic agents stop investing and aggregate investment no longer suggests a high technology draw, leading to sharp reversal in beliefs and a collapse of investment. We provide formal conditions under which these boom and bust episodes are guaranteed to arise.

Several features of these boom-bust cycles should be highlighted. First, they are not driven by a specific sequence of shocks. Instead, a unique initial draw is responsible for shaping the whole cycle. The crash, in particular, is not triggered by an exogenous shock but arises endogenously through the natural evolution of the model. In addition, the properties of the boom—its duration and magnitude—have a causal impact on the properties of the bust. More generally, the model provides a mechanism to generate bubble-like phenomena over the business cycle, which can be used for quantitative and policy analysis.

Through the observation of endogenous public signals, our model generates a form of *herding*: During the expansion phase, agents mistakenly follow the herd into an investment boom and a diminishing measure of agents use their private information to go against the crowd. In that sense, our paper relates to the original work by Banerjee (1992), Bikhchandani et al. (1992) and Chamley (2004) on herding and information cascades. The model we propose is, however, distinct along several dimensions. In previous herding models, decisions are made sequentially and idiosyncratic shocks govern the model dynamics. Both of these features, we believe, have prevented the introduction of herding features into macroeconomic models, with the exception of Loisel et al. (2012)

which we discuss below. In our model, instead, agents act simultaneously and learn from observing aggregates, providing a simpler mapping into standard macroeconomic models. Second, boom-bust cycles arise in traditional herding models only for specific sequences of shocks. In our model instead, because the common noise adds a dimension of uncertainty , the boom-bust pattern emerges endogenously in response to a single shock.

In the model, the amount of information that agents receive is endogenous and varies with the cycle, which opens the door to a form of *information cascades*. When the public signals received up to a certain date are very positive, most agents invest regardless of their private signals so that their private information is not encoded into the public signal. Similarly, after a series of bad news, many agents behave as if they disregarded their private information when deciding not to invest, making aggregate investment also uninformative. As a result, the model is able to generate sustained booms and rapid busts as periods with massive investment restrict the flow of information, but slight downturns can suddenly reveal more information on the true state of the world.

Due to this variable flow of public information, the model features an information externality: Agents do not internalize how their private investment decisions affect the flow of public information. We characterize the solution of a social planning problem and show that the planner pushes agents to invest less during booms and more during downturns so as to optimize the amount of information provided by aggregate investment. We also characterize the optimal investment tax that implements the efficient allocation and show that it displays a *leaning-against-the-wind* characteristic with investment taxes during booms and investment subsidies during downturns while the technology is uncertain.

To explore how the evolution of beliefs generated by our learning model can produce a general macroeconomic expansion followed by a recession, and to have a sense of the magnitude of the boom-bust cycles generated by the theory, we embed our main mechanism into a basic model of the business cycle, which models the technology adoption decision of entrepreneurs after the arrival of a new technology. The model features two types of capital, traditional and information technology capital (IT), and we assume that the new technology is more intensive in IT capital. As in the basic model, there is a form of social learning as agents learn from observing the measure of new-technology adopters. We use the New Keynesian framework as a backbone for two reasons. First, nominal rigidities provide a way to obtain positive comovements across macroeconomic aggregate out of belief shocks. Second, it allows us to discuss the implications of the model for monetary policy.

We provide a back-of-the-envelope calibration of the model to match various moments of the data that relate to the dot-com period. In particular, we discipline the amount of private information—a key moment for our mechanism—using dispersion in forecasts from the Survey of Professional Forecasters (SPF). We also use data from the SPF to discipline investors beliefs about the true value of the technology. Under our calibration, the model is able to generate a boom-bust cycle with positive comovement in consumption, investment, hours worked and output. The over-investment into IT capital during the boom period causes the economy to contract significantly when beliefs collapse as agents realize that resources were misallocated.

We also discuss our model's implications for the conduct of monetary policy in the face of boom-and-bust cycles and investigate whether a leaning-against-the-wind monetary policy would be desirable. We find that monetary policy can dampen the cycle but has little effect on the technology choice of the entrepreneurs and on the release of public information, in contrast to a technology-adoption tax. The downside of these policies is that they also slow down the adoption of good technologies.

1.1 Literature Review

Our paper relates to the original work on herding and information cascades of Banerjee (1992), Bikhchandani et al. (1992) and Chamley (2004). Our model differs from these traditional models of herding in the fact that agents act simultaneously and learn from observing the average action of other players. While most of the earlier literature focus on a strict type of permanent information cascade in which the flow of information is zero, our model features a "smooth" form of information cascades because the information flow varies smoothly over the cycle. This feature is important as it allows for the possibility of the economy endogenously exiting the cascade region, producing the busts in our simulations. Our learning model is closer to Vives (1997) who studies an environment in which agents with dispersed information learn by observing the average action across agents.[3] Chapter 4 of Chamley (2004) briefly reviews a model close to ours in which privately informed agents learn from the average action. As in our model, the amount of information released by the public signal varies with the public beliefs. Our model also relates to Avery and Zemsky (1998), who study herding in financial market and introduce multidimensional uncertainty to maintain the existence of information cascades. Cipriani and Guarino (2008) show that herding can occur in financial markets when investors have heterogeneous valuation for the asset.

To our knowledge, Loisel et al. (2012) is the only other macroeconomic model that features herding phenomena. Their paper presents a simple general equilibrium model with overlapping generations of finitely-lived entrepreneurs who are endowed with private signals and act sequentially to invest in a risky asset. As in traditional models of herding, individual adoption decisions are publicly observable and aggregate output in every period reflects investors' idiosyncratic noise. Boom-and-bust cycles arise in this model for specific exogenous sequences of shocks only.

This paper also relates to a literature in which the endogenous release of information generates sudden collapses in economic activity, as in Caplin and Leahy (1994) and Veldkamp (2005). It also relates to models in which the aggregation of private information leads to nonlinear aggregate

[3]Both Vives (1993) and Vives (1997) show that learning happens slowly in models in which privately informed agents learn from aggregates.

dynamics (Fajgelbaum et al., 2017).

A large literature on financial bubbles consider models in which an asset trades above its fundamental value (Samuelson, 1958; Tirole, 1985). In contrast, while our model provides a theory why bubble-like phenomena may endogenously emerge and burst, the equilibrium is unique and our model provides no distinction between fundamental versus "bubbly" equilibria. In that respect, our work is closer to Abreu and Brunnermeier (2003) who analyze how bubbles burst. Our work also relates to a strand of literature that studies the role of policy in environments with asset bubbles (Galí, 2014; Martin and Ventura, 2016; Asriyan et al., 2019).

Our paper is closely related to the literature on news or noise-driven business cycles (Beaudry and Portier, 2004, 2006; Lorenzoni, 2009; Jaimovich and Rebelo, 2009). Indeed, our model shares the view that boom-bust cycles may be due to false-positives. In the news-shock literature, beliefs are driven by the exogenous release of news at fixed dates. In contrast, in our model, the rise and fall in beliefs that generate boom-and-bust cycles is endogenously driven by model forces, allowing us to explore the model's predictions on the frequency and timing of such cycles and providing a greater role for stabilization policies. Christiano et al. (2008) considers the interaction of monetary policy and boom-bust cycles driven by news shocks. Closer to our work, Benhima (2019) builds a two-period model with dispersed private information in which an overly optimistic news shock about demand can create a boom in period 1 and a bust in period 2 when truth is revealed.

Finally, our model predicts the emergence of recurring aggregate cycles. Recent work by Beaudry et al. (2020) documents the existence of such cycles in U.S. data.

Section 2 introduces a simple version of our learning model that conveys the intuition of the mechanism. Section 3 presents our business cycle model. We calibrate our model in Section 4 and show several empirical implications of the mechanism. We also discuss the role of policy. The final section concludes.

2 Learning Model

We start by presenting our mechanism in a simplified dynamic investment game. This allows us to provide intuition for why social learning can lead to an endogenous herd-driven boom-bust cycle out of a single impulse shock. It also permits us to derive additional analytical results and discuss the model's optimal policy implications.

2.1 Notation

In what follows, whenever $F^x(\tilde{x}) = Pr(x \leq \tilde{x})$ denotes the cumulative distribution function (CDF) of some random variable x, f^x refers to its associated probability density function and $\overline{F}^x$ its complementary CDF, $\overline{F}^x(\tilde{x}) = Pr(x > \tilde{x})$.

2.2 Environment

Time is discrete and goes to infinity. The economy is populated by a unit measure of investors indexed by $j \in [0,1]$. Investors are risk-neutral and discount future consumption at rate $0 < \beta < 1$. Each investor has access to an investment technology that becomes available in period 0 and provides a period return

$$R_t = \theta + u_t,$$

that is identical across agents and where $\theta \in \{\theta_H, \theta_L\}$, $\theta_H > \theta_L$, is the permanent component of the technology, and u_t is an i.i.d transitory component drawn from the cumulative distribution F^u. Since u_t is an i.i.d component that agents cannot forecast, we refer to θ as the technology *fundamental* that agents try to learn about. Every period, investors must decide whether to invest in the technology ($i_{jt} = 1$) or not ($i_{jt} = 0$). We start with the assumption that the investment decision is binary but will relax it later in an extension. Investing is costly and requires the payment of a cost c, identical across agents, every period. We assume that agents have deep pockets and ignore any form of budget or financial constraints. The total return to an investor j in any given period t is therefore

$$y_{jt} = i_{jt} \left(R_t - c \right).$$

2.3 Information

The permanent component θ of the investment technology is randomly drawn once and for all at date 0. We denote by p_0 the ex-ante probability that $\theta = \theta_H$. Investors cannot observe θ directly but receive various private and public signals about it.

Private signals

First, we assume that agents receive a private signal s_j at date 0, upon the arrival of the new technology. Importantly, we allow these private signals to feature not only idiosyncratic noise but also common noise. Common noise in private information can be justified by the fact that agents use similar sources of information (mass media, internet) that may report noisy signals about the initial success of the investment technology (e.g., benchmark tests). Common noise is key to our mechanism as it introduces the possibility that the average belief about θ in the economy may vary for reasons orthogonal to the true value of the fundamental. In other words, common noise is what allows agents to be sometimes overly optimistic or pessimistic about the technology. It is crucial for investors to mistakenly attribute high investment patterns to high fundamentals even though the technology is bad in reality, fueling the initial stage of the boom and bust cycle.

This common noise is captured by the random variable ξ, distributed according to the CDF F^{ξ}. Formally, we assume that the private signal s_j is drawn from the CDF $F^s_{\theta+\xi}(s) = Pr\left(s_j \leq s\right)$, where $\{F_x^s\}_{x\in I}$ is a family of distributions. To prevent the possibility of trivial learning, we make the assumption that F_x^s has full support over $\mathbb{R}$, i.e., $f_{\theta+\xi}^s > 0$ everywhere. Finally, in order to guarantee some monotonicity in learning, we assume that the family $\{F_x^s\}_{x\in I}$ satisfies the *monotone likelihood ratio property* (MLRP). That is, for $x_1 < x_2 \in I$ and $s_1 < s_2$, we must have

$$\frac{f_{x_2}^s(s_2)}{f_{x_1}^s(s_2)} \geq \frac{f_{x_2}^s(s_1)}{f_{x_1}^s(s_1)}. \tag{MLRP}$$

Intuitively, the MLRP condition guarantees that a high signal s is more likely to be coming from a high realization of $x = \theta + \xi$. In other words, an investor observing a high private signal s_j will become more optimistic and will put a higher probability on the value of the technology θ, or the common noise ξ, being high.

Example. In most of our examples below, we will use *additive private signals* so that

$$s_j = \theta + \xi + v_j, \text{ with } v_j \sim \text{iid CDF } F^v. \tag{1}$$

Well-known distributions that satisfy the MLRP condition include the exponential, binomial, poisson or the Gaussian distributions.

Public signals

In addition to observing their initial private signal, investors collect public information over time by observing market activity and returns. We first assume that all agents observe the return on investment R_t. Since the transitory component in the investment return u_t is unknown to agents, the total return R_t offers an exogenous noisy signal about θ, which provides a constant amount of information over time. Second, and more importantly for our mechanism, we introduce a form of *social learning* in the economy by allowing investors to observe an endogenous signal which partially aggregates agents' private information. Specifically, we assume that investors receive a noisy measure of the total number of investing agents m_t, which we define as follows:

$$m_t = \int_0^1 i_{jt} dj + \varepsilon_t, \text{ with } \varepsilon_t \sim \text{iid CDF } F^\varepsilon. \tag{2}$$

The noise ε_t can be interpreted as coming either from measurement error or from the presence of noise traders that make aggregate variables less informative. The presence of noise is required in our setting to prevent agents form learning too quickly (or even immediately in some cases, as we discuss later). How much agents learn from investment activity in reality is a quantitative matter that we discuss in section 4.

Before going further, we would like to highlight some unusual features of signal m_t. In equilibrium, the decision to invest i_{jt} will be a nonlinear function of the investor's individual beliefs. In turn, beliefs will be a function of public information up to time t, $\{R_{t-1}, m_{t-1}, \dots, R_0, m_0\}$ and the private signal s_j. As a result, since public information is shared and can be filtered out to some extent, m_t partially aggregates the private information across the population of investors. To that extent, m_t has a useful information content regarding the fundamental θ and the common noise ξ. Versions of an endogenous signal like m_t have been studied in the literature, but usually under an assumption of linearity (Vives, 1993, 1997; Amador and Weill, 2012). What makes the signal m_t particularly interesting in our setup is that it is nonlinear. Hence, the amount of information contained in this signal will vary over time depending on the location in the state space (e.g., history of shocks), opening up the possibility of a form of *informational cascades* in which agents neglect their private information and aggregate investment becomes uninformative. Finally, because m_t is endogenous, its information content will react to changes in the environment providing a basis for government intervention.

