

DISCUSSION PAPER SERIES

DP16368

Herding Through Booms and Busts

Edouard Schaal and Mathieu Taschereau-
Dumouchel

MONETARY ECONOMICS AND FLUCTUATIONS

CEPR

Herding Through Booms and Busts

Edouard Schaal and Mathieu Taschereau-Dumouchel

Discussion Paper DP16368

Published 17 July 2021

Submitted 15 July 2021

Centre for Economic Policy Research
33 Great Sutton Street, London EC1V 0DX, UK
Tel: +44 (0)20 7183 8801
www.cepr.org

This Discussion Paper is issued under the auspices of the Centre's research programmes:

- Monetary Economics and Fluctuations

Any opinions expressed here are those of the author(s) and not those of the Centre for Economic Policy Research. Research disseminated by CEPR may include views on policy, but the Centre itself takes no institutional policy positions.

The Centre for Economic Policy Research was established in 1983 as an educational charity, to promote independent analysis and public discussion of open economies and the relations among them. It is pluralist and non-partisan, bringing economic research to bear on the analysis of medium- and long-run policy questions.

These Discussion Papers often represent preliminary or incomplete work, circulated to encourage discussion and comment. Citation and use of such a paper should take account of its provisional character.

Copyright: Edouard Schaal and Mathieu Taschereau-Dumouchel

Herding Through Booms and Busts

Abstract

This paper explores whether rational herding can generate endogenous aggregate fluctuations. We embed a tractable model of rational herding into a business cycle framework. In the model, technological innovations arrive with unknown qualities and agents have dispersed information about how productive the technology really is. Rational investors decide whether to invest based on their private information and the investment behavior of others. Herd-driven boom-bust cycles arise endogenously in this environment when the technology is unproductive but investors' initial information is unusually optimistic. Their overoptimism leads to high investment rates, which investors mistakenly attribute to good fundamentals, leading to a self-reinforcing pattern of higher optimism and higher investment until the economy reaches a peak, followed by a crash when agents ultimately realize their mistake. We calibrate the model to the U.S. economy and show that it can explain boom-and-bust cycles in line with episodes like the dot-com bubble of the 1990s.

JEL Classification: E32, D80

Keywords: Herding, booms and busts, bubbles, endogenous business cycles

Edouard Schaal - eschaal@crei.cat

Centre de Recerca en Economia Internacional (CREI), Universitat Pompeu Fabra and CEPR

Mathieu Taschereau-Dumouchel - mt763@cornell.edu

Cornell University

Herding Through Booms and Busts*

Edouard Schaal

CREI, ICREA, UPF, BGSE and CEPR

Mathieu Taschereau-Dumouchel

Cornell University

July 2021

Abstract

This paper explores whether rational herding can generate endogenous aggregate fluctuations. We embed a tractable model of rational herding into a business cycle framework. In the model, technological innovations arrive with unknown qualities and agents have dispersed information about how productive the technology really is. Rational investors decide whether to invest based on their private information and the investment behavior of others. Herd-driven boom-bust cycles arise endogenously in this environment when the technology is unproductive but investors' initial information is unusually optimistic. Their overoptimism leads to high investment rates, which investors mistakenly attribute to good fundamentals, leading to a self-reinforcing pattern of higher optimism and higher investment until the economy reaches a peak, followed by a crash when agents ultimately realize their mistake. We calibrate the model to the U.S. economy and show that it can explain boom-and-bust cycles in line with episodes like the dot-com bubble of the 1990s.

JEL Classification: E32, D80

*E-mail: eschaal@crei.cat, mt763@cornell.edu. This paper was previously circulated under the title "Herding Cycles". 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 participants at various seminars and conferences. Edouard Schaal acknowledges financial support from the Spanish Ministry of Economy and Competitiveness, through the Severo Ochoa Programme for Centres of Excellence in R&D (CEX2019-000915-S) and through the program Proyecto I+D Excelencia 2017 (grant ECO2017-87827-P-AEI/FEDER, UE), and support from the Generalitat de Catalunya, through CERCA and SGR Programme (2017-SGR-1393).

1 Introduction

New technologies sometimes lead to periods of massive investment that end in severe economic downturns. One salient example is the 1990s boom in information technologies that culminated in 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 clear in the 1990s, when extreme enthusiasm for the new technology led to large investments in communication networks, software, and IT equipment. This high volume of investment and the rising valuation of IT companies initially seemed to validate the optimistic outlook, but the crash eventually followed as some of the expected returns failed to materialize.¹ While there is no consensus about the deep drivers behind these fluctuations, a common view is that shifts in expectation played a key role in shaping the dot-com boom-bust cycles.

The idea that expectations contribute to aggregate economic fluctuations has a long tradition in macroeconomics. In his seminal work on the origin of economic fluctuations, [Pigou \(1927\)](#) emphasized the importance of beliefs in shaping the business cycle. In his view, booms can be caused by waves of optimism among business executives, and crashes arise when their lofty expectations turn out to have been mistaken. This hypothesis has been extensively studied in modern business cycle theory by the news-driven business cycle literature, pioneered by [Beaudry and Portier \(2004\)](#).² According to this view, agents receive news about future productivity, which sometimes turn out to be false. Boom-bust cycles arise after an initial sequence of positive news is later contradicted by experience.

These theories, however, remain mostly silent on the technological, social and psychological determinants that drive the evolution of beliefs. In most of these studies, the belief process obeys an exogenous law, and boom-bust cycles occur after a specific sequence of shocks. In contrast to the standard news literature, other explanations emphasize the social dimension of beliefs formation. Pigou, in the earlier mentioned volume, highlighted the important role that the behavior of others played in shaping these expectations.³ He suggested that executives infer business conditions by observing the decisions of other business leaders, so that they might be tempted to invest when they see their competitors expand their operations.⁴ This *herding* process can thus propel the economy

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

²See [Beaudry and Portier \(2014\)](#) for an overview of the empirical and theoretical research supporting the news view of the business cycle.

³In *Industrial Fluctuations (1927)*, Pigou states that “the varying expectations of business men [...] and not anything else, constitute the immediate and direct causes or antecedents of industrial fluctuation”. He also emphasizes the importance of the herding process: “the pioneers, who thus undertake and expand enterprises, at once fill a social need and lay treasure for themselves. Gradually, as no disaster happens to them, other less bold spirits follow their example; then others and yet other.”

⁴Other studies ([Blasco et al., 2012](#); [Galariotis et al., 2015](#)) document the presence of herding behavior along the

into a boom and precipitate the crash in a downturn, thereby providing a potential explanation for the emergence of technology-driven boom-bust cycles

This paper considers herding as a source of boom-and-bust cycles. We propose a theory in which a new technology of uncertain quality can trigger a slow-rising boom followed by a sudden crash, in line with the experience of the dot-com era. In the boom phase, the initial optimism of investors translates into high levels of aggregate investment, and that high investment, in turn, leads to further increases in optimism. This self-reinforcing process can fuel a long-lasting expansion of the economy, which comes to an end when the data no longer support an optimistic view of the technology. Agents then stop investing and the economy rapidly collapses. Our theory therefore captures the importance of expectations and the role of herding that were highlighted by [Pigou \(1927\)](#).

In the model, random technological innovations arrive over time and a group of rational 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. Importantly, to capture the idea that investors collect information from similar sources (news media, market reports, etc), we assume that private signals feature some common noise. This assumption is key as it allows the distribution of beliefs across investors to vary for reasons unrelated to the fundamental value of the technology. Investors do not initially know the extent of that bias but progressively learn about it over time.

Agents also receive public signals. First, they can learn by observing the exogenous return on investment, which provides noisy information about the technology. They can also learn from 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, in a non-linear fashion, 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”: bad realizations of the technology fundamental that are accompanied by unusually large and positive realizations of the common noise. False-positives may thus capture situations in which, for instance, excessively promising benchmark tests are widely advertised upon the introduction of the technology and lead to overly optimistic beliefs. 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 good, or because the common noise component of the private signals is high. Investors cannot tell these stories apart but, as long as the false-positive

development of several historical boom-bust cycles.

shocks are relatively rare, the high level of investment is mainly attributed to the good-technology state, whose posterior likelihood rises. 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. It is in that sense that our model displays a form of herding: agents mimic the behavior of others and sometimes mistakenly follow the herd into an investment boom, meanwhile a shrinking measure of agents use their private information to go against the crowd. 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 in their assessment. As a result, they keep track of the probability of being in a false-positive state, which appears increasingly likely as investment keeps falling slightly short of the most optimistic predictions. At some point, the most pessimistic agents stop investing and aggregate investment no longer supports a high-productivity scenario, leading to a reversal in beliefs and a collapse of investment. We provide formal conditions under which these boom and bust episodes are guaranteed to arise.

A distinguishing feature of our approach is that the boom-and-bust cycle emerges endogenously. 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. In contrast, this paper explores a fundamentally different view of aggregate fluctuations, in which both phases of the cycle are intrinsically linked, such that the “booms sow the seeds of the subsequent busts” (Beaudry et al., 2020). Indeed, our mechanism generates an *endogenous* boom-and-bust cycle out of the single impulse shock that is the arrival of the new technology.⁵ The crash, in particular, is not triggered by an exogenous shock but arises endogenously through the natural evolution of beliefs. As a result, the properties of the bust can be affected by what happened during the preceding boom, and policy interventions can have a particularly strong impact on the duration and magnitude of the boom as well as the timing and depth of the bust.

In the model, the mass of investing agent is a nonlinear aggregator of the dispersed information. As a result, 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 mass of investors. Similarly, after a series of bad news, agents rely less on their private information when deciding not to invest, making aggregate investment also less informative. As a result, the model is able to generate sustained booms

⁵By “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).

and rapid busts as periods when massive investment restricts the flow of information, but slight downturns can suddenly reveal more information on the true state of the world and trigger an economic collapse.

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 quantitative business cycle framework, 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, social learning takes place as agents observe the measure of new-technology adopters.

We calibrate 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 comovements in consumption, investment, hours worked and output. The overinvestment into IT capital during the boom period causes the economy to contract significantly when beliefs collapse as agents realize that resources were misallocated. Overall, our results suggest that rational herding among economic agent can account for significant fluctuations in macroeconomic aggregates.

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 it 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 is closely related to the literature on news or noise-driven business cycles (Beaudry and Portier, 2004; 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 generates 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. In addition the news literature does not consider the role of herding in driving fluctuations. Christiano et al. (2008) consider 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 the truth is revealed. Burnside et al. (2016) construct an epidemiological model in which the transmission of optimistic beliefs in a population about the housing market can create a boom and bust.