2.4 Belief Characterization

There are two aggregate shocks in this economy: the fundamental θ and the common noise ξ. The beliefs of an individual investor j are described by a joint probability distribution that we denote by

$$\pi_{jt}\left(\tilde{\theta}, \tilde{\xi}\right) = Pr\left(\theta = \tilde{\theta}, \xi \in \left[\tilde{\xi}, \tilde{\xi} + d\tilde{\xi}\right] \mid \mathcal{I}_{jt}\right),$$

in which we explicitly allow for ξ to take a continuum of values and where $\mathcal{I}_{jt}$ is agent j's information set at date t. Since investors receive different private signals, we should in principle keep track of the whole distribution of beliefs in the economy (i.e., a distribution over distributions). Fortunately, the information structure is simple enough that the model lends itself to a great simplification. As in Chapter 3 of Chamley (2004), it is enough to keep track of only one set of time-varying beliefs, the *public beliefs* $\pi_t\left(\tilde{\theta}, \tilde{\xi}\right) = Pr\left(\theta = \tilde{\theta}, \xi \in \left[\tilde{\xi}, \tilde{\xi} + d\tilde{\xi}\right] \mid \mathcal{I}_t\right)$, that correspond to the beliefs of an outside observer who only has access to public information $\mathcal{I}_t$ at time t, which is the collection of past investment returns and measures of investors:

$$\mathcal{I}_t = \{R_{t-1}, m_{t-1}, \dots, R_0, m_0\}.$$

In comparison to this outside observer, an investor's information set includes in addition the private signal s_j, so that $\mathcal{I}_{jt} = \mathcal{I}_t \cup \{s_j\}$. Investor's individual beliefs can easily be recovered from public beliefs using Bayes' rule and the private signal s_j, according to

$$\pi_{jt}\left(\tilde{\theta}, \tilde{\xi}\right) = \frac{\pi_t\left(\tilde{\theta}, \tilde{\xi}\right) f^s_{\tilde{\theta}+\tilde{\xi}}(s_j)}{\int \pi_t(\theta, \xi) f^s_{\theta+\xi}(s_j)\, d(\theta, \xi)}. \tag{3}$$

This simplification comes from the fact that only public information evolves over time. Indeed, since the private signal distribution $f^s_{\theta+\xi}$ is constant and known up to the realization of θ and ξ, it is easy to recover the entire distribution of private beliefs across investors for a given combination of (θ, ξ) at any point in time. As a result, the only object that we need to keep track of is the public belief function π_t.

2.5 Timing and Investment Decision

The timing is as follows. At date 0, the fundamental θ, the common noise component ξ and the private signals s_j are drawn once and for all. At date $t \geq 0$,

1. Each agent chooses whether to invest or not based on the individual beliefs π_{jt},
2. Investment returns are realized,
3. All agents observe $\{R_t, m_t\}$, update their beliefs and move to the next period.

The investment decision can be characterized in an easy way. Because returns accrue in the same period as the investment is made, the investment decision is a simple static problem. Investor j invests in period t if and only if

$$E\left[R_t \mid \mathcal{I}_{jt}\right] \geq c. \tag{4}$$

Defining

$$p_{jt} = Pr\left(\theta = \theta_H \mid \mathcal{I}_{jt}\right) = \int \pi_{jt}\left(\theta_H, \xi\right) d\xi \tag{5}$$

as the probability that investor j puts on being in the high technology state, the investment decision (4) is characterized by a cutoff rule $\hat{p}$ in the space of beliefs. That is, an agent invests if and only if $p_{jt} \geq \hat{p}$ where $\hat{p}$ is the belief of the marginal investor, such that

$$\hat{p}\theta_H + (1 - \hat{p})\,\theta_L = c. \tag{6}$$

The total measure of investing agents can then be expressed as

$$m_t = m^e\left(\pi_t, \theta, \xi\right) + \varepsilon_t \tag{7}$$

$$\text{where } m^e\left(\pi_t, \theta, \xi\right) = \int \mathbb{1}\left(p_j\left(\pi_t, s_j\right) \geq \hat{p}\right) f^s_{\theta+\xi}\left(s_j\right) ds_j. \tag{8}$$

The variable m^e is the expected measure of investing agents for a given state of the world (θ, ξ), excluding the noise traders. Importantly for what follows, m^e is an object that any agent in the economy can compute. Knowing the structure of the model, all agents agree on the cutoff $\hat{p}$. Second, thanks to the dichotomy between public beliefs and a fixed distribution of private signals $f^s_{\theta+\xi}$, all agents can compute the distribution of beliefs p_j given a realization of θ and ξ. This property is essential to tractably solve the inference problem to which we now turn.

2.6 Evolution of Beliefs

After characterizing the investment decision, we may now describe how beliefs are updated over time. Each end of period brings two new public signals for investors to process: R_t and m_t. The updating of information with R_t is straightforward as it is a simple exogenous signal. Applying Bayes' rule, we define the interim beliefs at the end of the period as

$$\pi_{t|R_t}\left(\tilde{\theta},\tilde{\xi}\right)=\frac{\pi_t\left(\tilde{\theta},\tilde{\xi}\right)f^u\left(R_t-\tilde{\theta}\right)}{\int\pi_t\left(\theta,\xi\right)f^u\left(R_t-\theta\right)d\left(\theta,\xi\right)}. \tag{9}$$

We now turn to incorporating the information contained in m_t. Solving the inference problem from an endogenous signal like m_t can be complicated in general because individual decisions need to be inverted to back out their information content about θ and ξ. Fortunately, and as highlighted at the end of the previous section, the inference problem is greatly simplified in our environment since the expected measure of investors m^e in every state of the world is a simple function of public beliefs π_t (known by everyone) and of the true realization of (θ,ξ). Investors solely differ in their assessment of the probability of each state (θ,ξ), encoded in π_{jt}, but there is no *infinite regress* problem arising from the necessity to forecast the beliefs of agents after any history of shocks.[4] Because of the equilibrium structure of signal (7), Bayes' rule gives us the simple updating equation

$$\pi_{t+1}\left(\tilde{\theta},\tilde{\xi}\right)=\frac{\pi_{t|R_t}\left(\tilde{\theta},\tilde{\xi}\right)f^\varepsilon\left(m_t-m^e\left(\pi_{t|R_t},\tilde{\theta},\tilde{\xi}\right)\right)}{\int\pi_{t|R_t}\left(\theta,\xi\right)f^\varepsilon\left(m_t-m^e\left(\pi_{t|R_t},\theta,\xi\right)\right)d\left(\theta,\xi\right)}. \tag{10}$$

2.7 Equilibrium

We are now ready to define an equilibrium in this economy.

Definition 1. An equilibrium consists of public beliefs π_t for all t, a distribution of private beliefs $\{p_{jt}\}_{j\in[0,1]}$ for all t, and an expected mass of investors m_t^e for all t, such that, given shocks, 1) the distribution of private beliefs is derived from the public beliefs through (3) and (5); 2) the expected mass of investors is consistent with investors decisions under their private beliefs as in (8); and 3) the public beliefs follow their law of motion (9)–(10) with m_t given by (7).

With that definition in hand, the following proposition characterizes the set of equilibria.

Proposition 1. *There exists a unique equilibrium.*

[4] In the absence of the simplifications from our information structure, learning from m_t would require to compute a hypothetical m_t and its probability in every state of the world after every history of shocks. Computing m_t, in turn, would require forecasting the beliefs of each individual at each date — themselves being the product of a sequence of individual inference problems. Townsend (1983) provides a famous example why this sort of inference often leads to an intractable infinite regress problem.

The proof of the proposition is straightforward. It shows that from a given distribution of public beliefs π_t, there is a unique mapping, given the realization of the shocks, to next period's public beliefs π_t. Starting from the initial π_0 we can therefore reconstruct the unique equilibrium sequence $\{\pi_0, \pi_1, \dots\}$. All other equilibrium quantities such as the measure of investors and distribution of private beliefs can be reconstructed from the public beliefs in a unique way.

2.8 Example: the 3-state model

We are now fully equipped to analyze the dynamics implied by the model. We start with a simple special case that conveys the intuition about the emergence of i) a smooth form of information cascades and ii) endogenous booms and busts.

We temporarily make the simplifying assumption that the pair (θ, ξ) can only take three different values, the minimal number of states required for endogenous boom-bust cycles to emerge in our model. Specifically, we assume

$$(\theta, \xi) \in \{(\theta_L, 0), (\theta_H, 0), (\theta_L, \Delta)\} \text{ with } \theta_L < \theta_L + \Delta < \theta_H.$$

We refer to $(\theta_L, 0)$ as the *low-technology* state, $(\theta_H, 0)$ as the *high-technology* state and (θ_L, Δ) as the *false-positive* state. The latter is the one of interest as it is the state that will trigger a boom-and-bust cycle by having investors mistakenly assess the technology to be of high quality before later realizing their mistake.

Having only three states reduces the number of state variables required to keep track of the belief distribution π_t. Public beliefs are now summarized by two variables

$$p_t \equiv \pi_t (\theta_H, 0) \text{ and } q_t \equiv \pi_t (\theta_L, \Delta),$$

and the corresponding updating rules can be found in Appendix A.1.

We now establish a first simple result. Under our assumptions, the individual beliefs about the probability of the high technology, $p_{jt} = \pi_{jt} (\theta_H, 0)$, is increasing in the private signal s_j. As a result, the investment decision can be further characterized by a cutoff rule $\hat{s}(p_t, q_t)$ in terms of private signals, which simplifies the expression of the expected measure of investing agents m^e as the following Lemma shows.

Lemma 1. *In the three-state model, for $\theta_L < \theta_L + \Delta < \theta_H$ and $\{F_x^s\}$ satisfying the MLRP condition, the optimal investment strategy is characterized by a cutoff rule in the private signal $\hat{s}(p_t, q_t)$, decreasing in p_t. That is, an agent invests if and only if $s_j \geq \hat{s}(p_t, q_t)$. The expected measure of investing agents is given by*

$$m^e (p_t, q_t, \theta, \xi) = \overline{F}^s_{\theta+\xi} (\hat{s}(p_t, q_t)).$$

Learning from m_t

To develop intuition on the way learning from the measure of investors works in this economy, we propose an example in Figure 1. Panel (a) displays the distribution of private signals s_j in the three states of the world. Due to the MLRP assumption, the three distributions are ordered in a first-order stochastic dominance sense. The expected measure of investing agents m^e is represented as the mass of agents located to the right of cutoff $\hat{s}_t$. We can see that m^e is small in the low-technology state $(\theta_L, 0)$ (in red), that agents expect more investment in the false-positive state (θ_L, Δ) (in green), and that it is at its largest in the high-technology state $(\theta_H, 0)$ (in blue).

The three measures m^e being computed, we then present in panel (b) the three potential distributions of m_t in the three states of the world assuming that the noise ε is normally distributed with mean 0. As the graph illustrates, agents expect very different distributions of investment m_t, each centered on their expected value m^e, in the different states of the world (θ, ξ). We can split the m_t-space into three regions that indicate which state is attributed more probability after observing m_t. For instance, for low m_t the likelihood of state θ_L is greater than that of the other states, so information updating will attribute it a higher probability. The two other states, (θ_L, Δ) and $(\theta_H, 0)$, have their own higher likelihood region that are also represented on the graph. Importantly for the emergence of our boom-and-bust cycles, beliefs about the high state will tend to increase after observing high realizations of m_t. It is in that sense that the model displays a form of "herding": agents become more optimistic (resp. pessimistic) after seeing high (resp. low) patterns of investment, leading them to make inefficient investment decisions, as we will see in our welfare analysis.

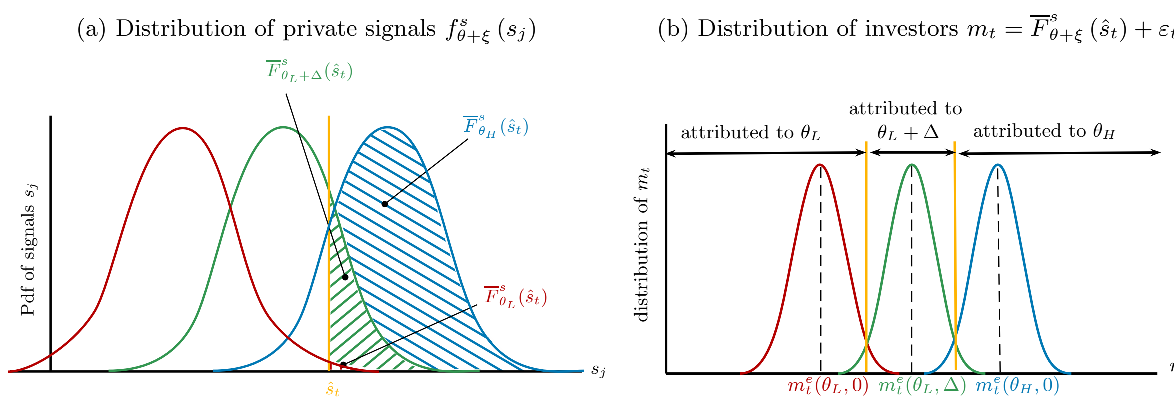

Notes: Panel (a) on the left displays the distribution of private signals s_j across the population in the three possible states of the world along with the corresponding expected measures of investing agents $m^e_t = \overline{F}^s_{\theta+\xi}(\hat{s}(p_t, q_t))$ for some public beliefs (p_t, q_t). Panel (b) on the right shows the distribution of $m_t = m^e_t + \varepsilon_t$ in the three states of the world assuming some Gaussian-like distribution F^ε with mean 0 and variance σ^2_ε.

Figure 1: Private beliefs and expected measure of investors

Signal-to-noise ratio and smooth information cascades

In the traditional herding literature (Banerjee, 1992; Bikhchandani et al., 1992), information cascades arise when public beliefs are so extreme (p_t extremely high or low, because of a particular sequence of investment decisions), that agents end up "disregarding" their own private information. That is, agents invest (or not) no matter what their private information is. As a result, observing previous investors' decisions becomes uninformative and the economy may end up being stuck in a situation with wrong beliefs forever.

Because social learning takes place through the observation of the continuous variable m_t rather than the sequence of binary decisions by previous investors, the emergence of information cascades is somewhat different in our setup. We show nonetheless that a similar form of "smooth" information cascades may arise depending on assumptions about signal distributions.

The bottom-right panel of Figure 2 represents how the measure of investing agents m_t varies in expectation, along with its ± 1-standard deviation error bands, as a function of the public belief p_t, holding q_t constant in the background. These curves are drawn by first connecting a given level of p_t in the bottom-right panel to the equilibrium signal threshold $\hat{s}\left(p_t, q_t\right)$ (upper-right panel), itself connected to the upper-left panel which shows how the measures $m^e = \overline{F}^s_{\theta+\xi}\left(\hat{s}\right)$ vary with the cutoff $\hat{s}$. As the bottom panel shows, the expected measure of investing agents m^e is a monotonic transformation of the CDF $F^s_{\theta+\xi}$ in the three different states.

The key feature to take away from this graph is that the *signal-to-noise* ratio in m_t varies nonmonotonically with the public beliefs. For intermediate values of p_t, the three expected measures m^e are far apart so that despite the noise ε_t, observing m_t is highly informative about the underlying state $\theta + \xi$ (i.e., the signal-to-noise ratio is high). For p_t large (resp. small) region, almost all (resp. no) agents invest, the three measures converge to $\lim_{\hat{s}\to\infty} \overline{F}^s_{\theta+\xi}\left(\hat{s}\right) = 1$ (resp. 0), so that the signal m_t is dominated by noise and becomes uninformative about the underlying fundamentals (i.e., the signal-to-noise ratio is low). Note that this result is not an artifact of specific distributions or functional forms but is instead a general feature of the model as long as $\hat{s}$ varies sufficiently on the support of $\overline{F}^s_{\theta+\xi}$.

The model offers a continuous and smooth analog to informational cascades when the equilibrium $\hat{s}$ reaches the extreme regions of the state space where learning is slow. Suppose for instance that public beliefs are optimistic(p_t high) so that $\hat{s}$ is very low. In such a situation, almost all agents act in the same way and invest in the new technology. Only few agents use their private information to "go against the crowd" and do not invest: the most pessimistic ones that have received a particularly low private signal. Unfortunately, their measure is so low that they are hard to detect when looking at the aggregate investment patterns. As a result, markets are nearly uninformative and beliefs can remain wrong for an extended period of time. The main difference with traditional herding models is that, under the assumption that private signals have full unbounded support,

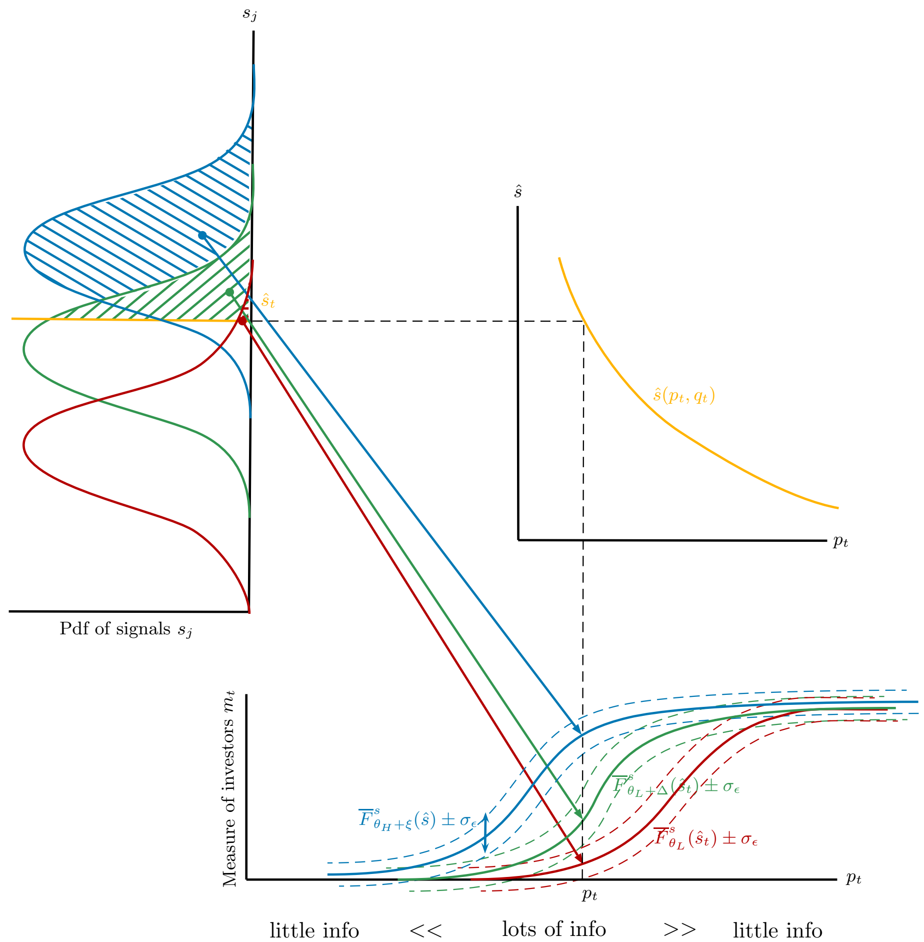

Notes: The top-left panel displays the distribution of private signals in the three states of the world along with the expected measure of investing agents $m_t^e = \overline{F}^s_{\theta+\xi}(\hat{s}_t)$, as previously represented in Figure 1 rotated by 90°. The top-right panel displays an example of equilibrium threshold $\hat{s}(p_t, q_t)$ as a function of public belief p_t. The bottom right panel shows how the measures of investors $m_t = m_t^e + \varepsilon$ varies with public belief p_t, keeping q_t constant in the background, under the assumption that $\varepsilon_t \sim \mathcal{N}(0, \sigma_\varepsilon)$. The mean m_t^e is represented with a continuous line and the corresponding ± 1-standard deviation σ_ε error bands with dashed lines.

Figure 2: Measure of investors m_t as a function of public belief p_t

the information flow is never exactly 0 so that there is always some learning taking place through m_t and R_t. Such a smooth form of information cascades is of interest to us for two reasons: i) it explains why the economy may remain for an extended period of time in the booming region, where agents understand that they could be wrong in their assessment of the true state of the world but invest nonetheless, ii) it opens the door to the economy endogenously exiting the information cascade and crashing when some threshold in beliefs is reached, as we will now describe.

Endogenous boom-and-bust cycle

We now present simulations of the model to illustrate its ability to generate endogenous boom-and-bust patterns out of a single impulse shock. We do not attempt to make a realistic calibration for the moment, but merely pick parameters so as to highlight the model's properties. We will examine later under what general conditions one should expect the highlighted phenomena to occur.

We present the impulse responses of the measure of investors (m_t) and the public beliefs (p_t, q_t), keeping all other shocks to their mean levels (e.g., ε_t, $u_t = 0$), when the economy is in the false-positive state $(\theta, \xi) = (\theta_L, \Delta)$, the case of interest for our purpose. Figures 9 and 10 in the Appendix show the economy's response to the high-technology and low-technology states.[5]

Figures 3 and 4 present two examples of endogenous boom-and-bust patterns that may arise in the model, depending on whether or not the economy falls into an information cascade. In both examples, the emergence of boom-and-bust patterns hinges on two key assumptions: (i) $\theta_L + \Delta$ needs to be sufficiently close to θ_H, so that the two states are hard to distinguish; (ii) the prior q_0 on the false-positive state (θ_L, Δ) needs to be sufficiently low relative to the true positive $(\theta_H, 0)$ for agents to initially attribute most of the rising investment pattern to the true positive state.

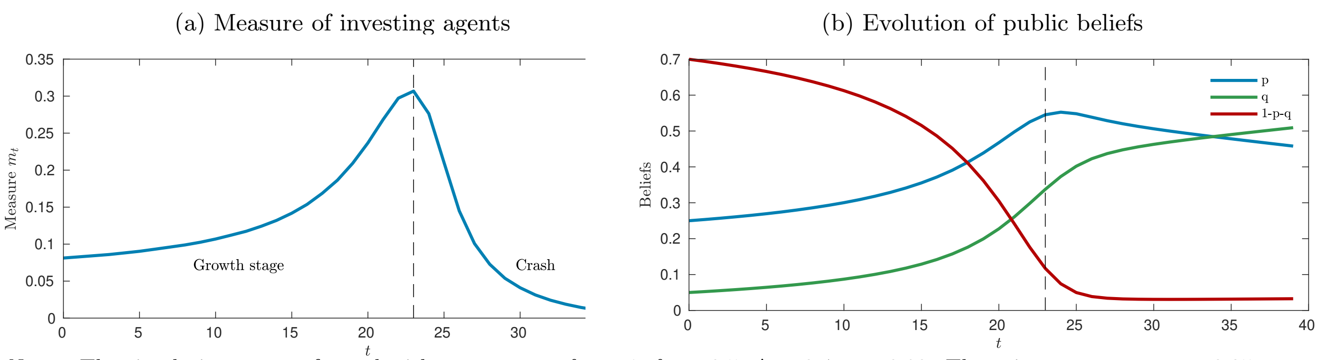

Notes: The simulation was performed with parameters: $\theta_H = 1$, $\theta_L = 0.5$, $\Delta = 0.4$, $c = 0.80$. The priors are set to $p_0 = 0.25$ and $q_0 = 0.05$. All the distributions are Gaussian: $F^s_{\theta+\xi} \sim \mathcal{N}(\theta + \xi, \sigma_s)$, $F^\varepsilon \sim \mathcal{N}(0, \sigma_\varepsilon)$ and $F^u \sim \mathcal{N}(0, \sigma_u)$ with standard deviations $\sigma_s = 0.5$, $\sigma_\varepsilon = 0.2$, $\sigma_u = 2.5$.

Figure 3: Slow boom, sudden crash

Figure 3 presents the evolution of an economy with a high cost of investing c. When the economy starts in period 0, the measure of investing agents (panel 3a) is low (because of the high cost c) but higher than expected. Seeing an unusually high investment rate, agents understand that it is unlikely to come from the low state and they reduce the probability assign to it (red curve in panel 3b). Agents understand that high investment in the model could arise either from the high-technology or the false-positive states. As a result, agents revise upward their probability assessments of both states (p_t and q_t rise). Importantly, however, given that agents start with a low prior on the false-positive state from happening, the observed high level investment is mostly attributed to the high-technology state, so the rise in p_t dominates their expectation. Consequently, agents become more optimistic overall, investment continues to grow, and the rising investment

[5]With our parametrization, these cases are relatively uninteresting: learning is fairly quick, and the dynamics are close to the full information case.

pattern, in turn, leads to further upward revisions in expectations, seemingly confirming the assessment that the economy is in the high state. We refer to this first stage of the cycle, characterized by the joint rising evolution of investment and beliefs (p_t, q_t), as the "growth stage".

Being rational, agents do understand the possibility that they may be mistaken and keep track of the probability of the false-positive state q_t in the background, which also rises throughout the growth stage. Since signals are unbiased along the impulse response path, the belief q_t rises in fact faster than p_t despite starting from a lower prior. Therefore, there comes a time when q_t is so high that agents become reluctant to invest and aggregate investment begins to decline. This is the beginning of the "crash" stage, which arises at an endogenous date without the need of an exogenous trigger. As investment reaches a peak of about 30% given our parametrization, the measure of investing agents m_t attains the intermediate region depicted in Figure 2 where it becomes more informative. As a consequence, agents learn the truth faster, investment drops, and the probability p_t starts declining until a belief reversal occurs later when the belief q_t takes over. Note that the truth is always learned in the end because of the strictly positive information flow and the law of large numbers.

This example shows that the model is able to generate asymmetric cycles. The growth stage is slow due to the low information flow when m_t is close to 0. The crash, on the other hand, is more sudden because it occurs in the region where i) uncertainty between the high state and the false-positive state is high (p_t and q_t are close, so beliefs are more responsive to new information), ii) the signal m_t is more informative at the peak.[6]

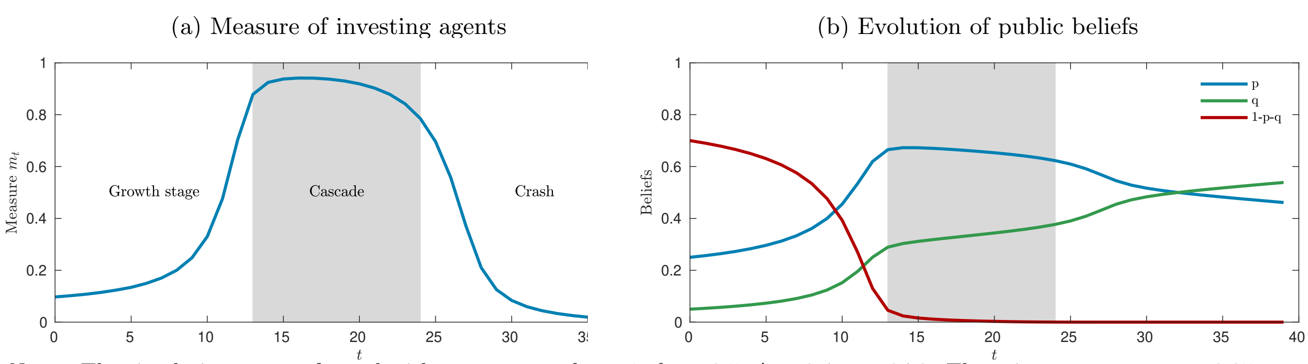

Notes: The simulation was performed with parameters: $\theta_H = 1$, $\theta_L = 0.5$, $\Delta = 0.4$, $c = 0.79$. The priors are set to $p_0 = 0.25$ and $q_0 = 0.05$. The distributions of all signals are Gaussian with standard deviations $\sigma_s = 0.5$, $\sigma_\varepsilon = 0.2$, $\sigma_u = 2.5$. The shaded area corresponds to the informational cascade period, loosely defined as periods when $m_t \geq 75\%$ in which learning is markedly slower, given our calibration.