As in Caplin and Leahy (1994) and Veldkamp (2005), the endogenous release of information in our model can generate sudden collapses in economic activity. Angeletos and La'O (2013) construct a model in which variations in sentiments can lead to aggregate fluctuations. Our work also contribute to a literature in which the aggregation of private information leads to nonlinear aggregate dynamics (Fajgelbaum et al., 2017).

Our paper also relates to the original work on herding and information cascades by Banerjee (1992), Bikhchandani et al. (1992) and Chamley (2004). Our model differs from these traditional models of herding in several dimensions. First, in previous herding models, agents make decisions sequentially and the dynamics of the model are governed by the gradual observation of these individual decisions. Both of these features do not sit well with standard macroeconomic models, which explains, we believe, why herding does not figure more prominently in this literature. In our setup, instead, agents act simultaneously and learn by observing aggregates, which allows for a smoother integration of herding mechanisms into standard macroeconomic models. Second, the source of agents' confusion is different. In traditional herding models, people are confused between the fundamental return and the idiosyncratic shocks that stem from a particular ordering of the investors. As a consequence, boom-bust cycles arise in those models only for specific orderings, or in other words for specific sequences of idiosyncratic shocks. In our model instead, agents are confused between between the fundamental and the common noise, which are drawn once and for all. Boom-bust patterns emerge endogenously through the natural evolution of beliefs and without any timing assumptions about shocks. This distinction with the existing literature is crucial to generate endogenous cycles. Finally, while most of the earlier literature on herding focused on a strict type of permanent information cascade in which the flow of information is exactly zero,

our model features a “smooth” form of information cascades in which the information flow varies continuously 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.⁶

Our learning model shares similarities with [Vives \(1997\)](#) who studies an environment in which agents with dispersed information learn by observing the average action across agents.⁷ Chapter 4 of [Chamley \(2004\)](#) briefly reviews a model 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. [Straub and Ulbricht \(2019\)](#) explore in a general setup the informativeness of nonlinear public signals and [Straub and Ulbricht \(2017\)](#) studies a particular application with financial constraints. None of these works consider the emergence of endogenous boom-bust cycles. Our model also relates to [Avery and Zemsky \(1998\)](#), who study herding in financial markets and introduce multidimensional uncertainty to allow for information cascades.

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 must invest in a risky asset. As in traditional models of herding, entrepreneurs act sequentially and individual investment decisions are publicly observable. Aggregate output fluctuates as single agents gradually make their investment decisions. Our paper extends this approach by offering a novel herding model, based on contemporaneous decisions and the observation of aggregate actions, which, we believe, can be more easily integrated into traditional macroeconomic models.

Our work also relates to a strand of literature that studies the role of bubbles in macroeconomic environments. These includes studies based on rational bubbles ([Galí, 2014](#); [Martin and Ventura, 2016](#); [Asriyan et al., 2019](#); [Guerrón-Quintana et al., 2020](#)) as well as bubbles due to financial constraints and others ([Kocherlakota, 1992](#); [Miao and Wang, 2012](#); [Barlevy, 2014](#); [Hirano and Yanagawa, 2016](#)).

Our paper also relates to the literature on endogenous deterministic cycles, which includes [Grandmont \(1985\)](#), [Boldrin and Woodford \(1990\)](#), [Benhabib \(1992\)](#), [Benhabib et al. \(2002\)](#) and [Matsuyama \(1999; 2013\)](#). More recent contributions to this literature include [Beaudry et al. \(2020\)](#), who provide empirical evidence in favor of endogenous cycles. In contrast to our setup, models in this literature feature deterministic limit cycles that can be periodic or chaotic.

Another literature views endogenous cycles from the point of view of equilibrium indeterminacy and sunspot dynamics. This includes [Benhabib and Farmer \(1994\)](#) and [Wen \(1998\)](#). More recent contributions include [Benhabib et al. \(2015\)](#), [Kaplan and Menzio \(2016\)](#), [Eeckhout and Lindenlaub \(2019\)](#) and [Golosov and Menzio \(2020\)](#). These studies typically feature multiple equilibria and ag-

⁶Our model can also generate permanent information cascades when the support of the private signal distribution is bounded.

⁷Both [Vives \(1993\)](#) and [Vives \(1997\)](#) show that learning happens slowly in models in which privately informed agents learn from aggregates.

gregate fluctuations are due to shifts in expectation triggered by sunspot shocks. Our model, instead, features a unique equilibrium and boom-bust cycles result from agents' gradual learning about the technology and the common noise.

Finally, a literature studies the emergence of boom-and-bust cycles in asset prices or in aggregate economic activity after departing from rationality or rational expectations. This includes the adaptive learning literature (Carceles-Poveda and Giannitsarou, 2008; Eusepi and Preston, 2011; Adam et al., 2017) and, more recently, a literature that uses diagnostic expectations (Bordalo et al., 2021).

Section 2 introduces a simple learning model that conveys the intuition for the mechanism. The following section describes the forces at work in the model and discusses its welfare implications. Section 4 presents our business cycle model. We calibrate the model in Section 5 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. We also use this simplified model to derive analytical results and discuss the 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 \bar{F}^x its complementary CDF, $\bar{F}^x(\tilde{x}) = Pr(x > \tilde{x})$.

2.2 Environment

Time is discrete and goes on forever, $t = 0, 1, \dots$. 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

$$R_t = \theta + u_t, \tag{1}$$

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 .⁸ We refer to θ as the technology *fundamental*. As we describe below, agents do not observe θ directly but learn about it through a variety of signals.

Every period, investors must decide whether to invest in the technology ($i_{jt} = 1$) or not ($i_{jt} = 0$).⁹ Investing is costly and requires the payment every period of a cost c that is identical across agents. 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

$$y_{jt} = i_{jt} (R_t - c). \quad (2)$$

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 do not 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. This common noise might come, for instance, from sources of information shared by agents (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 varies for reasons that are orthogonal to the true value of the fundamental. In other words, common noise is what allows agents, as a group, to be sometimes overly optimistic or pessimistic about the technology. As a consequence, it is a source of confusion for the investors who must figure out whether high investment rates result from either good fundamentals or overoptimism. Boom-and-bust cycles are then caused by investors mistakenly attributing high investment rates to good fundamentals when the technology is bad in reality.

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 of agent j is drawn from the CDF $F_{\theta+\xi}^s(s) = Pr(s_j \leq s)$, where $\{F_x^s\}_{x \in I}$ is a family of distributions that admit the probability density functions $\{f_x^s\}_{x \in I}$. 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.¹⁰ Finally, in order to guarantee monotonicity in learning,

⁸While we model the fundamental θ as a shock to the productivity of an investment, it is straightforward to modify the model so that the fundamental corresponds to other shocks such as demand, changes in taxation, etc.

⁹We assume in the main text that the investment decision is binary, but we show in Appendix A.4 that the model also generates boom-bust cycles when firms can invest along an intensive margin.

¹⁰This full-support assumption rules out permanent information cascades in which, for instance, the public information is so optimistic that even the most pessimistic agent prefers to invest. In that case, aggregate investment

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

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

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 becomes more optimistic and puts a higher probability on the value of the technology θ and the common noise ξ being high.

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

$$s_j = \theta + \xi + v_j, \text{ with } v_j \sim \text{iid CDF } F^v. \quad (3)$$

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

Public signals

In addition to their initial private signal, investors collect public information over time by observing market activity and investment returns. We first assume that all agents observe the return R_t , given by (1). Since the transitory component in the investment return u_t is unknown to the agents, the total return R_t provides an exogenous noisy signal about θ , which offers 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 the private information of agents. This is to capture the type of information that agents learn by observing aggregate quantities or prices, which result from the aggregation of individual decisions.¹¹ Specifically, we assume that investors receive a noisy measure of the total number of investing agents m_t , which we define as

$$m_t = \int_0^1 i_{jt} dj + \varepsilon_t, \text{ with } \varepsilon_t \sim \text{iid CDF } F^\varepsilon. \quad (4)$$

The noise ε_t can be interpreted as coming either from measurement errors 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 from learning too quickly (or even immediately in some cases, as we discuss later).

does not reveal any private information and social learning can stop altogether. See Appendix A.5 for an example.

¹¹In particular, in our full business cycle model from Section 4, observing aggregate quantities or prices will provide a public signal of the same form.

Before going further, we would like to highlight some particular features of the signal m_t . In equilibrium, the decision to invest i_{jt} is a nonlinear function of the investor's individual beliefs. In turn, beliefs are a function of public information up to time t , $\{R_{t-1}, m_{t-1}, \dots, R_0, m_0\}$, and of the private signal s_j . As a result, since public information is shared and can be filtered out, m_t partially aggregates the private information across the population of investors. To that extent, m_t contains useful information regarding the fundamental θ and the common noise ξ .

As we explain in more details below, it is the presence of this endogenous signal that will allow for *herding* to occur in our environment. As m_t aggregates the dispersed information of private agents, a particularly high draw of m_t will be interpreted as being indicative of a strong fundamental θ and a high common noise ξ . As a result, agents will update their beliefs in favor of these states, which might lead them to invest more, which will increase m_t further, and so on.

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 it contains varies over time depending on the economy's location in the state space (e.g., history of shocks), opening up the possibility of a form of *informational cascades* in which agents rely less on their private information and aggregate investment becomes less informative. Finally, the endogeneity of m_t is the source of an information externality, which provides a basis for government intervention, as we discuss in section 3.3.

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

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

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 useful 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* $\Lambda_t(\tilde{\theta}, \tilde{\xi}) = Pr\left(\theta = \tilde{\theta}, \xi \in [\tilde{\xi}, \tilde{\xi} + d\tilde{\xi}] \mid \mathcal{I}_t\right)$. These public beliefs 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 also includes 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

$$\Lambda_{jt}(\tilde{\theta}, \tilde{\xi}) = \frac{\Lambda_t(\tilde{\theta}, \tilde{\xi}) f_{\tilde{\theta}+\tilde{\xi}}^s(s_j)}{\int \Lambda_t(\theta, \xi) f_{\theta+\xi}^s(s_j) d(\theta, \xi)}. \quad (5)$$

This simplification comes from the fact that only public information evolves over time. Indeed, since the private signals distribution $f_{\theta+\xi}^s$ 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 Λ_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 their individual beliefs Λ_{jt} ,
2. Investment returns R_t 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

$$E[R_t | \mathcal{I}_{jt}] \geq c. \quad (6)$$

Defining

$$p_{jt} = Pr(\theta = \theta_H | \mathcal{I}_{jt}) = \int \Lambda_{jt}(\theta_H, \xi) d\xi \quad (7)$$

as the probability that investor j puts on being in the good-technology state, the investment decision (6) is characterized by a cutoff rule p^* in the space of beliefs. That is, an agent invests if and only if¹² $p_{jt} \geq p^*$ where p^* is the belief of the marginal investor such that

$$p^* \theta_H + (1 - p^*) \theta_L = c. \quad (8)$$

¹²To break indifference, we assume that indifferent agents invest in the technology. This assumption is innocuous if $F_{\theta+\xi}^s$ has no mass points.