Figure 4: Endogenous boom and bust with information cascade

Whether the growth stage gives way to a sudden collapse depends on the parametrization of the model. Figure 4 depicts an example of a cycle that reaches the low-informativeness region, in which

[6] A similar mechanism is at work in Veldkamp (2005) where crashes, as they happen when information flows more rapidly, occur suddenly but in response to exogenous shocks.

case the economy goes through an information cascade before the crash. This simulation uses the same parametrization as Figure 3, but with a slightly lower cost c so that investment rises faster and reaches higher levels than in the previous example.As a result, there comes a time at the end of the growth stage when agents are so optimistic that m_t reaches the extreme right region in the bottom panel of Figure 2 where it becomes uninformative. The economy thus enters a period akin to an information cascade, as described earlier, where almost all agents invest except for a few who prefer to stay away due to their particularly pessimistic private information. Through this mechanism, the economy may remain stuck for a long period of time with wrong beliefs and excessive investment. Because the flow of information is never exactly zero, the economy eventually exits the cascade. This event occurs when the belief about the false-positive state q_t reaches a threshold at which a sufficient fraction of agents stop investing, bringing back the economy to the region where m_t is informative. The crash takes place in a manner similar to the previous example: because of the high flow of information, beliefs converge more quickly to their true values and a belief reversal occurs in the later stages. The way the economy exits the cascade is reminiscent of the mechanism proposed by Caplin and Leahy (1994) and its reinterpretation in Chapter 4 of Chamley (2004).

2.9 Continuous Case

How general are the phenomena highlighted in the 3-state model? In this section, we discuss under what conditions endogenous boom-and-bust cycles may arise in a less restrictive environment.

First, we relax the three-state assumption and return to the specification where ξ can take any value, including a continuum. Second, we wish to understand how our two key conditions, (i) $\theta_L + \xi$ close to θ_H and (ii) low q_0, translate to the more general case. To build intuition on this issue, Figure 5 shows the impulse responses of the economy in the continuous-ξ case assuming that ξ is independent of θ and is normally distributed with mean 0 and standard deviation σ_ξ. As in the previous section, we present the response of the economy in the low-technology state $\theta = \theta_L$ but we vary the size of the ξ shock. Four cases are represented as multiples of the standard deviation so that $\xi = n\sigma_\xi$ with $n \in \{1.5, 1.9, 2, 2.1\}$. The figure shows very distinct behaviors depending on the size of the shock. When the shock is relatively small, $\xi = 1.5\sigma_\xi$ (yellow curve), the economy does not experience any herding behavior in which the high initial investments leads to rising optimism. Agents put a sufficiently high likelihood on this event and are, consequently, able to detect it relatively quickly. Things start to differ as we increase the size of the shock. For an intermediate-sized shock, $\xi = 1.9\xi$ (red curve), the economy begins to experience a boom-bust cycle of the sort we described earlier. Because of the low probability of experiencing a shock close to two standard deviations, agents are initially fooled by the high investment rates and the economy enters a growth stage with rising optimism and investment. The growth stage is slow and the crash occurs around date $t = 18$ without experiencing a cascade, as in Figure 3. When the size of the shock is larger, $\xi \geq 2\sigma_\xi$ (green and blue curves), the rise in investment is so large that the economy goes through an information cascade after experiencing a short growth stage, as in Figure 4. The economy exits the cascade endogenously at a date which is further delayed as the size of the shock increases.

(a) Measure of investing agents m_t (b) Evolution of public beliefs p_t

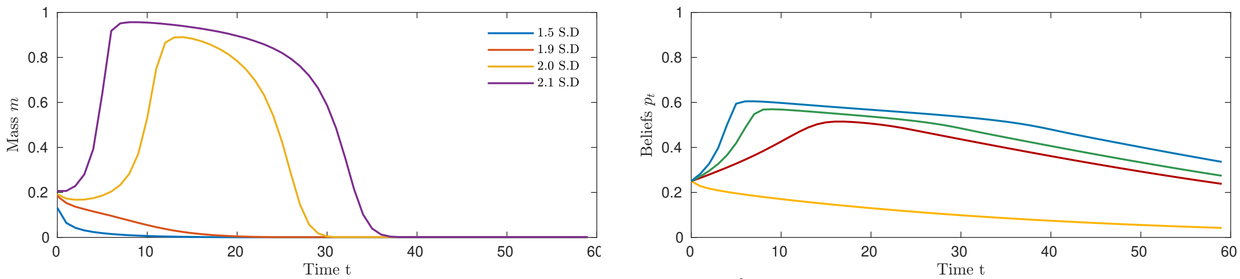

Notes: Panel (b) shows the overall probability of the high state $p_t = \int \pi_t(\theta_H, \xi)\, d\xi$. The simulation was performed with parameters: $\theta_H = 1$, $\theta_L = 0.5$, $\Delta = 0.4$, $c = 0.75$. The priors are set to $p_0 = 0.25$ and $q_0 = 0.05$. All the distributions are Gaussian as in Figure 3 with the additional assumption that $\xi \sim \mathcal{N}(0, \sigma_\xi)$ with standard deviations $\sigma_s = 0.5$, $\sigma_\varepsilon = 0.2$, $\sigma_u = 2.5$ and $\sigma_\xi = 0.25$.

Figure 5: Boom-and-bust cycles in the continuous case

These simulations show that the dynamics depicted in the examples of Figures 3 and 4, in the previous section, are not mere curiosities but regular fixtures of the more general model. Indeed, the simulations show that the endogenous boom-and-bust phenomenon occurs whenever the shock to ξ is unusually large, sufficiently so that agents underestimate its likelihood and initially attribute the observation of high investment to the high-technology state. This shows a limit of the theory: because their are rational, agents cannot make systematic mistakes in their assessment of state probabilities, and this model only offers a theory of *infrequent* boom-and-busts.

The following proposition establishes this result formally in the Gaussian case. We show that there always exists a sufficiently large shock in ξ to trigger a boom-and-bust cycle in beliefs, as long as the exogenous signal coming from the observation of R_t is not too precise.

Proposition 2. *In the Gaussian case, i.e.,* $F^\xi \sim \mathcal{N}\left(0, \sigma_\xi^2\right)$, $F^s|\theta, \xi \sim \mathcal{N}\left(\theta + \xi, \sigma_s^2\right)$, $F^\varepsilon \sim \mathcal{N}\left(0, \sigma_\varepsilon^2\right)$, $F^u \sim \mathcal{N}\left(0, \sigma_u^2\right)$, *for* θ *and* ξ *independent and signal* R_t *sufficiently uninformative* (σ_u *low), there exists a* $\underline{\xi}$ *such that all shocks* $\xi \geq \underline{\xi}$ *generate a boom-and-bust cycle in the impulse response of beliefs* p_t *to a false-positive shock* (θ_L, ξ)*.*

2.10 Welfare

We now turn to the analysis of welfare in this economy. Since investors do not internalize that their investment decisions affect the release of public information, the equilibrium is in general not efficient and policy interventions can be beneficial. To show this formally, we introduce a social planner that maximizes aggregate welfare under limited information. Specifically, we assume that the planner only observes the public signals and cannot rely on the private information of the investors when making decisions. We impose these restrictions so that problem of the planner is not trivial and that it resembles that of a government trying to design policy under uncertainty about the true value of the new technology.

As in Angeletos and Pavan (2007), we assume that the planner seeks to maximize the sum of the investors' expected utility, where the expectation is computed according to the investor's private beliefs. Written in recursive form, the problem of the social planner is

$$V(\mathcal{I}) = \max_{\hat{p}} E_{\theta,\xi}\left[\int_{p_j \geq \hat{p}} E\left[\theta - c \mid \mathcal{I}_j\right] dF^{p_j}_{\theta+\xi}(p_j) \mid \mathcal{I}\right] + \beta E_{\theta,\xi}\left[V\left(\mathcal{I}'\right) \mid \mathcal{I}\right] \tag{11}$$

where $\mathcal{I}'$ is public information next period, which evolves according to the law of motion (10), and where $F^{p_j}_{\theta+\xi}(p_j)$ is the CDF of the agents' subjective probability that $\theta = \theta_H$ when the true state of the world is $\theta + \xi$. The first term in (11) captures the current-period returns on investment of letting only agents with private beliefs above $\hat{p}$ invest. To compute that term, the planner first uses the public beliefs $\mathcal{I}$ to evaluate the likelihood of being in a given state $\theta + \xi$. Since the planner knows the structure of the economy, it can then reconstruct the distribution $F^{p_j}_{\theta+\xi}(p_j)$ of private beliefs in that state, which is needed to compute the mass of investors above $\hat{p}$. The second term in (11), the continuation value, captures the impact of a given investment threshold $\hat{p}$ on the future public information. It is this term that creates a gap between the equilibrium and the efficient allocation. In the equilibrium, the actions of an individual investor have a negligible impact on the release of information, so that each investor disregards that channel when making their decisions. The planner, on the other hand, understands that by changing the cutoff $\hat{p}$, the mass of investors also changes and so does the release of public information.

Taking the derivative with respect to $\hat{p}$ in (11) the first-order condition is

$$E_{\theta,\xi}\left[\left(\hat{p}\theta_H + (1-\hat{p})\,\theta_L - c\right) f^{p}_{\theta+\xi}(\hat{p}) \mid \mathcal{I}\right] = \beta \frac{\partial E_{\theta,\xi}\left[V(\mathcal{I}') \mid \mathcal{I}\right]}{\partial \hat{p}}. \tag{12}$$

The left-hand side of this equation reflects the expected cost of increasing the threshold $\hat{p}$ at the margin. If the true state is $\theta + \xi$, increasing $\hat{p}$ at the margin pushes a mass $f^{p}_{\theta+\xi}(\hat{p})$ of agents away from investing, each of which loses $\hat{p}\theta_H + (1-\hat{p})\,\theta_L - c$ in expected returns. The planner takes the expectation of these terms over all $\theta + \xi$ since it does not know the true state. The right-hand side of the equation reflects the impact of increasing $\hat{p}$ on the flow of public information that will be released at the end of the period. By changing $\hat{p}$, the planner can increase the gap between the expected realizations of m in different states of the world. When it does so, m becomes more informative since the exogenous noise ε is less able to drown the signal.

We can solve the planner's problem numerically and go back to the boom-bust cycles explored in Figure 4 to see how a social planner would deviate from the equilibrium. Figure 6 compares the efficient allocation (bold lines) with the equilibrium (thin lines). We see from Panel (a) that in period 0 the planner is more aggressive than the private agents in pursuing the investment opportunity: the initial mass of investors is 0.25 in the efficient allocation while it is only 0.15 in the equilibrium. The planner behaves in that way because it understands that the additional investment releases valuable information. Indeed, we can see in Panel (b) that the public beliefs move more rapidly in the efficient allocation. In particular, the planner learns quickly that the low-technology state can be completely ruled out, as $1 - p - q$ declines sharply in the first few periods. Since many agents invests under the planner's chosen threshold $\hat{s}$, the likelihood of being in the high-technology state increases, and the (mistaken) boom begins even though the the true fundamental is θ_L. This example shows that while the planner behaves in a way that releases more information, it is still fooled into believing that the new technology is great. As in the equilibrium, the efficient allocation features herd-driven boom-bust cycles, but the busts happen faster here. When the beliefs the public beliefs are very optimistic the planner pushes for less investment than in the equilibrium so that information keeps being released. As a result, agents learn faster that the true state of the world is θ_L and the boom comes to an end sooner.

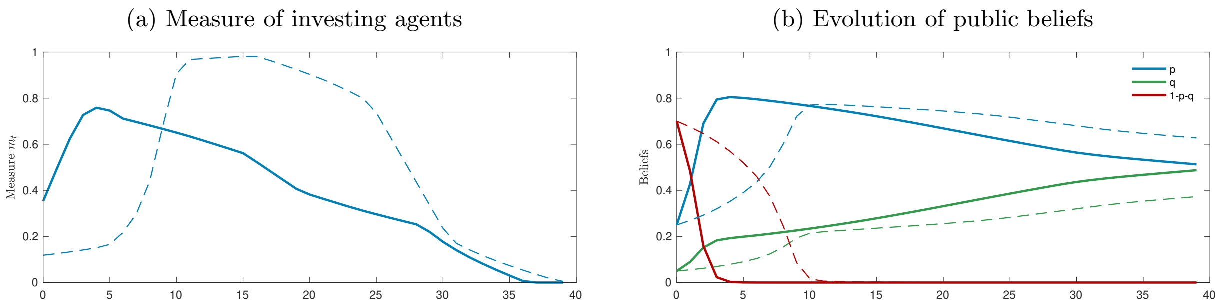

Notes: Bold lines correspond to the efficient allocation and thin lines correspond to the equilibrium. The true value of fundamental is $(\theta, \xi) = (\theta_L, \Delta)$, the false-positive state. The simulation was performed with parameters: $\theta_H = 1$, $\theta_L = 0.5$, $\Delta = 0.4$, $c = 0.79$, $\beta = 0.5$. The priors are set to $p_0 = 0.25$ and $q_0 = 0.05$. The distributions of all signals are Gaussian with standard deviations $\sigma_s = 0.4$, $\sigma_\varepsilon = 0.2$, $\sigma_u = 2.5$.

Figure 6: Endogenous boom-and-bust in the efficient allocation

To better understand how the planner behaves in this economy, we can construct an optimal investment tax that makes the equilibrium and the efficient allocation coincide, as the next proposition shows.

Proposition 3. *The efficient allocation can be implemented as an equilibrium by an investment tax*

$$\tau^* = \left(E_{\theta,\xi} \left[f^p_{\theta+\xi} (\hat{p}) \mid \mathcal{I} \right] \right)^{-1} \beta \frac{\partial E_{\theta,\xi} \left[V \left(\mathcal{I}' \right) \mid \mathcal{I} \right]}{\partial \hat{p}}, \tag{13}$$

and a lump-sum transfer to all investors.

The optimal tax τ^* balances the distortion in investment it creates (first term in the product) with the potential benefit on information acquisition (second term). Figure 7 shows this optimal tax over time in our 3-state example. At time $t = 0$, investors are quite pessimistic about the new technology and would rather not invest. But some amount of investment would release valuable information so the optimal tax τ^* is initially negative. Soon after, between periods $t = 3$ and $t = 5$, the planner faces the opposite problem. Investors believe that the new technology is good and invest too much. The optimal tax is therefore positive to increase the flow of information. At $t = 5$, the planner is sufficiently convinced of being in state θ_H that acquiring more information loses its value and the planner slowly remove the tax. Things change at about $t = 8$. The newly released information suggests that we might not be in θ_H after all and the likelihood of the false-positive state rises. Indeed we can see that p decreases and q increases in Panel (b) of Figure 6. As a result, the value of information increases as well and the optimal tax increases to incentive agents to invest less. The tax starts to decline again around $t = 17$ once the planner is getting convinced that we are actually in the false-positive state and the value of information shrinks. Finally, the tax turns negative once again as agents massively refuse to invest while some investment would provide valuable information. After $t = 35$, the public beliefs are sufficiently certain that the true state is (θ_L, ξ), there is no value in obtaining additional information and the optimal tax goes to zero.