The total measure of investing agents can then be expressed as

$$m_t = m^e(\Lambda_t, \theta, \xi) + \varepsilon_t \quad (9)$$

$$\text{where } m^e(\Lambda_t, \theta, \xi) = \int \mathbb{I}(p_j(\Lambda_t, s_j) \geq p^*) f_{\theta+\xi}^s(s_j) ds_j. \quad (10)$$

The variable m^e is the expected measure of investing agents for a given state of the world (θ, ξ) , excluding the noise ε_t . Importantly for what follows, m^e is an object that any agent in the economy can compute. To see this, note that since they know the structure of the model and the public beliefs, all agents agree on the cutoff p^* . Second, thanks to the dichotomy between public beliefs and the fixed distribution of private signals $f_{\theta+\xi}^s$, 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 from the endogenous public signal, to which we now turn.

2.6 Evolution of Beliefs

After characterizing the investment decision, we can 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

$$\Lambda_{t|R_t}(\tilde{\theta}, \tilde{\xi}) = \frac{\Lambda_t(\tilde{\theta}, \tilde{\xi}) f^u(R_t - \tilde{\theta})}{\int \Lambda_t(\theta, \xi) f^u(R_t - \theta) d(\theta, \xi)}. \quad (11)$$

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 the public beliefs Λ_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 Λ_{jt} , but there is no *infinite regress* problem arising from the necessity to forecast the beliefs of agents after any history of shocks.¹³ Because of the equilibrium structure of the signal (9), Bayes' rule gives us the simple updating equation

¹³In 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.

$$\Lambda_{t+1}(\tilde{\theta}, \tilde{\xi}) = \frac{\Lambda_{t|R_t}(\tilde{\theta}, \tilde{\xi}) f^\varepsilon(m_t - m^e(\Lambda_t, \tilde{\theta}, \tilde{\xi}))}{\int \Lambda_{t|R_t}(\theta, \xi) f^\varepsilon(m_t - m^e(\Lambda_t, \theta, \xi)) d(\theta, \xi)}. \quad (12)$$

2.7 Equilibrium

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

Definition 1. An equilibrium consists of history-contingent public beliefs Λ_t , a distribution of private beliefs $\{\Lambda_{jt}\}_{j \in [0,1]}$ and a measure of investors m_t for all t , such that, 1) the distribution of private beliefs is derived from the public beliefs through (5) and (7); 2) the measure of investors is consistent with investors decisions under their private beliefs as in (9); and 3) the public beliefs follow the laws of motion (11)–(12).

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

Proposition 1. *There exists a unique equilibrium.*

The proof of the proposition is straightforward. It shows that from a given distribution of public beliefs Λ_t , there is a unique mapping, given the realization of the shocks, to next period's public beliefs Λ_{t+1} . Starting from the initial Λ_0 we can therefore reconstruct the unique equilibrium sequence $\{\Lambda_0, \Lambda_1, \dots\}$. All other equilibrium quantities such as the measure of investors and the distribution of private beliefs can then be reconstructed from the public beliefs in a unique way.

3 Endogenous Booms and Busts

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 then show that these results extend to a more general setup. We also discuss the welfare properties of the model and describe how the efficient allocation can be implemented using an investment tax.

3.1 The 3-state model

To simplify the exposition, 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

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

We refer to $(\theta_L, 0)$ as the *bad-technology* state, $(\theta_H, 0)$ as the *good-technology* state and $(\theta_L, \bar{\xi})$ as the *false-positive* state. The latter is the state of interest as it is the one 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 Λ_t . Public beliefs are now summarized by the two variables

$$p_t \equiv \Lambda_t(\theta_H, 0) \text{ and } q_t \equiv \Lambda_t(\theta_L, \bar{\xi}),$$

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

We now establish a first result. Under our assumptions, the individual beliefs about the probability of the good technology, $p_{jt} = \Lambda_{jt}(\theta_H, 0)$, is increasing in the private signal s_j . As a result, the investment decision can be further characterized by a cutoff rule $s^*(p_t, q_t)$ 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, the optimal investment strategy is characterized by a cutoff rule in the private signal $s^*(p_t, q_t)$ that is decreasing in p_t . That is, an agent invests if and only if $s_j \geq s^*(p_t, q_t)$. The expected measure of investing agents is given by*

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

Learning from m_t

To develop intuition on the way agents learn from the measure of investors, 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 the sense of first-order stochastic dominance. The expected measure of investing agents m^e is represented as the mass of agents located to the right of the cutoff s_t^* . We can see that m^e is small in the bad-technology state $(\theta_L, 0)$ (in red), that agents expect more investment in the false-positive state $(\theta_L, \bar{\xi})$ (in green), and that it is at its largest in the good-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 the state θ_L is greater than that of the other states, so information updating will attribute it a higher probability. The two other states, $(\theta_L, \bar{\xi})$ and $(\theta_H, 0)$, have their own higher likelihood region that are also represented on the graph. Importantly for the emergence of boom-and-bust cycles, beliefs about the good state 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.

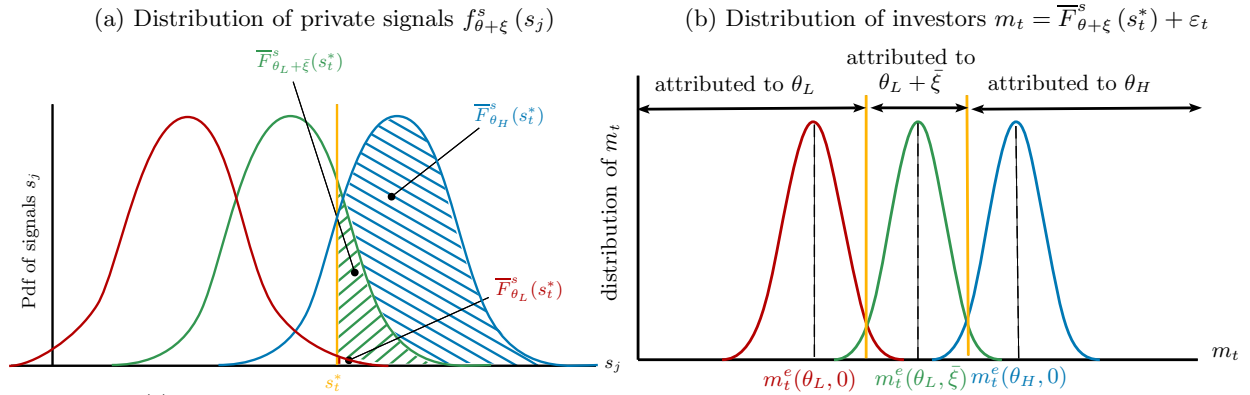



Figure 1: Private beliefs and expected measure of investors

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 history of public signals), that agents end up neglecting 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 mistaken 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 the distributions of signals.

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. These curves are drawn by first connecting a given level of p_t in the bottom-right panel to the equilibrium signal threshold $s^*(p_t, q_t)$ (upper-right panel), itself connected to the upper-left panel which shows how the measures $m^e = \bar{F}_{\theta+\xi}^s(s^*)$ vary with the cutoff s^* . As the bottom panel shows, the expected measure of investing agents m^e is a monotonic transformation of the CDF $F_{\theta+\xi}^s$ in the three different states.



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 $m_t^e = \bar{F}_{\theta+\xi}^s(s_t^*)$, as previously represented in Figure 1 rotated by 90° . The top-right panel displays an example of equilibrium threshold $s^*(p_t, q_t)$ as a function of public belief p_t . The bottom right panel shows how the measure 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 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), almost all (resp. no) agents invest, the three measures converge to $\lim_{s^* \rightarrow -\infty} \bar{F}_{\theta+\xi}^s(s^*) = 1$ (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 s^* varies sufficiently on the support of $\bar{F}_{\theta+\xi}^s$.¹⁴

¹⁴Figure 2 may give the wrong impression that the nonmonotonicity result highly depends on the sigmoidal shape

The model offers a continuous and smooth analog to informational cascades when the equilibrium 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 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 particularly low private signals. Unfortunately, their measure is so small 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, 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

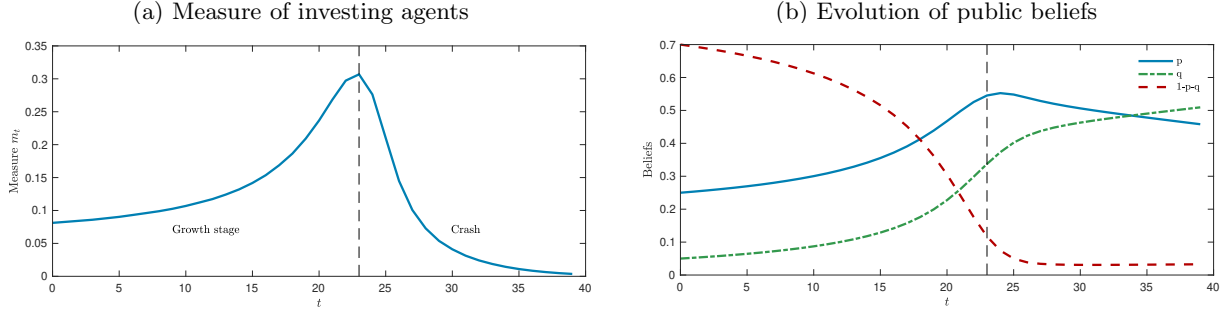
We now present simulations of the model to illustrate its ability to generate endogenous boom-bust patterns out of a single impulse shock. We do not attempt to make a realistic calibration but merely pick parameters so as to highlight the model’s properties. We will examine later under what general conditions one should expect the boom-bust cycles 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 $(\theta, \xi) = (\theta_L, \bar{\xi})$, the case of interest for our purpose.¹⁵

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) $\theta_L + \bar{\xi}$ needs to be sufficiently close to θ_H , so that the two states are hard to distinguish; and (ii) the prior q_0 on the false-positive state $(\theta_L, \bar{\xi})$ needs to be sufficiently small relative to the true positive $(\theta_H, 0)$ for agents to initially attribute most of the rising investment pattern to the true positive state.

of the CDFs $F_{\theta+\xi}^s$. While the regions with higher signal-to-noise ratio may change with the distribution, a robust prediction for any distribution is that the measure m_t is less informative for extreme public beliefs, when agents herd on the same action, since $\bar{F}_{\theta+\xi}^s(s^*) \rightarrow 1$ (resp. $\bar{F}_{\theta+\xi}^s(s^*) \rightarrow 0$) when p_t gets close to 1 (resp. 0) and the cutoff s^* goes to minus infinity (resp. infinity) for any signal distribution.