As we can see, the tax incentivizes agents to behaves differently than the crowd in what amounts to a *leaning-against-the-wind* pattern: The tax is negative when no agent wants to invest, and positive when agents invest massively. We will see in our quantitative model that these same forces have important consequences for the conduct of monetary policy.

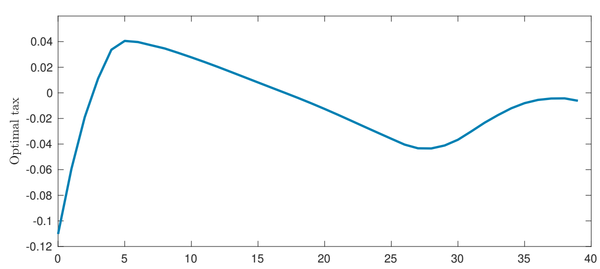

Notes: Optimal tax to implement the efficient allocation after the same shock as (6). The simulation was performed with parameters: $\theta_H = 1$, $\theta_L = 0.5$, $\Delta = 0.4$, $c = 0.79$, $\beta = 0.5$. The priors are set to $p_0 = 0.25$ and $q_0 = 0.05$. The distributions of all signals are Gaussian with standard deviations $\sigma_s = 0.4$, $\sigma_\varepsilon = 0.2$, $\sigma_u = 2.5$.

Figure 7: Optimal tax over time

3 Business Cycle Model with Herding

After exploring the mechanism in the simple model, we now embed the same economic forces in a business cycles framework. Our objective is threefold. First, our previous setup is highly stylized and we want to examine the robustness of the mechanism in a more realistic environment that involves more moving parts (e.g., prices and constraints). Second, we want to investigate under what conditions the evolution of beliefs characterized in the previous section may lead to a macroeconomic expansion followed by a contraction deep enough to go below the trend. This requires additional ingredients as we will now discuss. Finally, a more realistic setup is required if we want to explore the quantitative implications of the theory.

3.1 Foreword

Generating business cycle fluctuations out of belief shocks has been the focus of the *news* (or *noise*)-driven business cycle literature since Beaudry and Portier (2004) and recently reviewed in Beaudry and Portier (2014). A key lesson from this literature is that standard models have difficulty generating positive comovements across macroeconomic aggregates out of sheer optimism, particularly between consumption and investment.

The failure to generate positive comovements stems from two main reasons. First, there is a *static* problem, originally identified by Barro and King (1984), due to the intratemporal labor market equilibrium: when agents become more optimistic about technology, the expected higher income encourages agents to cut on their labor supply which leads to a contraction in output. Second, there is a *dynamic* problem arising from standard parametrizations with intertemporal elasticity of substitution of consumption less than 1 (e.g., CRRA utility function with relative risk aversion greater than 1): anticipating higher wealth, agents smooth consumption by moving resources from the future to the present and disinvest in response to a positive belief shock.

To circumvent the first difficulty, we follow Lorenzoni (2009) and introduce price stickiness. With price rigidities, a positive belief shock can be expansionary as long as monetary policy is sufficiently accommodative. In that case, firms fulfill the higher demand and a fall in the real interest helps sustain the expansion in demand due to optimistic expectations. We solve the second difficulty by proposing a model of technology adoption with two types of capital: a new-technology-specific capital (e.g., IT capital) and a traditional form of capital. Assuming that the new technology is intensive in IT capital, a rise in IT investment is a requisite for agents to benefit from the innovation. The addition of capital adjustment costs, as in Jaimovich and Rebelo (2009), to prevent a steep decline in traditional capital suffices to guarantee a joint increase in aggregate consumption and investment.

3.2 Household

There are four types of agents: i) a representative household, ii) entrepreneurs, who face a technology adoption choice, iii) retailers, who are the only agents facing price rigidities, and iv) a monetary authority. The household lives forever, consumes, supplies labor and is the owner of all the firms and capital stocks in the economy. The preferences of the household are given by

$$E\left[\sum \beta^t \left(\frac{C_t^{1-\gamma}}{1-\gamma} - \frac{L_t^{1+\psi}}{1+\psi}\right)\right], \quad \gamma \geq 1, \psi \geq 0,$$

where C_t is the consumption of the final good and L_t is labor. The household can save in a risk-free one-period nominal bond, B_t, and in two different forms of capital: a traditional type (T) in quantity K_t^T and a new-technology-specific capital (IT) in quantity K_t^{IT}. The household is subject to the real budget constraint

$$C_t + \sum_{i=T,IT} I_t^i + \frac{B_t}{P_t} = w_t L_t + \sum_{i=T,IT} z_t^i K_t^i + \frac{1+r_{t-1}}{1+\pi_t}\frac{B_{t-1}}{P_{t-1}} + \Pi_t,$$

where I_t^i, $i = T, IT$, is the investment in each capital type, z_t^i the corresponding real rental rate, w_t the real wage, Π_t the total profits, r_{t-1} is the nominal interest rate on government debt issued at date $t-1$, P_t is the nominal price level and $1 + \pi_t = P_t/P_{t-1}$ the inflation rate.

For the reasons invoked earlier, agents face adjustment costs in capital in the same form as in Christiano et al. (2005). The law of motion for each type of capital, $i = T$ or IT, is given by

$$K_{t+1}^i = (1-\delta) K_t^i + I_t^i \left(1 - S\left(\frac{I_t^i}{I_{t-1}^i}\right)\right), \; S(x) = \frac{\kappa}{2}(x-1)^2,$$

where $\kappa \geq 0$ is a parameter that determines the size of the adjustment costs.

3.3 Technology

There are four sectors: i) an entrepreneur sector, ii) a wholesale sector, iii) a retail sector and iv) a final good sector. The most important one, the entrepreneur sector, is the analog of the investment model from section 2.

Entrepreneur sector

There is a unit continuum of entrepreneurs indexed by $j \in [0, 1]$ who are monopolistic producers of differentiated varieties sold to the wholesale sector. Until date 0, entrepreneurs have access to a unique "old" production technology, which is Cobb-Douglas in some capital bundle K_{jt}^o, to be described shortly, and labor L^o_{jt},

$$Y^o_{jt} = A^o \left(K^o_{jt}\right)^{\alpha} \left(L^o_{jt}\right)^{1-\alpha}, \, 0 \leq \alpha \leq 1.$$

In order to abstract from standard real business cycle-like fluctuations, we assume that the "old" TFP, A^o, is constant over time. Unexpectedly, at date 0, a "new" technology arrives with production function

$$Y^n_{jt} = A^n_t \left(K^n_{jt}\right)^{\alpha} \left(L^n_{jt}\right)^{1-\alpha}, \, 0 \leq \alpha \leq 1.$$

The TFP of the new technology A^n_t is characterized by a constant fundamental $\theta \in \{\theta_H, \theta_L\}$, $\theta_H > \theta_L$, whose value is initially unknown. Importantly, the new technology is not immediately productive. We make the assumption that the new technology is initially as productive as the old one, $A^n_t = A^o$, until it matures with some fixed probability $\lambda > 0$. Upon maturation, the true nature of the technology is revealed. Maturation is a one-time event and all uncertainty is resolved afterwards. That is,

$$A^n_t = \begin{cases} A^o & \text{before maturation} \\ \theta & \text{after maturation.} \end{cases}$$

In addition to differing in TFP, the two technologies differ in the capital bundle they use as input. The capital bundle used by each technology $i = o, n$ is given by

$$K^i_{jt} = \left(\omega_i \left(K^{IT}_{it}\right)^{\frac{\zeta-1}{\zeta}} + (1-\omega_i)\left(K^T_{it}\right)^{\frac{\zeta-1}{\zeta}}\right)^{\frac{\zeta}{\zeta-1}}, \, \zeta > 0, \tag{14}$$

with the assumption that the intensity in IT capital is greater for the new than the old technology, $1 \geq \omega_n > \omega_o \geq 0$.

After date $t \geq 0$, entrepreneurs face a technology choice problem. We assume that a fraction $0 \leq \mu \leq 1$ of entrepreneurs are "noise entrepreneurs", that is, they are clueless regarding technological adoption and behave randomly. Specifically, we assume that a fraction ε_t of them adopt the new technology, where ε_t is i.i.d, distributed according to CDF F^{ε} over support $[0,1]$. The remaining $1-\mu$ of entrepreneurs are rational and choose the best of the two technologies based on public and private information. There is no cost of switching, so the decision for entrepreneur j is static:

$$i_{jt} = \underset{i_{jt} \in \{0,1\}}{\operatorname{argmax}} \; (1 - i_{jt}) \, E\left[\Pi^o_t | \mathcal{I}_{jt}\right] + i_{jt} E\left[\Pi^n_t | \mathcal{I}_{jt}\right],$$

where Π^i_t, $i = o, n$, are the profits from using technology i, $\mathcal{I}_{jt}$ is the information set of entrepreneur j, and i_{jt} is a dummy capturing the technology adoption decision.

Wholesale sector

The wholesale, retail sectors and final good sectors play no major role in the model other than separating price rigidities from the technology choice problem of entrepreneurs.

The wholesale sector is modeled as a representative firm which produces a wholesale good with CES technology

$$Y_t^w = \left(\int_0^1 Y_{jt}^{\frac{\sigma-1}{\sigma}} dj \right)^{\frac{\sigma}{\sigma-1}}, \sigma \geq 0, \tag{15}$$

where Y_{jt} is the quantity of inputs it purchases from the monopolistic entrepreneurs. The wholesale sector is perfectly competitive, giving rise to the demand schedule, $Y_{jt} = (P_{jt}/P_t^w)^{-\sigma} Y_t^w$, where $P_t^w = \left(\int_0^1 P_{jt}^{1-\sigma} dj \right)^{\frac{1}{1-\sigma}}$ is the price of the wholesale good and P_{jt} the price of each differentiated entrepreneur good.

Retail sector

The retail sector is composed of a unit continuum of monopolistic producers who buy the wholesale good at P_t^w and costlessly differentiate it using a one-to-one technology. Retail sector firms are the only ones to face price rigidities. We assume that they face Calvo-style frictions: firms can only reset their price with probability $1 - \chi$, leading to a standard Phillips curve.

Final good sector

The final good sector, similar to the wholesale sector, is modeled as a representative firm that operates under perfect competition and produces the final good, used for consumption and investment, using inputs from the retail sector. It uses the CES technology,

$$Y_t = \left(\int_0^1 \left(Y_{jt}^r \right)^{\frac{\sigma-1}{\sigma}} dj \right)^{\frac{\sigma}{\sigma-1}},$$

where Y_{jt}^r is the quantity purchased from each retail firm and σ is the same elasticity of substitution as in (15).

3.4 Monetary Authority

To close the model, we need to specify the policy followed by the monetary authority. As is common in the literature, we assume that the central bank follows a Taylor rule,

$$\frac{1 + r_t}{1 + \overline{r}} = \left(\frac{1 + \pi_t}{1 + \overline{\pi}} \right)^{\phi_\pi} \left(\frac{Y_t}{\overline{Y}} \right)^{\phi_y}, \tag{16}$$

where $\overline{r}$, $\overline{\pi}$ and $\overline{Y}$ correspond to the values of, respectively, the nominal interest rate, inflation and output at the steady state around which we linearize the economy, which we define in the next section.

3.5 Information

The information structure for entrepreneurs mimics the one in the simple model of Section 2. The true technology parameter θ is drawn once-and-for-all at date 0. The ex-ante probability that $\theta = \theta_H$ is denoted by p_0. Agents in the economy cannot observe θ directly but receive various private and public signals about it. After maturation, we assume that θ is revealed to everyone, so the economy operates under full information afterwards. Other than the true realization of θ, there is no uncertainty in the economy. In particular, the productivity of the old technology, A^o, is known and there is common knowledge about the distributions of the various shocks.

Private information

We assume that entrepreneurs receive a private signal s_j about θ at date 0 when the new technology appears. Replicating the same information structure as Section 2, entrepreneurs draw their signals from the CDF $F^s_{\theta+\xi}$ where $\xi \sim F^\xi$ captures common noise and the family $\{F^s_x\}_{x\in I}$ satisfies the same conditions as before, including the MLRP property.

Public information

In addition to observing their private signals, entrepreneurs and all other agents in the economy (household, central bank, retailers, etc.) collect public information over time. As in the simplified model, social learning takes place through the observation of market activity. In particular, we assume that agents observe the measure of entrepreneurs that adopt the new technology:

$$m_t = \int_0^1 i_{jt} dj = \underbrace{\int_0^{1-\mu} i_{jt} dj}_{\text{rational entrepreneurs}} + \underbrace{(1-\mu)\,\varepsilon_t}_{\text{noise entrepreneurs}}, \text{ with } \varepsilon_t \sim \text{iid CDF } F^\varepsilon. \tag{17}$$

The measure of the new-technology adopters m_t in (17) is almost identical to the measure of investors in (7) with the exception that we now take a stand on the origin of the noise by assuming that it arises from a fraction μ of noise entrepreneurs.[7] Apart from that point, the informational content of m_t is identical to the earlier model up to some rescaling.

[7] In this general equation model, m_t is a real quantity that enters many equations (resource constraints, market clearing equations, etc.). For simplicity and to make sure that agents learn the same from observing prices and aggregate quantities as from the direct observation of m_t, we choose to interpret ε_t as "real" noise coming from the actions of noise entrepreneurs that contaminates the true quantity of technology adopters rather than pure informational noise.

Because the productivity of the new technology is identical to A^o until maturation, there is no other source of information in the economy. In equilibrium, prices and aggregate quantities will solely be functions of m_t and of public information up to time t. As a result, prices and quantities provide no other information than already contained in m_t.