¹⁵Figures 11 and 12 in the Appendix show the economy’s response to the good-technology and bad-technology states. With our parametrization, these cases are relatively uninteresting: learning is fairly quick, and the dynamics are close to the full information case.



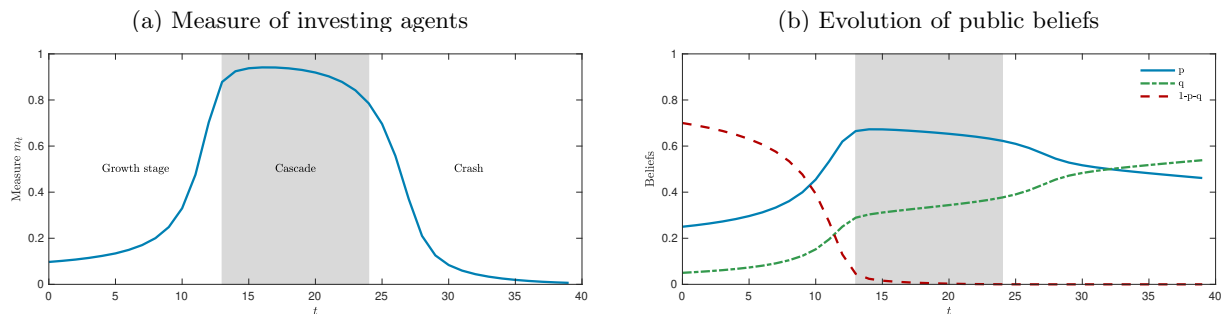
Notes: The simulation was performed with parameters: $\theta_H = 1$, $\theta_L = 0.5$, $\bar{\xi} = 0.4$, $c = 0.80$. The priors are set to $p_0 = 0.25$ and $q_0 = 0.05$. All the distributions are Gaussian: $F_{\theta+\xi}^s \sim \mathcal{N}(\theta + \xi, \sigma_s)$, $F^\varepsilon \sim \mathcal{N}(0, \sigma_\varepsilon)$ 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 $t = 0$, the measure of investing agents (panel 3a) is small (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 bad state and they reduce the probability assigned to it (red curve in panel 3b). Agents also understand that the high investment rate could arise from either the good-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, the observed high level of investment is mostly attributed to the good-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 pattern, in turn, leads to further upward revisions in expectations, seemingly confirming the assessment that the economy is in the good state. We refer to this first stage of the cycle, characterized by the joint increase of investment rates 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, a time comes 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.

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 good state and the false-positive state is high (p_t and q_t are close, so beliefs are more responsive to new information), and ii) the signal m_t is more informative at the peak.¹⁶



Notes: The simulation was performed with parameters: $\theta_H = 1$, $\theta_L = 0.5$, $\bar{\xi} = 0.4$, $c = 0.79$. 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

Information cascades

Whether the growth stage gives way to a sudden collapse or not depends on the parametrization of the model. Figure 4 depicts an example of a cycle in which the economy grows so rapidly at first that it reaches the low-informativeness region associated with a high m_t . In that 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 they herd on investing and m_t becomes uninformative. The economy thus enters a period akin to an information cascade, as described earlier, where almost all agents invest due to overly optimistic public beliefs, and in which markets almost cease to provide 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

¹⁶A 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.

more quickly to their true values and a belief reversal occurs in the later stages.¹⁷

Our model can also generate permanent information cascades, in line with the theoretical models of herding in the spirit of Banerjee (1992) and Bikhchandani et al. (1992). Under the conditions that the support of the private signals s_j is bounded to a set $[\underline{s}, \bar{s}]$ and that the exogenous public signal is uninformative ($\sigma_u \rightarrow \infty$), a situation can arise in which the public beliefs are so optimistic that the investment threshold s^* reaches \underline{s} . In this case, all agents invest regardless of their private signal. As a result, the endogenous public signal m_t does not reveal any of the dispersed information and the observation of the mass of investors is completely uninformative about the true state of the world. The economy is then trapped in a constant state of massive investment even though the true quality of the technology is bad. Similarly, if the public beliefs are sufficiently pessimistic, the economy can remain in a permanently depressed state in which no investment takes place. Appendix A.5 shows how a permanent information cascade can arise in this sample economy.

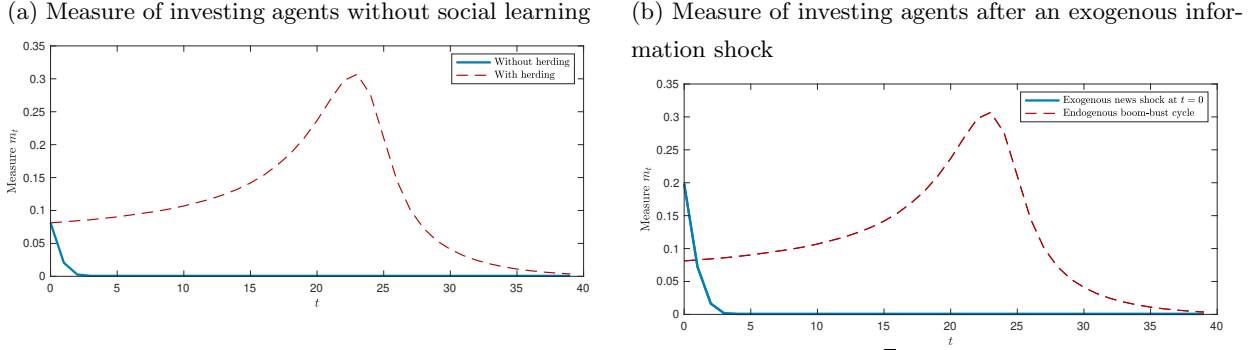
The importance of herding

To better understand the importance of the social learning channel for the dynamics of the economy, we provide a simulation in which that mechanism is turned off ($\sigma_\varepsilon \rightarrow \infty$) so that agents no longer learn by observing m_t and, as a result, no herding takes place. That simulation is presented in panel (a) of Figure 5. As we can see, without the herding mechanism the economy does not initially grow into a boom, despite being in the false positive state and investors having an unusually optimistic prior. In our previous examples, it was the slow diffusion of the information contained in private signals into the public beliefs that led to an economic boom, but since this channel is shut down here, the boom does not happen and the mass of investors m_t quickly converges to zero.

To further highlight how our endogenous learning mechanism differs from more traditional exogenous shocks, we provide another simulation, in panel (b), in which the economy is hit by an exogenous u_t shock immediately before $t = 0$ and the social learning channel remains shut down ($\sigma_\varepsilon \rightarrow \infty$). Upon impact, this shock leads to an increase in m_t , as expected, but the dynamics it triggers is qualitatively different from that generated by the herding mechanism (red dashed curve). As the figure illustrates, investment peaks immediately then gradually fades out as more information is collected. The propagation is weak and does not lead to the feedback loop highlighted in the case of herding, which showed a slow rising pattern of self-reinforcing investment and optimism.¹⁸

¹⁷The way the economy exits the cascade is reminiscent of the “wisdom after the fact” mechanism proposed by Caplin and Leahy (1994) and its reinterpretation in Chapter 4 of Chamley (2004).

¹⁸Only a specific sequence of increasingly positive then negative exogenous shocks u_t can replicate the type of dynamics observed in Figure 3. In contrast, in our model, the boom and bust dynamic is the natural outcome of the economic forces at work. The two models also have very different policy implications, as we discuss in section 3.3.



Notes: Panel (a) The simulation was performed with parameters: $\theta_H = 1$, $\theta_L = 0.5$, $\bar{\xi} = 0.4$, $c = 0.80$. The true state is $(\theta, \xi) = (\theta_L, \bar{\xi})$. The priors are set to $p_0 = 0.25$ and $q_0 = 0.05$. All the distributions are Gaussian: $F_{\theta+\xi}^s \sim \mathcal{N}(\theta + \xi, \sigma_s)$, $F^\varepsilon \sim \mathcal{N}(0, \sigma_\varepsilon)$ and $F^u \sim \mathcal{N}(0, \sigma_u)$ with standard deviations $\sigma_s = 0.5$, $\sigma_\varepsilon = 20$ (uninformative endogenous signal), $\sigma_u = 2.5$. Panel (b) has the same parameters and shows the impact of a shock to u_{-1} so that $p_0 = 0.35$. In both panels the dashed red lines represent m_t from (3).

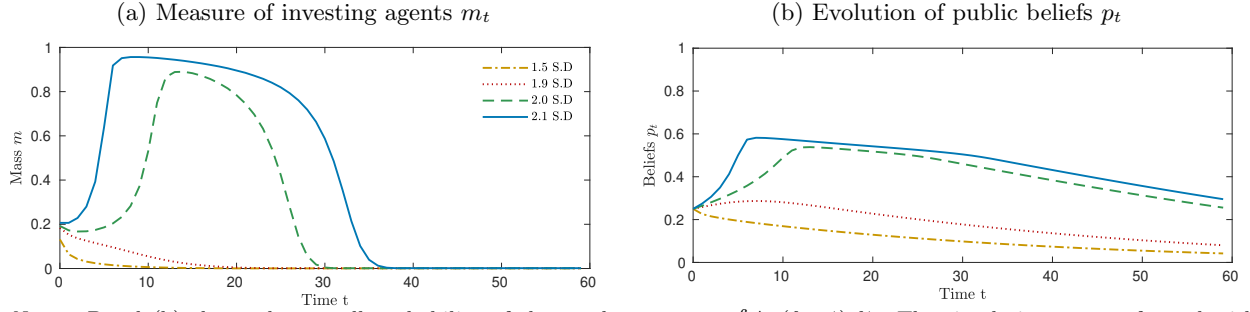
Figure 5: The role of endogenous learning and the impact of exogenous news shocks

3.2 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 on a continuum of values. 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 6 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 bad-technology state $\theta = \theta_L$ but we vary the size of the ξ shock. Four shocks of various sizes are represented, with ξ expressed as a multiple of the standard deviation, namely $\xi = k\sigma_\xi$, $k \in \{1.5, 1.9, 2, 2.1\}$. 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 dash-dotted line), 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 ξ draw 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, $\xi = 2\sigma_\xi$ (green dashed line), the economy begins to experience a boom-bust cycle of the sort 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 $t = 20$, as in Figure 3. When the size of the shock is larger, $\xi > 2\sigma_\xi$ (blue continuous line), 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.