Beliefs

As in the simple model, we denote by $\mathcal{I}_t = \{m_{t-1}, \ldots, m_0\}$ the public information available to non-entrepreneur agents (households, monetary authority, retailers and outside observers). Public beliefs are captured by the joint distribution

$$\pi_t\left(\tilde{\theta}, \tilde{\xi}\right) = Pr\left(\theta = \tilde{\theta}, \xi \in \left[\tilde{\xi}, \tilde{\xi} + d\tilde{\xi}\right] \mid \mathcal{I}_t\right).$$

Finally, we denote by $\mathcal{I}_{jt} = \mathcal{I}_t \cup \{s_j\}$ the information set of entrepreneur j and her beliefs by the joint distribution $\pi_{jt}\left(\tilde{\theta}, \tilde{\xi}\right) = Pr\left(\theta = \tilde{\theta}, \xi \in \left[\tilde{\xi}, \tilde{\xi} + d\tilde{\xi}\right] \mid \mathcal{I}_{jt}\right)$.

3.6 Timing

Before date 0, the economy is in a deterministic, no-inflation steady state using the old technology. At date 0, the new technology fundamental θ, the common noise component ξ and the private signals s_j are drawn once-and-for-all. For all date $t \geq 0$,

1. Entrepreneurs choose whether to adopt the new technology or not based on their individual beliefs π_{jt},
2. The measure of technology adopters m_t is realized,
3. The new technology matures with probability λ,
4. Simultaneously:
   (a) All agents observe m_t and update their information,
   (b) Production takes place,
   (c) The household chooses consumption, investment and labor supply,
   (d) The monetary authority sets the policy rate,
   (e) Markets clear.

3.7 Investment Decision

The technology adoption decision is more complicated than in the simple model because of the presence of general equilibrium effects. When choosing whether to use the new technology, agents have to forecast the profits from either technology. Profits in equilibrium depend not only on productivity but also on the level of demand from wholesalers Y_t^w, prices and the real marginal costs $mc_t^i = \frac{1}{A_{jt}}\left(\frac{z_t^i}{\alpha}\right)^\alpha \left(\frac{w_t}{1-\alpha}\right)^{1-\alpha}$ from using each technology $i = o, n$:

$$\Pi_t^i = \left(P_{jt} - P_t mc_t^i\right) Y_{jt} = \frac{1}{\sigma}\left(\frac{\sigma}{\sigma-1}\right)^{-\sigma} P_t^{1-\sigma} \left(P_t^w\right)^\sigma \left(mc_t^i\right)^{1-\sigma} Y_t^w.$$

Solving the model by linearizing the equations of the DSGE model, entrepreneur j ultimately chooses to invest if and only if

$$E\left[\hat{A}_t^n - \alpha \hat{z}_t^n \,|\, \mathcal{I}_{jt}\right] \geq E\Big[\underbrace{\hat{A}_t^o}_{=0} - \alpha \hat{z}_t^o \,|\, \mathcal{I}_{jt}\Big], \tag{18}$$

where the hatted variables are log-deviations from a steady state that we define in the next section and z_t^i, $i = n, o$ is, as before, the rental rates on the capital bundles (14). As equation (18) demonstrates, entrepreneurs not only have to forecast the technology A_t^n but also factor prices, as they are now competing for the same inputs. This introduces a novel dimension that may break down the monotonicity of the investment decision with respect to private signals and, hence, will require special attention in the resolution method.

3.8 Belief Updating

The information structure in this general model is essentially the same as in the simple model and we obtain the same simplification that allows us to split the static private part of beliefs from their dynamic, time-varying public part. As a consequence, the technology decision i_{jt} is a simple function of the aggregate state variables of the economy $\Omega_t = \left(K_t^{IT}, K_t^T, I_t^{IT}, I_t^T\right)$, public beliefs π_t and the private signals s_j. The measure of new technology adopters is given by

$$m_t = (1-\mu)\, m^e\left(\Omega_t, \pi_t, \theta, \xi\right) + \mu \varepsilon_t \tag{19}$$

$$\text{with } m^e\left(\Omega_t, \pi_t, \theta, \xi\right) = \int \mathbb{I}\left(i_j\left(\Omega_t, \pi_t; s_j\right) = 1\right) f_{\theta+\xi}^s\left(s_j\right) ds_j. \tag{20}$$

In turn, the belief updating equation (10) needs to be amended in the following way

$$\pi_{t+1}\left(\tilde{\theta}, \tilde{\xi}\right) = \frac{\pi_t\left(\tilde{\theta}, \tilde{\xi}\right) f^\varepsilon\left(\frac{1}{\mu}\left(m_t - (1-\mu)\, m^e\left(\Omega_t, \pi_t, \tilde{\theta}, \tilde{\xi}\right)\right)\right)}{\int \pi_t\left(\theta, \xi\right) f^\varepsilon\left(\frac{1}{\mu}\left(m_t - (1-\mu)\, m^e\left(\Omega_t, \pi_t, \theta, \xi\right)\right)\right) d\left(\theta, \xi\right)}. \tag{21}$$

4 Quantitative Exercise

We turn to the quantitative evaluation of our general macroeconomic model. After a brief discussion of our resolution method, we calibrate the model to a specific episode in US history and examine the ability of the model to endogenously generate a pattern of macroeconomic expansion followed by a contraction. We finally explore some of the model's implications for the conduct of monetary policy.

4.1 Resolution Method

Our model can only be solved numerically. We follow a strategy, common in the information friction literature and linearize the equations of the model that are unrelated to the updating of beliefs (Woodford, 2003; Angeletos and La'O, 2013). The main benefit of this approach is greater tractability, allowing us to focus on the nonlinearities implied by the learning model while putting aside the (usually weak) nonlinearities of the DSGE model. We carry out the linearization around the non-stochastic zero inflation steady state that preceded period 0—before the new technology is introduced.

Granted the benefits of the linearization approach, one difficulty remains in the need to keep track of the potentially infinite-dimensional public belief $\pi_t(\theta, \xi)$. We use a simplification proposed by Kozlowski et al. (2019) which exploits the fact that, due of the Law of Iterated Expectations, beliefs follow a *martingale*, that is, $E_t\left[\pi_{t+1}\left(\tilde{\theta}, \tilde{\xi}\right) \mid \mathcal{I}_t\right] = \pi_t\left(\tilde{\theta}, \tilde{\xi}\right)$. The martingale property implies that any equilibrium condition of the form $E_t\left[f\left(x_t, x_{t+1}, \pi_t, \pi_{t+1}\right)\right] = 0$, where f is a nonlinear function and x_t a vector of model variables can be approximated to a first-order as

$$E\left[f\left(x_t, x_{t+1}, \pi_t, \pi_{t+1}\right) \mid \mathcal{I}_t\right] \simeq E\left[f\left(x_t, x_{t+1}, \pi_t, \pi_t\right) \mid \mathcal{I}_t\right].$$

This implies that the model can be solved in each period as if current beliefs were constant going forward. As a consequence, we solve the model every period using a standard linear solver, compute the adoption threshold$\hat{s}$ and the evolution of beliefs in a nonlinear way, then repeat in the next period under the new beliefs.

4.2 Calibration (Preliminary)

As we argued before, our model offers a theory of infrequent endogenous booms-and-busts. For that reason, we do not expect our theory to explain general business cycle patterns in the absence of other shocks, but rather to provide a narrative for certain episodes. We thus focus our calibration exercise on a particular episode in recent US history that best fits the description of a technology-driven boom and bust cycle: the late 1990's Dot-Com bubble. We map the new technology in our model to the introduction of IT technologies in the 1990s and we focus more specifically on the late part of the cycle which covers the period that preceded the stock market collapse in the NASDAQ composite index starting from a trough in 1998Q4 to the crash in 2001Q1.

The model is solved at the quarterly frequency. Table 1 lists a first set of standard parameters that we take from the literature. Labor intensity α is set to target a standard labor share of 36%. The discount factor β is set to match an annual real interest rate of about 4%. The

Parameter	Value	Target
α	0.36	Labor share
β	0.99	4% annual interest rate
γ	1	risk aversion (log)
ψ	2	Frisch elasticity of labor supply (Chetty et al., 2011)
θ_p	0.75	1 year price duration
σ	10	Markups of about 11%
ϕ_y	0.125	Clarida et al. (2000)
ϕ_π	1.5	Clarida et al. (2000)
κ	9.11	Schmitt-Grohé and Uribe (2012)
ζ	1.71	Elasticity between types of capital (Boddy and Gort, 1971)

Table 1: Standard parameters

household's preference over consumption is logarithmic and the Frisch elasticity is set to 2, within the range of standard macro-level estimates (Chetty et al., 2011). The Calvo price-setting parameter θ_p is set to yield an average price duration of 1 year. The elasticity of substitution between varieties σ is set to 10 to match an average markup of 11%. The Taylor rule parameters are within the estimates of Clarida et al. (2000). The capital adjustment cost parameter is estimated in Schmitt-Grohé and Uribe (2012). Finally, we pick the elasticity ζ between the different types of capital within the firm from early estimates by Boddy and Gort (1971).

Table 2 below lists the more important parameters that attempt to match features of the Dot-Com Bubble. We set the IT-capital shares ω_i, $i = o, n$, to match an IT investment share in total investment of 3.36% before the introduction of the new technology in 1995 and 3.56% in 2005 (OECD Factbook, 2008). The probability of maturation for new technologies λ is set to 1/10 to match an average waiting time of 10 quarters, corresponding to the period 1998Q4-2001Q1. We now turn to the technology parameters. A^o is normalized to 1. We use the Survey of Professional Forecaster (SPF) mean real GDP growth forecast over the current quarter. Under the assumption that factors are fixed in the short run, this identifies changes in the productivity parameter θ. The highest forecast for growth was 4.19% in 2000Q2 in annualized terms. Correcting for a mean growth trend in GDP of 2.4% over 1991-1998, this yields $\theta_H = 1.05$. Similarly, targeting the lowest growth forecast of 0.80% in 2001Q1, we obtain an estimate $\theta_L = 0.96$. The distribution of private signals is assumed to be Gaussian, centered on $\theta + \xi$ with standard deviation σ_s. To set the dispersion σ_s, we target the average dispersion of growth forecasts in the SPF over 1998-2001. Finally, we must assume a distribution for the fraction of noise traders that adopt the new technology with support over $[0, 1]$. We choose a beta distribution of parameters (2,2).[8]

[8]The distribution Beta$(1, 1)$ is uniform and produces a flat learning response. As a result, we pick a Beta$(2, 2)$

Parameter	Value	Target
ω_o	0.34	IT investment in GDP pre-1995 (2.86%)
ω_n	0.36	IT investment post-2005 (3.56%)
λ	1/10	Duration of NASDAQ boom-bust 1998Q4-2001Q1
θ_h	1.05	SPF highest growth forecast over 1998-2001
θ_l	0.96	SPF lowest growth forecast over 1998-2001
s_j	$\mathcal{N}\left(\theta+\xi, 0.137\right)$	SPF avg. dispersion in forecasts over 1998-2001
ε	Beta$(2,2)$	Non-uniform distribution over $[0,1]$
μ	5%	Fraction of noise traders
ξ	$\mathcal{N}\left(0,\sigma_\xi^2\right)$	See text

Table 2: Dot-Com episode related parameters

Two parameters remain to calibrate for which there does not exist widely accepted estimates or natural targets. The first one is the fraction of noise entrepreneurs μ which controls the informativeness of the social learning channel. While some estimates exist in the literature regarding the informativeness of markets (see for instance David et al., 2016), these estimates do not cover social learning about new technologies. We conduct sensitivity analysis on this parameter but start with a benchmark value of $\mu = 5\%$. Finally, we assume that the common noise shock ξ is normally distributed with standard deviation σ_ξ. Without long time series on the frequencies and magnitudes of boom-bust episodes, we cannot reliably identify the parameter σ_ξ. Instead, we explore different values of σ_ξ and compute the corresponding frequency of boom-bust cycles. Our numerical experiments suggest a nonmonotonic relationship: for low values of σ_ξ, the probability that large enough shocks $\xi > \underline{\xi}$ produce boom-bust cycles is low; for high values of σ_ξ, agents put sufficient probability on high realization of ξ that boom-bust cycles are rarely triggered. The figures we report are computed with some intermediate value of σ_ξ that maximizes the frequency of boom-bust cycles.

4.3 Boom-and-Bust Cycles

Figure 8 presents the impulse responses of the economy to a false-positive shock with $\theta = \theta_L$ but ξ large enough to trigger a boom-and-bust cycle. The measure of high-technology adopters m_t and beliefs evolve in a way that mimics the type of boom-bust cycle highlighted in Figure 4. In this example, the shock ξ is large enough to send the economy very quickly in the region with a large m. Because a false-positive shock of that size is initially deemed unlikely, beliefs about the high state jump on impact pushing the economy in the slow-learning region. As time unfolds, agents observe a high rate of technology adoption that is more or less consistent with the high state, but more

distribution with full support over $[0,1]$ and mode at 0.5.

so with the false positive one. As a result, the high-state posterior probability p_t slowly declines and the false-positive probability rises over time until a crash occurs around period 6, for reasons identical to those we described in the simple learning model.