Notes: Panel (b) shows the overall probability of the good state $p_t = \int \Lambda_t(\theta_H, \xi) d\xi$. The simulation was performed with parameters: $\theta_H = 1$, $\theta_L = 0.5$, $c = 0.75$. 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 6: 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 investments to the good-technology state.

We now 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., $F^\xi \sim \mathcal{N}(0, \sigma_\xi^2)$, $F^s | \theta, \xi \sim \mathcal{N}(\theta + \xi, \sigma_s^2)$, $F^\varepsilon \sim \mathcal{N}(0, \sigma_\varepsilon^2)$, $F^u \sim \mathcal{N}(0, \sigma_u^2)$, 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) .*

Note that the above discussion shows a restriction imposed by the theory: because they are rational, agents cannot make systematic mistakes in their assessment of the probability of each state. Hence, boom-bust cycles can only arise for shocks that have a low enough probability of occurring. Our model thus offers a theory of *infrequent* booms-and-busts. Going beyond this limitation may require the introduction of deviations from rationality.

3.3 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 signals that are publicly available and cannot rely on the private

information of the investors when making decisions. We impose these restrictions so that the 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 a new technology.

We follow [Angeletos and Pavan \(2007\)](#) in assuming that the planner seeks to maximize the sum of the investors' expected utility, where the expectation is computed according to the investors' private beliefs. In each period, the planner picks an investment threshold p_t^* such that agents with beliefs $p_{jt} \geq p_t^*$ invest. Written in recursive form, the problem of the social planner is

$$V(\mathcal{I}) = \max_{p^*} E_{\theta, \xi} \left[\int_{p_j \geq p^*} E[\theta - c \mid \mathcal{I}_j] dF_{\theta+\xi}^{p_j}(p_j) \mid \mathcal{I} \right] + \beta E_{\theta, \xi} [V(\mathcal{I}'(p^*)) \mid \mathcal{I}], \quad (13)$$

where \mathcal{I}' is public information next period, which evolves according to the law of motion (12), and where $F_{\theta+\xi}^{p_j}(p_j)$ is the CDF of the agents' subjective probability that $\theta = \theta_H$ when the true state of the world is $\theta + \xi$. The expectation $E_{\theta, \xi}$ is then taking over these states using the public beliefs.

The first term in (13) captures the current-period returns from letting agents with private beliefs above 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 $F_{\theta+\xi}^{p_j}(p_j)$ of private beliefs in that state, which is needed to compute the mass of investors above p^* . The second term in (13), the continuation value, captures the impact of a given investment threshold p^* on the future public information. It is this term that creates a gap between the equilibrium and the efficient allocation. In the competitive equilibrium, individual investors are atomistic. Hence, investors disregard the impact that their individual actions have on the release of public information. The planner, on the other hand, understands that by changing the cutoff p^* , the mass of investors also changes, which affects the informativeness of public signals.

The first-order condition of the planner with respect to p^* can be written as


$$E_{\theta, \xi} \left[(p^* \theta_H + (1 - p^*) \theta_L - c) f_{\theta+\xi}^p(p^*) \mid \mathcal{I} \right] = \beta \frac{\partial E_{\theta, \xi} [V(\mathcal{I}') \mid \mathcal{I}]}{\partial p^*}. \quad (14)$$

The left-hand side of this equation reflects the expected cost of increasing the threshold p^* at the margin. If the true state is $\theta + \xi$, increasing p^* slightly pushes a mass $f_{\theta+\xi}^p(p^*)$ of agents away from investing, each of which loses $p^* \theta_H + (1 - p^*) \theta_L - c$ in expected returns. The planner takes the expectation of these losses over all the states $\theta + \xi$. The right-hand side of the equation reflects the impact of increasing p^* on the flow of public information that is released at the end of the period. By changing p^* , the planner can, for instance, increase the gap between the expected realizations of m in different states of the world. When it does so, m becomes more informative as the signal-to-noise ratio increases. Notice that when $\beta = 0$ the first-order condition (14) collapses to the equilibrium cut-off rule (8), such that the efficient allocation coincide with the equilibrium.

Example: Efficiency in the 3-state model

To better understand the sources of inefficiencies in this environment, we now go back to the example explored in Figure 4 and look at the solution of the planner’s problem. Again, we suppose that the economy is in the false-positive state, and we plot the impulse responses of the mass of investing agents and the public beliefs in the efficient allocation. The results are presented in Figure 7, which shows the impulse responses in the efficient allocation (bold lines) and the equilibrium (thin dashed lines). We see from Panel (a) that in period $t = 0$ the planner is more aggressive than the private agents in pursuing the investment opportunity. The planner behaves in that way because pushing the measure of investors towards intermediate values makes the signal more informative. Indeed, we can see in Panel (b) that the public beliefs move more rapidly in the efficient allocation. In other words, there is an initial phase of “experimentation”. In this particular example, the planner quickly learns that the bad-technology state can be ruled out, as $1 - p - q$ declines sharply in the first few periods.

But one surprising consequence of this faster learning is that the planner is also more rapidly “fooled” into believing that the true state of the technology is $\theta = \theta_H$. Since the public signal m is more informative, and that this signal is consistent with the good technology state, p increases faster than in the competitive equilibrium. As a result, the efficient allocation can still feature boom-bust cycles and information cascades.¹⁹ In turn, the bust tends to happen sooner in the planner’s allocation, as Panel (a) illustrates. Indeed, during the boom, when public beliefs are very optimistic, the planner attempts to curb the amount of investment in comparison to the competitive equilibrium in order to collect more information. 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 the fundamental is $(\theta, \xi) = (\theta_L, \bar{\xi})$, the false-positive state. The simulation was performed with parameters: $\theta_H = 1$, $\theta_L = 0.5$, $\bar{\xi} = 0.4$, $c = 0.79$, $\beta = 0.5$. 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: Endogenous boom-and-bust in the efficient allocation

¹⁹It might seem surprising that the planner rapidly pushes the economy into a boom, even though the true state of the technology is bad, but recall that the planner, like the private agents, does not know the true state of the world and that this early boom would be very beneficial in the good-technology state of the world.

Optimal taxation

To gain further insight into the nature of the model’s inefficiencies, it is useful to look at a particular implementation of the efficient allocation using an investment tax (or subsidy) τ^* that increases the effective cost of investment to $c + \tau^*$. The next proposition characterizes τ^* .

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

$$\tau^* = \left(E_{\theta, \xi} \left[f_{\theta+\xi}^p(p^*) \mid \mathcal{I} \right] \right)^{-1} \beta \frac{\partial E_{\theta, \xi} [V(\mathcal{I}') \mid \mathcal{I}]}{\partial p^*}, \quad (15)$$

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).


We plot in Figure 8 how this optimal tax varies with the public beliefs p in the example of Figure 7.²⁰ We see that the tax tends to be negative for low values of p and positive for larger values. As we described in Section 3.1, when agents are pessimistic (low p), few of them invest and the endogenous public signal does not reveal much information. The planner therefore sets $\tau^* < 0$ to encourage entry and make the observed mass of investors a more precise signal. The opposite happens when many investors invest (high p). In this case, the planner sets $\tau^* > 0$ to discourage investment, once again to make the endogenous public signal more informative. The tax for intermediate values of p reflects these information concerns. Note, however, that the planner phases out the tax completely as p approaches 0 or 1. In these cases, the public beliefs are so extreme that there is very little uncertainty about the true state of the world. Since there is not much more to learn, the planner sets the tax close to zero to minimize the distortion in investment.

As we can see, the tax incentivizes agents to behave against 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.

4 A 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

²⁰To draw this plot, we fix q to some arbitrary value $q = 0.01$ and plot m in the good-technology state.



Notes: Mass of investors and optimal tax as a function of p with the same parameters as 7. These curves are generated by fixing $q = 0.01$. m is computed as $1 - F_{\theta_H}(\hat{s}(p, q))$.

Figure 8: Optimal tax as a function of the mass of investing agents p

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 discuss below. Finally, a more realistic setup is required to explore the quantitative implications of the theory, which we do in the next section.

4.1 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 \(2004\)](#). 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 the 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 future income, 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 [Jaimovich and Rebelo \(2009\)](#) and assume GHH preferences ([Greenwood et al., 1988](#)) to remove the positive income effect on labor supply. 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 prerequisite for agents to benefit from the innovation, and we can observe a joint increase in aggregate consumption and investment along the booming phase of the cycle. To finally ensure that aggregate output may expand in response to this surge in demand, we introduce price stickiness. Under a sufficiently accommodative monetary policy, a muted response of the real interest rate helps sustain the expansion in demand

due to optimistic expectations.

4.2 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

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

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 IT capital 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 is the corresponding real rental rate, w_t is the real wage, Π_t is 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 $1+\pi_t = P_t/P_{t-1}$ is the inflation rate. As usual, the law of motion for each type of capital, $i = T$ or IT , is given by

$$K_{t+1}^i = (1-\delta) K_t^i + I_t^i,$$

where δ is the depreciation rate.

4.3 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_{jt}^o ,

$$Y_{jt}^o = A^o (K_{jt}^o)^\alpha (L_{jt}^o)^{1-\alpha}, \quad 0 \leq \alpha \leq 1.$$

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 becomes available with production function

$$Y_{jt}^n = A_t^n (K_{jt}^n)^\alpha (L_{jt}^n)^{1-\alpha}.$$

The TFP of the new technology A_t^n 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. To further isolate the economy from exogenous productivity fluctuations and focus on belief-driven cycles, we make the assumption that the new technology is initially as productive as the old one, $A_t^n = A^o$, until it matures with some fixed probability $\lambda > 0$ each period. Upon maturation, the true nature of the technology is revealed. Maturation is a one-time event and all uncertainty is resolved afterwards. That is,

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

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

$$K_{jt}^i = \kappa_i \left(\omega_i (K_{it}^{IT})^{\frac{\zeta-1}{\zeta}} + (1 - \omega_i) (K_{it}^T)^{\frac{\zeta-1}{\zeta}} \right)^{\frac{\zeta}{\zeta-1}}, \quad \zeta > 0, \quad (16)$$

where $\kappa_i = \left(\omega_i^\zeta + (1 - \omega_i)^\zeta \right)^{-\frac{1}{\zeta-1}}$ and with the assumption that the intensity in IT capital is greater for the new than for the old technology, $0 \leq \omega_o < \omega_n \leq 1$.²¹ We denote by z_t^i , $i = o, n$, the rental price of each bundle.

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 a CDF F^ε with support $[0, 1]$. The remaining $1 - \mu$ entrepreneurs are rational and choose the best of the two technologies, based on public and

²¹The value of the parameter κ_i is set such that permanent changes in the measure of new technology adopters m_t have no effect on steady-state output for equal productivities.

private information. There is no cost of switching, so the decision for entrepreneur j is static:

$$i_{jt} = \operatorname{argmax}_{i_{jt} \in \{0,1\}} (1 - i_{jt}) E [\Pi_t^o | \mathcal{I}_{jt}] + i_{jt} E [\Pi_t^n | \mathcal{I}_{jt}],$$

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

Wholesale sector

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

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

$$Y_t^w = \left(\int_0^1 Y_{jt}^{\frac{\sigma-1}{\sigma}} dj \right)^{\frac{\sigma}{\sigma-1}}, \quad \sigma \geq 0, \quad (17)$$

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,

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

where Y_{jt}^r is the quantity purchased from each retail firm j and σ is the same elasticity of substitution as in (17).

4.4 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,

$$\frac{1 + R_t}{1 + \bar{R}} = \left(\frac{1 + \pi_t}{1 + \bar{\pi}} \right)^{\phi_\pi} \left(\frac{Y_t}{\bar{Y}} \right)^{\phi_y}, \quad (18)$$

where \bar{R} , $\bar{\pi}$ and \bar{Y} correspond to the values of, respectively, the target nominal interest rate, inflation and output, which we define later.

4.5 Information

The information structure for entrepreneurs mimics the one in the simple model of Section 2. The true technology 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 becomes common knowledge, so the economy operates under full information from then on. 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. As in the earlier model, entrepreneurs draw their signals from the CDF $F_{\theta+\xi}^s$ where $\xi \sim F^\xi$ is the common noise term and the family $\{F_x^s\}_{x \in I}$ 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 with the investment return in the simplified model, agents learn by observing an exogenous public signal $s_t^u = \theta + u_t$, centered around θ with noise distributed according to CDF F^u and standard deviation σ_u . Social learning takes place through the observation of endogenous market activity. In particular, we assume that agents observe the measure of entrepreneurs that adopt the new technology:

$$m_t = \int_0^1 i_{jt} dj = \underbrace{\int_0^{1-\mu} i_{jt} dj}_{\text{rational entrepreneurs}} + \underbrace{\mu \varepsilon_t}_{\text{noise entrepreneurs}}, \quad \text{with } \varepsilon_t \sim \text{iid CDF } F^\varepsilon. \quad (19)$$

The measure of the new-technology adopters m_t in (19) is almost identical to the measure of investors in (9) 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.²² Apart from that point, the informational content of m_t is identical to the earlier model up to some rescaling.

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 the public information up to time t . As a result, prices and quantities provide no other information than is already contained in m_t .

Beliefs

As in the simple model, we denote by $\mathcal{I}_t = \{m_{t-1}, s_{t-1}^u, \dots, m_0, s_0^u\}$ the public information available to non-entrepreneur agents (households, monetary authority, retailers and outside observers). Public beliefs are captured by the joint distribution

$$\Lambda_t(\tilde{\theta}, \tilde{\xi}) = Pr(\theta = \tilde{\theta}, \xi \in [\tilde{\xi}, \tilde{\xi} + d\tilde{\xi}] | \mathcal{I}_t).$$

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

4.6 Timing

Before date 0, the economy is in a deterministic, no-inflation initial steady state that corresponds to the economy before the introduction of the new 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 based on the capital stocks (K_t^{IT}, K_t^T) and their information set (Λ_t, s_j) (Stage A),
2. The measure of technology adopters m_t is realized,
3. The new technology matures with probability λ ,
4. Simultaneously (Stage B),
 - (a) All agents observe $\{m_t, s_t^u\}$ and update their information,
 - (b) The household chooses consumption, investment and labor supply,
 - (c) Production takes place,

²²In this general 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 information 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.

- (d) The monetary authority sets the policy rate,
- (e) Markets clear.

The notation “Stage A” and “Stage B” is used in Appendix B.1 to identify when decisions are made in the full equilibrium definition.

4.7 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 $i = o, n$:

$$\Pi_t^i = (P_{jt} - P_t mc_t^i) Y_{jt}. \quad (20)$$

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

$$E \left[\hat{A}_t^n - \alpha \hat{z}_t^n | \mathcal{I}_{jt} \right] \geq E \left[\underbrace{\hat{A}_t^o}_{=0} - \alpha \hat{z}_t^o | \mathcal{I}_{jt} \right], \quad (21)$$

where the hatted variables are log-deviations from a steady state that we define in the next section and z_t^i , $i = n, o$, are the rental rates on the capital bundles (16). As equation (21) 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.

4.8 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 $\Omega_t = (K_t^{IT}, K_t^T, \Lambda_t)$ and the private signals s_j . The measure of new technology adopters is given by

$$m_t = (1 - \mu) m^e(\Omega_t, \theta, \xi) + \mu \varepsilon_t \quad (22)$$

$$\text{with } m^e(\Omega_t, \theta, \xi) = \int \mathbb{1}(i_j(\Omega_t; s_j) = 1) f_{\theta+\xi}^s(s_j) ds_j. \quad (23)$$

In turn, the belief updating equation (12) needs to be amended in the following way

$$\Lambda_{t+1}(\tilde{\theta}, \tilde{\xi}) = \frac{\Lambda_{t|s_t^u}(\tilde{\theta}, \tilde{\xi}) f^\varepsilon\left(\frac{1}{\mu}\left(m_t - (1 - \mu) m^e(\Omega_t, \tilde{\theta}, \tilde{\xi})\right)\right)}{\int \Lambda_{t|s_t^u}(\theta, \xi) f^\varepsilon\left(\frac{1}{\mu}\left(m_t - (1 - \mu) m^e(\Omega_t, \theta, \xi)\right)\right) d(\theta, \xi)}, \quad (24)$$

where $\Lambda_{t|s_t^u}$ are the posterior beliefs after observing the public signal s_t^u .

This concludes the exposition of the full quantitative model. A complete equilibrium definition is provided in Appendix B.1 along with the full set of equations in Appendix B.2.

5 Quantitative Exercise

We now 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.

5.1 Resolution Method

Our model can only be solved numerically. We follow a strategy common in the information friction literature which consists in linearizing 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 precedes 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 $\Lambda_t(\theta, \xi)$. 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, $E_t[\Lambda_{t+1}(\tilde{\theta}, \tilde{\xi}) | \mathcal{I}_t] = \Lambda_t(\tilde{\theta}, \tilde{\xi})$. The martingale property implies that any equilibrium condition of the form $E_t[f(x_t, x_{t+1}, \Lambda_t, \Lambda_{t+1})] = 0$, where f is a nonlinear function and x_t a vector of model variables can be approximated to a first-order as

$$E[f(x_t, x_{t+1}, \Lambda_t, \Lambda_{t+1}) | \mathcal{I}_t] \simeq E[f(x_t, x_{t+1}, \Lambda_t, \Lambda_t) | \mathcal{I}_t].$$

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 technology adoption decision and the evolution of beliefs in a nonlinear way, then repeat in the

next period under the new beliefs. Section B.3 in the appendix provides additional details about our resolution method.

5.2 Calibration

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 1995Q4 to the crash in 2001Q1.

Our goal with this exercise is not to show that the mechanism can precisely replicate the behavior of the economy during the dot-com period. Rather, we want to determine whether a reasonable calibration of the model is able to generate boom-bust cycles that are similar in terms of magnitude and comovements to what we see in the data. Since we focus on a single historical episode, there is limited data to pin down certain parameters with confidence. In these cases, we rely on the best data available and provide extensive robustness tests in the appendix to show that the mechanism does not hinge on specific parameter values. We are, however, careful to properly discipline moments that are key for the mechanism such as the dispersion of private beliefs, as we explain in more details below.

Parameter	Value	Target
α	0.36	Labor share
β	0.99	4% annual interest rate
ψ	2	Frisch elasticity of labor supply (Chetty et al., 2011)
χ	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)
ζ	1.71	Elasticity between types of capital (Boddy and Gort, 1971)

Table 1: Standard parameters

The model is solved at a quarterly frequency. Table 1 lists a first set of standard parameters that we take from the literature. The 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 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 χ is set to yield a standard average price duration of 1 year (Basu and Bundick, 2017). The elasticity of substitution between varieties σ is set to 10 to match an average markup of 11%. The Taylor rule parameters (ϕ_y, ϕ_π) are within the estimates of Clarida et al. (2000). Finally, we pick the elasticity ζ between the different types of capital within the firm from early estimates by Boddy and Gort (1971).²³

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 a share of IT capital of 3% before the introduction of the new technology in 1991 and 14% in 2007 (Strauss and Samkharadze, 2011). The probability of maturation for new technologies λ is set to $1/22$ to match an average waiting time of 22 quarters, corresponding to length of our period of interest 1995Q4-2001Q1. 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 to discipline θ_H and θ_L . 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 $\theta_H = 1.099$. Similarly, targeting the lowest growth forecast of 0.80% in 2001Q1, we obtain an estimate of $\theta_L = 0.912$.²⁴

For the common noise, we adopt the same structure as in Section 3.1 and assume that that the pair (θ, ξ) can only take the three values

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

where $(\theta_L, 0)$ is the *bad-technology* state, $(\theta_H, 0)$ is the *good-technology* state and $(\theta_L, \bar{\xi})$ is the *false-positive* state. As before, we let p_t and q_t denote the public beliefs about the good-technology and false-positive states. 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 1995Q4-2001Q1, which yields $\sigma_s = 0.156$. 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 with parameters (2,2).²⁵

Five 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

²³We find that ζ does not have a large impact on our simulations so that, for instance, a more neutral Cobb-Douglas parametrization ($\zeta = 1$) leads to similar impulse responses.

²⁴See Appendix B.4 for more details about how we calibrate θ_L and θ_H .

²⁵The distribution Beta(1, 1) is uniform and produces a flat learning response. As a result, we pick a Beta(2, 2) distribution which is symmetric around its mode of 0.5. This assumption is relatively unimportant, while the important parameter is μ , which governs the informativeness of the endogenous public signal.