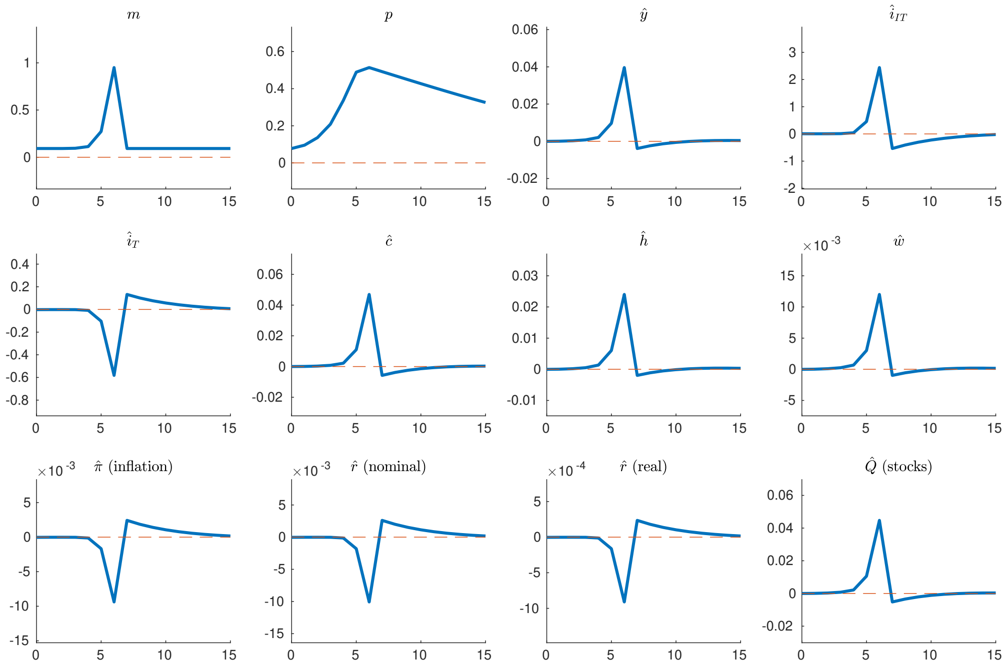

Notes: The impulse responses are reported in log-deviations from the initial non-stochastic steady state. Variable $\hat{h}_t$ captures hours and $\hat{v}_n$ denotes the value of a firm operating the high-technology.

Figure 8: Impulse response to a common noise shock $\xi = 0.96\,(\theta_H - \theta_L)$

More specific to the current exercise is how this pattern of technology adoption and the evolution of beliefs translate to other macroeconomic variables. As agents become more optimistic after observing people rushing to adopt the new technology, the household anticipates higher productivity growth in the future and higher income, resulting in upward pressure on consumption due to a positive income effect. With expectations of higher productivity from the new technology, the demand for IT capital rises and the household responds by increasing IT investment. The new technology being less intensive in the other form of capital, the demand for traditional capital falls and so does traditional investment. The rise in consumption and investment in IT capital, despite being accompanied by a moderate decline in traditional investment, contribute to an overall rise in aggregate demand. This is where price rigidities play an important role. In a real business cycle model, the rise in aggregate demand should be offset by a sharp rise in the real interest rate. With sticky prices, the interest rate response is muted if the monetary authority is sufficiently accommodative. As a result, aggregate demand keeps rising. Firms, satisfying demand, respond by increasing output and employment, reversing the negative pressure on labor supply that arises because of the income effect. Consequently, wages increase, but inflation remains low because firms anticipate greater productivity and lower marginal costs in the future. As evidenced by variable $\hat{v}_n$, which captures the value of new-technology firms, IT companies experience a stock market boom along the expansion. These dynamic effects are reversed when the crash occurs and agents realize that the new technology is actually of low quality. While agents mostly abandon the new technology, a recession occurs because agents wake up after having invested too much in IT capital and not enough in the traditional capital. This creates misallocation and a negative income effect which puts downward pressure on aggregate demand, despite a recovery in traditional investment.

Several comments are in order at this point. In line with our objective, the model is able to generate comovement across macroeconomic aggregates (e.g., c, i, h and y), thus providing a theoretical narrative for this type of boom and bust cycles. Less appealing is, however, the fact that the current calibration is unable to generate a slow boom/sudden crash pattern. Given our current parametrization, the economy swings very rapidly between extreme regions where m is either close to 0 or 1. Second, while the model is able to generate a recession with a significant peak-to-trough gap (about 1.5%), it remains smaller than the one in the data (about 3%). This result seems, however, a feature of belief-driven cycles that our model shares with the news/noise-driven business cycle literature. We also compute the frequency at which boom-and-bust cycle arise in our model. While the existing consensus is that such cycles are rare in models with rational agents (for instance, Avery and Zemsky (1998) in a traditional model of herding must push their probability to 10^{-6}), we find that boom-bust cycles can arise in our framework at a frequency as high as 16% when varying σ_ξ. We view this number as quite encouraging for the ability of rational herding models in explaining the data.

5 Conclusion

This paper explores whether rational herding can generate endogenous business cycle fluctuations. We propose a novel theory of herding which captures many essential features of more traditional models (Banerjee, 1992; Bikhchandani et al., 1992; Chamley, 2004), while being tractable enough to be embedded into a general equilibrium business cycle framework.

We show that the model is able to endogenously generate a boom-and-bust pattern out of a single belief shock without the need for a particular sequence of shocks. Our model has predictions on the frequency, the timing and the conditions under which such cycles emerge or burst. It can thus be used to analyze the role of stabilization policy, including investment-specific taxes or monetary policy.

We have restricted our attention to technology-driven boom-and-bust cycles for the sake of precision, butthe implications of the theory go beyond this context and we believe our herding model can be used in other environments to analyze herding behavior following any sort of innovation, be it financial innovations or innovations to the demand for certain types of goods (new products, housing).

Several extensions are worth investigating. First, our current macroeconomic model ignores the role of debt. An interesting extension would be to study how the rising pattern of optimism during the growth stage of the cycle could relax financial constraints and lead to an expansion in credit, triggering a wave of bankruptcies at the time of the crash. Another natural extension would be to consider a financial market application of our herding model and examine, in particular, the role of speculation. We leave these ideas to future research.

References

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

Amador, M. and P.-O. Weill (2012). Learning from private and public observations of others' actions. *Journal of Economic Theory 147*(3), 910–940.

Angeletos, G.-M. and J. La'O (2013). Sentiments. *Econometrica 81*(2), 739–779.

Angeletos, G.-M. and A. Pavan (2007). Efficient use of information and social value of information. *Econometrica 75*(4), 1103–1142.

Asriyan, V., L. Fornaro, A. Martin, and J. Ventura (2019). Monetary policy for a bubbly world. *working paper*.

Avery, C. and P. Zemsky (1998). Multidimensional uncertainty and herd behavior in financial markets. *American economic review*, 724–748.

Banerjee, A. V. (1992). A simple model of herd behavior. *The quarterly journal of economics 107*(3), 797–817.

Barro, R. J. and R. G. King (1984). Time-separable preferences and intertemporal-substitution models of business cycles. *The Quarterly Journal of Economics 99*(4), 817–839.

Beaudry, P., D. Galizia, and F. Portier (2020). Putting the cycle back into business cycle analysis. *American Economic Review 110*(1), 1–47.

Beaudry, P. and F. Portier (2004). An exploration into pigou's theory of cycles. *Journal of monetary Economics 51*(6), 1183–1216.

Beaudry, P. and F. Portier (2006). Stock prices, news, and economic fluctuations. *American Economic Review 96*(4), 1293–1307.

Beaudry, P. and F. Portier (2014). News-driven business cycles: Insights and challenges. *Journal of Economic Literature 52*(4), 993–1074.

Benhabib, J. (1992). *Cycles and chaos in economic equilibrium.* Princeton University Press.

Benhima, K. (2019). Booms and busts with dispersed information. *Journal of Monetary Economics 107*, 32–47.

Bikhchandani, S., D. Hirshleifer, and I. Welch (1992). A theory of fads, fashion, custom, and cultural change as informational cascades. *Journal of political Economy 100*(5), 992–1026.

Boddy, R. and M. Gort (1971). The substitution of capital for capital. *The Review of Economics and Statistics*, 179–188.

Boldrin, M. and M. Woodford (1990). Equilibrium models displaying endogenous fluctuations and chaos: a survey. *Journal of Monetary Economics 25*(2), 189–222.

Caplin, A. and J. Leahy (1994). Business as usual, market crashes, and wisdom after the fact. *The American economic review*, 548–565.

Chamley, C. P. (2004). *Rational herds: Economic models of social learning.* Cambridge University Press.

Chetty, R., A. Guren, D. Manoli, and A. Weber (2011). Are micro and macro labor supply elasticities consistent? a review of evidence on the intensive and extensive margins. *American Economic Review 101*(3), 471–75.

Christiano, L. J., M. Eichenbaum, and C. L. Evans (2005). Nominal rigidities and the dynamic effects of a shock to monetary policy. *Journal of political Economy 113*(1), 1–45.

Christiano, L. J., C. L. Ilut, R. Motto, and M. Rostagno (2008). Monetary policy and stock market boom-bust cycles.

Cipriani, M. and A. Guarino (2008). Herd behavior and contagion in financial markets. *The BE Journal of Theoretical Economics 8*(1).

Clarida, R., J. Gali, and M. Gertler (2000). Monetary policy rules and macroeconomic stability: evidence and some theory. *The Quarterly journal of economics 115*(1), 147–180.

David, J. M., H. A. Hopenhayn, and V. Venkateswaran (2016). Information, misallocation, and aggregate productivity. *The Quarterly Journal of Economics 131*(2), 943–1005.

Fajgelbaum, P. D., E. Schaal, and M. Taschereau-Dumouchel (2017). Uncertainty traps. *The Quarterly Journal of Economics 132*(4), 1641–1692.

Galí, J. (2014). Monetary policy and rational asset price bubbles. *American Economic Review 104*(3), 721–52.

Guesnerie, R. and M. Woodford (1992). Endogenous fluctuations. *Advances in economic theory 2*, 289–412.

Jaimovich, N. and S. Rebelo (2009). Can news about the future drive the business cycle? *American Economic Review 99*(4), 1097–1118.

Kozlowski, J., L. Veldkamp, and V. Venkateswaran (2019). The tail that wags the economy: Beliefs and persistent stagnation. *Journal of Political Economy forthcoming*.

Loisel, O., A. Pommeret, and F. Portier (2012). Monetary policy and herd behavior: leaning against bubbles.

Lorenzoni, G. (2009). A theory of demand shocks. *American Economic Review 99*(5), 2050–84.

Martin, A. and J. Ventura (2016). Managing credit bubbles. *Journal of the European Economic Association 14*(3), 753–789.

Samuelson, P. A. (1958). An exact consumption-loan model of interest with or without the social contrivance of money. *Journal of political economy 66*(6), 467–482.

Schmitt-Grohé, S. and M. Uribe (2012). What's news in business cycles. *Econometrica 80*(6), 2733–2764.

Tirole, J. (1985). Asset bubbles and overlapping generations. *Econometrica: Journal of the Econometric Society*, 1499–1528.

Townsend, R. M. (1983). Forecasting the forecasts of others. *Journal of Political Economy 91*(4), 546–588.

Veldkamp, L. L. (2005). Slow boom, sudden crash. *Journal of Economic theory 124*(2), 230–257.

Vives, X. (1993). How fast do rational agents learn? *The Review of Economic Studies 60*(2), 329–347.

Vives, X. (1997). Learning from others: a welfare analysis. *Games and Economic Behavior 20*(2), 177–200.

Woodford, M. (2003). Imperfect common knowledge and the effects of monetary policy. *Knowledge, Information, and Expectations in Modern Macroeconomics: In Honor of Edmund S. Phelps*, 25.

Xiong, W. (2013). *Bubbles, Crises, and Heterogeneous Beliefs*, pp. 663–713. Cambridge University Press.

A Appendix of Section 2

A.1 Equations for the three-state model

This section provides the specific model equations that characterize beliefs in the three-state model. Equation (3) that builds private beliefs from the public ones becomes

$$\begin{aligned} p_{jt} &= p_j\left(p_t, q_t, s_j\right) = \frac{p_t f^s_{\theta_H}\left(s_j\right)}{p_t f^s_{\theta_H}\left(s_j\right) + q_t f^s_{\theta_L+\Delta}\left(s_j\right) + \left(1-p_t-q_t\right) f^s_{\theta_L}\left(s_j\right)}, \\ q_{jt} &= q_j\left(p_t, q_t, s_j\right) = \frac{q_t f^s_{\theta_L+\Delta}\left(s_j\right)}{p_t f^s_{\theta_H}\left(s_j\right) + q_t f^s_{\theta_L+\Delta}\left(s_j\right) + \left(1-p_t-q_t\right) f^s_{\theta_L}\left(s_j\right)}. \end{aligned} \tag{22}$$

Equation (9) that defines the interim beliefs after observing R_t is simply

$$\begin{aligned} p_{t|R_t} &= \frac{p_t f^u\left(R_t-\theta_H\right)}{p_t f^u\left(R_t-\theta_H\right) + \left(1-p_t\right) f^u\left(R_t-\theta_L\right)}, \\ q_{t|R_t} &= \frac{q_t f^u\left(R_t-\theta_L\right)}{p_t f^u\left(R_t-\theta_H\right) + \left(1-p_t\right) f^u\left(R_t-\theta_L\right)}. \end{aligned}$$

Finally, in the three state model, the optimal investment strategy characterized by Equation (10) that defines the law of motion of beliefs after observing m_t becomes

$$\begin{aligned} p_{t+1} &= \frac{p_{t|R_t} f^\varepsilon\left(m_t - \overline{F}^s_{\theta_H}\left(\hat{s}\left(p_t,q_t\right)\right)\right)}{p_{t|R_t} f^\varepsilon\left(m_t - \overline{F}^s_{\theta_H}\left(\hat{s}\left(p_t,q_t\right)\right)\right) + q_{t|R_t} f^\varepsilon\left(m_t - \overline{F}^s_{\theta_L+\Delta}\left(\hat{s}\left(p_t,q_t\right)\right)\right) + \left(1-p_{t|R_t}-q_{t|R_t}\right) f^\varepsilon\left(m_t - \overline{F}^s_{\theta_L}\left(\hat{s}\left(p_t,q_t\right)\right)\right)}, \\ q_{t+1} &= \frac{q_{t|R_t} f^\varepsilon\left(m_t - m_e\left(p_{t|R_t}, q_{t|R_t}, \theta_L, \Delta\right)\right)}{p_{t|R_t} f^\varepsilon\left(m_t - \overline{F}^s_{\theta_H}\left(\hat{s}\left(p_t,q_t\right)\right)\right) + q_{t|R_t} f^\varepsilon\left(m_t - \overline{F}^s_{\theta_L+\Delta}\left(\hat{s}\left(p_t,q_t\right)\right)\right) + \left(1-p_{t|R_t}-q_{t|R_t}\right) f^\varepsilon\left(m_t - \overline{F}^s_{\theta_L}\left(\hat{s}\left(p_t,q_t\right)\right)\right)}. \end{aligned}$$