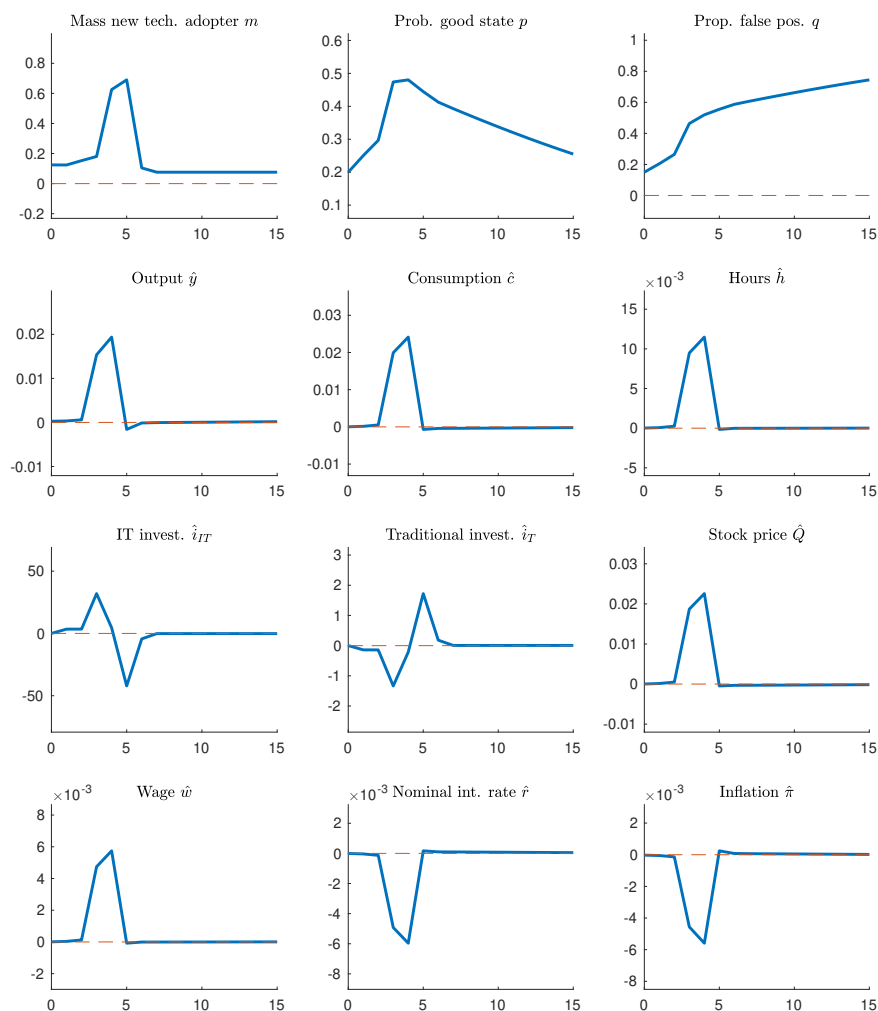
Parameter	Value	Target
ω_o	0.11	Share of IT capital 1991 (3%)
ω_n	0.26	Share of IT capital 2007 (14%)
λ	1/22	Duration of NASDAQ boom-bust 1995Q4-2001Q1
θ_H	1.099	SPF highest growth forecast over 1998-2001
θ_L	0.912	SPF lowest growth forecast over 1998-2001
s_j	$\mathcal{N}(\theta + \xi, 0.156)$	SPF avg. dispersion in forecasts over 1998-2001
ε	Beta(2, 2)	Non-uniform distribution over [0, 1]
μ	15%	Fraction of noise traders
p_0	0.20	Prior on the “good technology” state
q_0	0.15	Prior on the “false positive” state
$\bar{\xi}$	$0.95(\theta_h - \theta_l)$	See text
σ^u	$3 \times \sigma^s$	See text

Table 2: Dot-Com episode related parameters

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 $\mu = 15\%$. We must also specify the priors p_0 and q_0 that agents associate with the states of the world “good technology” and “false positive”. There is no obvious moment that we can target for these parameters given that we focus on one historical episode. We start with a benchmark parametrization that assigns relatively low values to those parameters and do some robustness analysis in the appendix. Our benchmark parametrization attempts to capture the idea that if new technologies can be invented frequently, only few of them lead to deep economic transformations like the ones considered in this paper. This suggests a small value for p_0 . Similarly, the false-positive signal ξ cannot happen too frequently, otherwise agents would distrust their private signals and the boom-bust cycles would never arise. As a benchmark, we therefore pick $p_0 = 0.2$ and $q_0 = 0.15$ but we show in [Appendix B](#) that boom-bust episodes are robust to some variation in these values. We must also set a value for the common noise term $\bar{\xi}$. Again, neither the literature nor the data provides much guidance. Since one of our goals is to evaluate the potential of the model in generating boom-bust cycles, we pick a relatively large shock and set $\bar{\xi} = 0.95(\theta_H - \theta_L)$. Finally, we need a value for the exogenous public signal about θ . For the mechanism to operate, we need that signal to be mostly uninformative, otherwise agents know the true value of θ and boom-bust cycles can obviously never arise. We set $\sigma_u = 3 \times \sigma_s$ in the benchmark simulations so that a private signal has about the same information as three exogenous public signals. [Appendix B](#) conducts sensitivity analysis over all the parameters mentioned in this paragraph. The results are fairly robust.

5.3 Boom-and-Bust Cycles

Figure 9 presents the impulse responses of the economy to a false-positive shock.²⁶ In period $t = 0$, a new technology is discovered and entrepreneurs receive encouraging private signals about the true value of that technology. Since the false-positive state is initially deemed unlikely, entrepreneurs begin to adopt the new technology and the economy goes through a growth phase with m_t moving upward. While this growth in m_t is initially consistent with both the “good technology” and the “false-positive” state, there comes a point at which agents start to realize that the data is more consistent with $\theta = \theta_L$, and the likelihood p_t starts to decline. As result, the mass of entrepreneurs who adopt the new technology collapses around $t = 5$, pushing the economy into a crash.



Notes: The impulse responses are reported in log-deviations from the initial non-stochastic steady state. All shocks after the realization of (θ, ξ) are set to 0 ($\varepsilon_t = 0$ and the technology never matures exogenously).

Figure 9: Impulse response in the false-positive state of the world

²⁶We set all shocks after the realization of (θ, ξ) to 0, i.e., $\varepsilon_t = 0$ and $u_t = 0$ for all t . Similarly, the technology never matures through the λ shock and the bust we observe is purely endogenous.

While the behavior of the mass of adopters m_t is similar to what we have observed in the simplified model of Section 3.1, Figure 9 shows how this pattern 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. Price rigidities play an important role in turning this surge in demand into a general macroeconomic boom. 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 remains high. Firms, satisfying demand, respond by raising output and employment. Because of a higher labor demand, wages increase, but inflation remains low because firms anticipate greater productivity and lower marginal costs in the future. As evidenced by the variable \hat{Q} , which captures the value of the firms, the economy also experiences 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 abandon the new technology, a recession occurs, with GDP falling below trend, because agents wake up after having invested too much in IT capital and not enough in traditional capital. This misallocation of resources creates a negative income effect which puts downward pressure on aggregate demand, despite a recovery in traditional investment.

A few comments are in order at this point. First, while the model is able to generate a recession with a significant peak-to-trough gap (about 2%), 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 most of the news/noise-driven business cycle literature. Second, while the model is able to generate a recession with output falling below trend, that effect is somewhat weak in our simulation. Given that the capital stocks can be adjusted rapidly, the misallocation channel responsible for the dip does not lead to a large drop in output. The introduction of debt and bankruptcy in the model would provide another channel through which the crash could result in a deeper recession. Third, we can 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,²⁷ our benchmark calibration suggests that boom-bust cycles may arise in our calibrated model at the fairly high frequency of $q_0 = 15\%$ after the introduction of a new technology. We view this number as quite

²⁷Chamley (2004) suggests that a boom-bust cycles arises with a probability of 10^{-6} in the traditional model of herding of Avery and Zemsky (1998).

encouraging for the ability of rational herding models in explaining the data.

Overall, the impulse responses of Figure 9 suggest that it is possible to generate a realistic macroeconomic boom-bust cycle that is entirely driven by the internal forces of the model—without the need of an exogenous shock to trigger the bust after the initial introduction of a new technology.

5.4 Monetary Policy Implications

We conclude this quantitative exercise with a discussion of the role of monetary policy during boom-bust cycles. It has long been the subject of policy debates whether monetary authorities should narrowly focus on output and inflation stabilization or, instead, intervene in response to bubble-like phenomena in investment or in the stock market.

As we explained in the context of the simplified learning model, our theory justifies the use of interventions that *lean against the wind* because of the presence of an informational externality. While evaluating optimal monetary policy in our setup is beyond the scope of this paper, we ask whether a monetary policy rule that leans against the wind would be desirable in our economy.²⁸ The answer is mostly negative and we illustrate our argument with a simple example.

We augment our Taylor rule with an explicit mandate to increase interest rates when investment in the new technology I_t^{IT} rises:²⁹

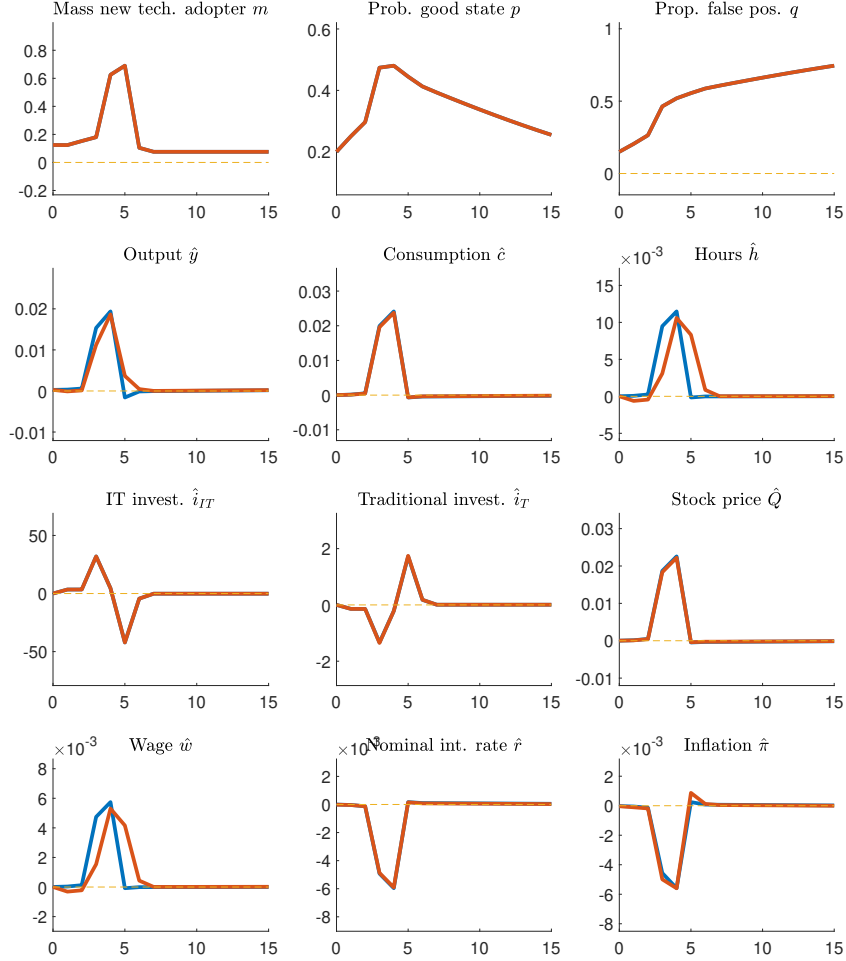
$$\frac{1+r_t}{1+\bar{r}} = \left(\frac{1+\pi_t}{1+\bar{\pi}}\right)^{\phi_\pi} \left(\frac{Y_t}{\bar{Y}}\right)^{\phi_y} \left(\frac{I_t^{IT}}{\bar{I}^{IT}}\right)^{\phi_i}, \quad \phi_i \geq 0. \quad (25)$$

We compute ex-ante welfare using Monte-Carlo simulations with all the model shocks and compute the coefficient ϕ_i that maximizes welfare. Note that the model is solved around the efficient steady state corrected for monopolistic distortions. Any welfare gain thus arises mainly because of the information externality and the suboptimality of the initial monetary policy. Figure 10 reports the impulse responses to a false-positive shock with the benchmark case in blue and the welfare-maximizing leaning-against-the-wind Taylor rule in red.