A.2 Propositions

Proposition 1. *There exists a unique equilibrium.*

Proof. The threshold $\hat{p}$ is uniquely determined by (6). Fix the fundamental (θ, ξ) and the realization of the shocks $\{u_0, \varepsilon_0, u_1, \varepsilon_1, \dots\}$. Given public beliefs π_t, (3) and (5) yield a unique distribution of private beliefs $\{p_{jt}\}_{j\in[0,1]}$. Given these, there is a unique m^e_t, derived from (8) and, therefore a unique m_t from (7). As a result, updating beliefs through (9) and (10) yield unique $\pi_{t|R}$ and π_{t+1}. We have shown that the updating of public beliefs yields a unique π_{t+1} from any π_t. Starting from arbitrary public beliefs π_0, there is therefore a unique equilibrium path $\{\pi_0, \pi_1, \dots\}$, and all other quantities can be uniquely determined from it. □

Lemma 1. *In the three-state model, for $\theta_L < \theta_L + \Delta < \theta_H$ and $\{F^s_x\}$ satisfying the MLRP condition, the optimal investment strategy in characterized by a cutoff rule in the private signal $\hat{s}\left(p_t, q_t\right)$, decreasing in p_t. That is, an agent invests if and only if $s_j \geq \hat{s}\left(p_t, q_t\right)$. The expected measure of investing agents is given by*

$$m^e\left(p_t, q_t, \theta, \xi\right) = \overline{F}^s_{\theta+\xi}\left(\hat{s}\left(p_t, q_t\right)\right).$$

Proof. The proof is straightforward. Under the above conditions, rewrite the individual probability of the high-technology state as

$$p_j\left(p_t,q_t,s_j\right)=\frac{p_t}{p_t+q_t\frac{f^s_{\theta_L+\Delta}\left(s_j\right)}{f^s_{\theta_H}\left(s_j\right)}+\left(1-p_t-q_t\right)\frac{f^s_{\theta_L}\left(s_j\right)}{f^s_{\theta_H}\left(s_j\right)}}.$$

Under the assumption of MLRP and $\theta_L<\theta_L+\Delta<\theta_H$, p_j is clearly increasing in s_j. Hence, for all (p_t,q_t), there exists a cutoff $\hat{s}\left(p_t,q_t\right)\in\mathbb{R}\cup\{-\infty,\infty\}$ such that $s_j\geq\hat{s}\left(p_t,q_t\right)\Leftrightarrow p_j\left(p_t,q_t,s_j\right)\geq\hat{p}$. Also, because p_j is increasing in p_t, the implicit function theorem ensures that $\hat{s}\left(p_t,q_t\right)$ is decreasing in p_t. The measure of investing agents is thus

$$m^e\left(p_t,q_t,\theta,\xi\right)=\int\mathbb{I}\left(p_j\left(p_t,q_t,s_j\right)\geq\hat{p}\right)f^s_{\theta+\xi}\left(s_j\right)ds_j=\overline{F}^s_{\theta+\xi}\left(\hat{s}\left(p_t,q_t\right)\right).$$

□

Proposition 2. *In the Gaussian case, i.e.,* $F^\xi\sim\mathcal{N}\left(0,\sigma_\xi^2\right)$, $F^s|\theta,\xi\sim\mathcal{N}\left(\theta+\xi,\sigma_s^2\right)$, $F^\varepsilon\sim\mathcal{N}\left(0,\sigma_\varepsilon^2\right)$, $F^u\sim\mathcal{N}\left(0,\sigma_u^2\right)$, *for* θ *and* ξ *independent and signal* R_t *sufficiently uninformative* (σ_u *low), there exists a large enough* $\underline{\xi}$ *such that all shocks* $\xi\geq\underline{\xi}$ *generate a boom-and-bust cycle in the impulse response of beliefs* p_t *to a false-positive shock* (θ_L,ξ).

Proof. Our strategy is to show that there exists a sufficiently large $\underline{\xi}$, such that for all shock $\xi\geq\underline{\xi}$ the public beliefs about the high state in date 1, p_1, increases after observing m_0. Since beliefs must converge to the truth in the long-run, due to the strictly positive flow of information, and the law of large numbers, this guarantees the existence of a boom-and-bust cycle. We start under the assumption that R_t is totally uninformative, $\sigma_u=\infty$.

First, we establish that the optimal strategy in the Gaussian case follows a cutoff strategy in $\hat{s}$. The probability that individual j puts on the high state is given by

$$p_j\left(p_0,s_j\right)=\frac{\int\pi_0\left(\theta_H,\xi\right)f^s_{\theta_H+\xi}\left(s_j\right)d\xi}{\int\pi_0\left(\theta_H,\xi\right)f^s_{\theta_H+\xi}\left(s_j\right)d\xi+\int\pi_0\left(\theta_L,\xi\right)f^s_{\theta_L+\xi}\left(s_j\right)d\xi}.$$

Since ξ is independent from θ, $\pi_0\left(\theta_H,\xi\right)=p_0f^\xi\left(\xi\right)$ and $\pi_0\left(\theta_L,\xi\right)=\left(1-p_0\right)f^\xi\left(\xi\right)$. Notice, then, that $\int f^\xi\left(\xi\right)f^s_{\theta+\xi}\left(s_j\right)d\xi$ is the pdf of s_j given θ, which is a normal, $s_j|\theta\sim\mathcal{N}\left(\theta,\sigma_\xi^2+\sigma_s^2\right)$. Denote ϕ the pdf of a unit normal, we have:

$$p_j\left(p_0,s_j\right)=\frac{1}{1+\frac{\left(1-p_0\right)\int f^\xi(\xi)f^s_{\theta_L+\xi}\left(s_j\right)d\xi}{p_0\int f^\xi(\xi)f^s_{\theta_H+\xi}\left(s_j\right)d\xi}}=\frac{1}{1+\frac{1-p_0}{p_0}\phi\left(\frac{s_j-\theta_L}{\sqrt{\sigma_\xi^2+\sigma_s^2}}\right)/\phi\left(\frac{s_j-\theta_H}{\sqrt{\sigma_\xi^2+\sigma_s^2}}\right)}.$$

Since the Gaussian family satisfies the MLRP property, p_j is increasing in s_j. Hence, the optimal investment strategy at date 0 takes a cutoff form $\hat{s}_0$.

Under the assumption that R_t is uninformative, the public belief about the high state at the beginning of period 1, p_1, is given by

$$p_1=\int\pi_1\left(\theta_H,\xi\right)d\xi=\frac{\int\pi_0\left(\theta_H,\xi\right)f^\varepsilon\left(m_0-m^e\left(\pi_0,\theta_H,\xi\right)\right)d\xi}{\int\pi_0\left(\theta_H,\xi\right)f^\varepsilon\left(m_0-m^e\left(\pi_0,\theta_H,\xi\right)\right)d\xi+\int\pi_0\left(\theta_L,\xi\right)f^\varepsilon\left(m_0-m^e\left(\pi_0,\theta_L,\xi\right)\right)d\xi}.$$

Using the independence property between θ and ξ and the cutoff property, the above formula can be rewritten as

$$p_1=\frac{1}{1+\frac{1-p_0}{p_0}\frac{\int f^\xi(\xi)f^\varepsilon\left(m_0-\overline{F}^s_{\theta_L+\xi}(\hat{s}_0)\right)d\xi}{\int f^\xi(\xi)f^\varepsilon\left(m_0-\overline{F}^s_{\theta_H+\xi}(\hat{s}_0)\right)d\xi}}.$$

Denoting ξ_0 the true shock, the impulse response in m_t yields $m_0=\overline{F}^s_{\theta_L+\xi_0}\left(\hat{s}_0\right)$, which goes to 1 as $\xi_0\to\infty$. Because the MLRP property implies first-order stochastic dominance, we have $\overline{F}^s_{\theta_L+\xi}\left(\hat{s}_0\right)<\overline{F}^s_{\theta_H+\xi}\left(\hat{s}_0\right)$. Since $f^\varepsilon\left(\varepsilon\right)$ is decreasing for $\varepsilon\geq0$, we have

$$f^\varepsilon\left(m_0-\overline{F}^s_{\theta_L+\xi}\left(\hat{s}_0\right)\right)<f^\varepsilon\left(m_0-\overline{F}^s_{\theta_H+\xi}\left(\hat{s}_0\right)\right)$$

for all $\xi \leq \theta_L - \theta_H + \xi_0$. Decompose the difference between the denominator and numerator can be written

$$\int f^{\xi}(\xi) f^{\varepsilon}\left(m_0 - \overline{F}^s_{\theta_H+\xi}(\hat{s}_0)\right) d\xi - \int f^{\xi}(\xi) f^{\varepsilon}\left(m_0 - \overline{F}^s_{\theta_L+\xi}(\hat{s}_0)\right) d\xi$$

$$\underset{\xi_0 \to \infty}{\longrightarrow} \int_{-\infty}^{\infty} f^{\xi}(\xi)\left[f^{\varepsilon}\left(1 - \overline{F}^s_{\theta_H+\xi}(\hat{s}_0)\right) - f^{\varepsilon}\left(1 - \overline{F}^s_{\theta_L+\xi}(\hat{s}_0)\right)\right] d\xi > 0$$

The difference converges to a strictly positive term. Thus, there exists $\underline{\xi}$ such that for all $\xi_0 > \underline{\xi}$

$$\int f^{\xi}(\xi) f^{\varepsilon}\left(\overline{F}^s_{\theta_L+\xi_0}(\hat{s}_0) - \overline{F}^s_{\theta_L+\xi}(\hat{s}_0)\right) d\xi < \int f^{\xi}(\xi) f^{\varepsilon}\left(\overline{F}^s_{\theta_L+\xi_0}(\hat{s}_0) - \overline{F}^s_{\theta_H+\xi}(\hat{s}_0)\right) d\xi$$

and $p_1 > p_0$. The shock is large enough for agents to attribute it mostly to the high state, initiating the growth stage of the cycle. By continuity of the belief updating equations in σ_u. There must also exists a sufficiently large σ_u (R_t sufficiently uninformative) for which $p_1 > p_0$ after $\xi_0 \geq \underline{\xi}$. □

Proposition 3. *The efficient allocation can be implemented as an equilibrium by an investment tax*

$$\tau^* = \left(E_{\theta,\xi}\left[f^p_{\theta+\xi}(\hat{p}) \mid \mathcal{I}\right]\right)^{-1} \beta \frac{\partial E_{\theta,\xi}\left[V(\mathcal{I}') \mid \mathcal{I}\right]}{\partial \hat{p}}, \qquad ((13))$$

and a lump-sum transfer to all investors.

Proof. We consider a tax τ that makes the effective cost of investing $c + \tau$. Under that tax, (6) shows that the marginal investor $\hat{p}$ is such that $\hat{p}\theta_H + (1-\hat{p})\theta_L = c + \tau$. Combining with (12) and reorganizing yields (13). □

A.3 Figures

(a) Measure of investing agents (b) Evolution of public beliefs

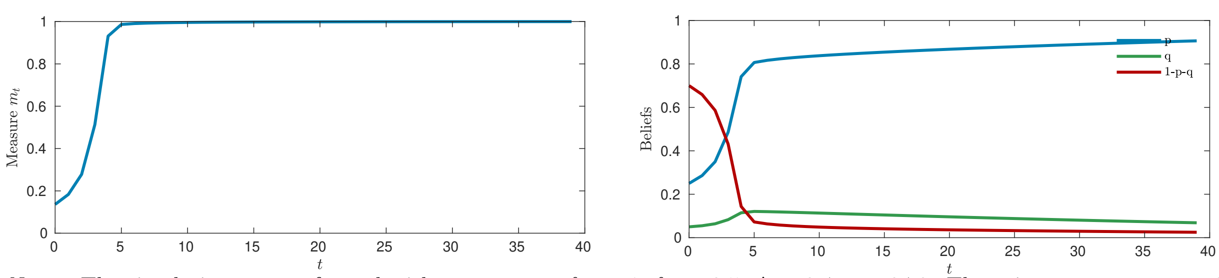

Notes: The simulation was performed with parameters: $\theta_H = 1$, $\theta_L = 0.5$, $\Delta = 0.4$, $c = 0.79$. The priors are set to $p_0 = 0.25$ and $q_0 = 0.05$. The distributions of all signals are Gaussian with standard deviations $\sigma_s = 0.4$, $\sigma_\varepsilon = 0.2$, $\sigma_u = 2.5$. The shaded area corresponds to the informational cascade period, loosely defined as periods when $m_t \geq 85\%$ in which learning is markedly slower, given our calibration.

Figure 9: Impulse response in the case of a true positive

(a) Measure of investing agents

(b) Evolution of public beliefs

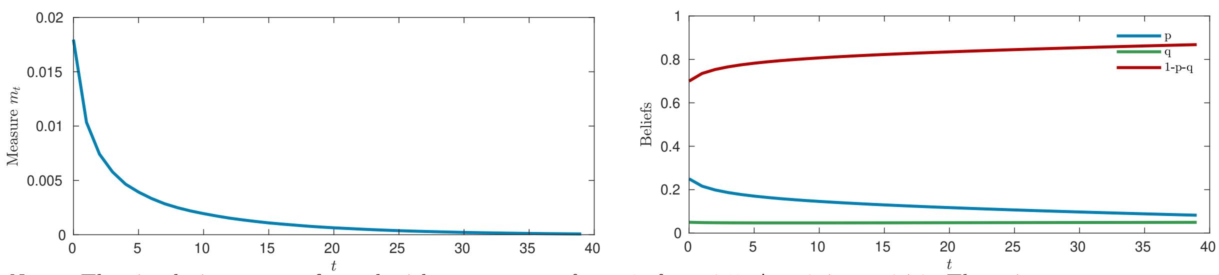

Notes: The simulation was performed with parameters: $\theta_H = 1$, $\theta_L = 0.5$, $\Delta = 0.4$, $c = 0.79$. The priors are set to $p_0 = 0.25$ and $q_0 = 0.05$. The distributions of all signals are Gaussian with standard deviations $\sigma_s = 0.4$, $\sigma_\varepsilon = 0.2$, $\sigma_u = 2.5$. The shaded area corresponds to the informational cascade period, loosely defined as periods when $m_t \geq 85\%$ in which learning is markedly slower, given our calibration.

Figure 10: Impulse response in the case of a true negative