As the figure illustrates, the differences manifest themselves mostly in output, hours and wages, and to a lesser extent in the inflation rate. The other series are nearly identical. The augmented policy (25) with $\phi_i > 0$ manages to reduce the volatility of output in response to the shock: the

²⁸Computing the optimal monetary policy or the optimal investment tax would require to solve the model using an entirely different, less tractable method. Since the nonlinearities in learning are essential for the result on leaning-against-the-wind policies, the resolution method cannot rely on linearization and the state space reduction result described in subsection 5.1 no longer holds.

²⁹We also experimented with a Taylor rule of the form $\frac{1+r_t}{1+\bar{r}} = \left(\frac{1+\pi_t}{1+\bar{\pi}}\right)^{\phi_\pi} \left(\frac{Y_t}{\bar{Y}}\right)^{\phi_y} \left(\frac{m_t}{\bar{m}}\right)^{\phi_m}$, $\phi_m \geq 0$, but found no parameter $\phi_m > 0$ that increased welfare, due to the fact that m_t lags the cycle by one period as it tracks the dynamics of IT capital.



Notes: The impulse responses are reported in log-deviations from the initial steady state. The blue curve represents the response of the economy in our benchmark calibration. The red curve corresponds to the Taylor rule described in (25) with the coefficient $\phi_i = 0,3.84.10^{-5}$ that maximize ex-ante welfare.

Figure 10: Impact of leaning-against-the-wind monetary policy

economy expands slightly less during the growth stage of the cycle and the recession that would normally follow is averted. Hours are similarly less volatile. The optimal coefficient $\phi_i = 3.84 \times 10^{-5}$ is fairly small because this policy comes at the cost of slowing booms driven by good technologies, which the monetary authority is reluctant to do because the probability of a false-positive episode remains low. Accordingly, the welfare gain is small: +0.002% in consumption equivalent relative to the benchmark calibration. Our results thus provide little support for this type of interventions.

Beyond these considerations, another reason why a monetary policy that leans against the wind is not more effective is that it is an inappropriate tool to address the information externality highlighted in subsection (3.3). We know from that section that the information externality shows up as a distortion in s^* and m_t , which govern the amount of information released by the social learning channel. In the current context, we see in the first panel of Figure 10 that monetary policy

has virtually no effect on m_t . The reason is that the technology adoption decision (21) is entirely driven by beliefs about θ and the relative rental rates of capital z_t^i , $i = o, n$. Since the monetary authority cannot affect beliefs directly (they have no additional information to provide to investors), the only channel by which it could operate is through the rental rates differential $\hat{z}_t^n - \hat{z}_t^o$. We find this channel quantitatively weak: by raising interest rates, the household's incentives to accumulate both types of capitals are affected, but in similar ways: the cost of the two capital bundles change by an almost identical amount and technology adoption is left nearly unchanged. We thus find little scope for the introduction of leaning-against-the-wind motives in monetary policy in a context like ours and conclude that other policy instruments with a direct impact on technology adoption, like the investment taxes and subsidies explored in subsection 3.3, are more promising.

6 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 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, but the 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, etc).

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

- Adam, K., A. Marcet, and J. Beutel (2017). Stock price booms and expected capital gains. *American Economic Review* 107(8), 2352–2408.

- 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.
- Barlevy, G. (2014). A leverage-based model of speculative bubbles. *Journal of Economic Theory* 153, 459–505.
- 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.
- Basu, S. and B. Bundick (2017). Uncertainty shocks in a model of effective demand. *Econometrica* 85(3), 937–958.
- 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 (2014, dec). 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.
- Benhabib, J. and R. E. Farmer (1994). Indeterminacy and increasing returns. *Journal of Economic Theory* 63(1), 19–41.
- Benhabib, J., S. Schmitt-Grohé, and M. Uribe (2002). Chaotic interest-rate rules. *American Economic Review* 92(2), 72–78.
- Benhabib, J., P. Wang, and Y. Wen (2015). Sentiments and aggregate demand fluctuations. *Econometrica* 83(2), 549–585.
- 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.
- Blasco, N., P. Corredor, and S. Ferreruella (2012). Market sentiment: a key factor of investors' imitative behaviour. *Accounting & Finance* 52(3), 663–689.

- 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.
- Bordalo, P., N. Gennaioli, A. Shleifer, and S. J. Terry (2021). Real credit cycles. Technical report, National Bureau of Economic Research.
- Burnside, C., M. Eichenbaum, and S. Rebelo (2016). Understanding booms and busts in housing markets. *Journal of Political Economy* 124(4), 1088–1147.
- Caplin, A. and J. Leahy (1994). Business as usual, market crashes, and wisdom after the fact. *The American economic review*, 548–565.
- Carceles-Poveda, E. and C. Giannitsarou (2008). Asset pricing with adaptive learning. *Review of Economic dynamics* 11(3), 629–651.
- 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., C. L. Ilut, R. Motto, and M. Rostagno (2008). Monetary policy and stock market boom-bust cycles.
- 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.
- Eeckhout, J. and I. Lindenlaub (2019). Unemployment cycles. *American Economic Journal: Macroeconomics* 11(4), 175–234.
- Eusepi, S. and B. Preston (2011). Expectations, learning, and business cycle fluctuations. *American Economic Review* 101(6), 2844–72.
- Fajgelbaum, P. D., E. Schaal, and M. Taschereau-Dumouchel (2017). Uncertainty traps. *The Quarterly Journal of Economics* 132(4), 1641–1692.
- Galariotis, E. C., W. Rong, and S. I. Spyrou (2015). Herding on fundamental information: A comparative study. *Journal of Banking & Finance* 50, 589–598.
- Galí, J. (2014). Monetary policy and rational asset price bubbles. *American Economic Review* 104(3), 721–52.
- Golosov, M. and G. Menzio (2020). Agency business cycles. *Theoretical Economics* 15(1), 123–158.
- Grandmont, J.-M. (1985). On endogenous competitive business cycles. *Econometrica: Journal of the Econometric Society*, 995–1045.

- Greenwood, J., Z. Hercowitz, and G. W. Huffman (1988). Investment, capacity utilization, and the real business cycle. *The American Economic Review*, 402–417.
- Guerrón-Quintana, P., T. Hirano, and R. Jinnai (2020). Recurrent bubbles and economic growth. *Available at SSRN 3350097*.
- Guesnerie, R. and M. Woodford (1992). Endogenous fluctuations. *Advances in economic theory* 2, 289–412.
- Hirano, T. and N. Yanagawa (2016). Asset bubbles, endogenous growth, and financial frictions. *The Review of Economic Studies* 84(1), 406–443.
- Jaimovich, N. and S. Rebelo (2009). Can news about the future drive the business cycle? *American Economic Review* 99(4), 1097–1118.
- Kaplan, G. and G. Menzio (2016). Shopping externalities and self-fulfilling unemployment fluctuations. *Journal of Political Economy* 124(3), 771–825.
- Kocherlakota, N. R. (1992). Bubbles and constraints on debt accumulation. *Journal of Economic Theory* 57(1), 245–256.
- 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.
- Matsuyama, K. (1999). Growing through cycles. *Econometrica* 67(2), 335–347.
- Matsuyama, K. (2013). The good, the bad, and the ugly: An inquiry into the causes and nature of credit cycles. *Theoretical Economics* 8(3), 623–651.
- Miao, J. and P. Wang (2012). Bubbles and total factor productivity. *American Economic Review* 102(3), 82–87.
- Pigou, A. (1927). Industrial fluctuations.
- Straub, L. and R. Ulbricht (2017). Endogenous uncertainty and credit crunches. *Available at SSRN 2668078*.
- Straub, L. and R. Ulbricht (2019). Endogenous second moments: A unified approach to fluctuations in risk, dispersion, and uncertainty. *Journal of Economic Theory* 183, 625–660.
- Strauss, H. and B. Samkharadze (2011). Ict capital and productivity growth. *EIB papers* 16(2), 8–28.
- 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.
- Wen, Y. (1998). Capacity utilization under increasing returns to scale. *Journal of Economic theory* 81(1), 7–36.
- 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 (5) that builds private beliefs from the public ones becomes

$$\begin{aligned} p_{jt} &= p_j(p_t, q_t, s_j) = \frac{p_t f_{\theta_H}^s(s_j)}{p_t f_{\theta_H}^s(s_j) + q_t f_{\theta_L + \bar{\xi}}^s(s_j) + (1 - p_t - q_t) f_{\theta_L}^s(s_j)}, \\ q_{jt} &= q_j(p_t, q_t, s_j) = \frac{q_t f_{\theta_L + \bar{\xi}}^s(s_j)}{p_t f_{\theta_H}^s(s_j) + q_t f_{\theta_L + \bar{\xi}}^s(s_j) + (1 - p_t - q_t) f_{\theta_L}^s(s_j)}. \end{aligned} \quad (26)$$

Equation (11) that defines the interim beliefs after observing R_t is simply

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

Finally, in the three state model, the optimal investment strategy characterized by Equation (12) that defines the law of motion of beliefs after observing m_t becomes

$$\begin{aligned} p_{t+1} &= \frac{p_{t|R_t} f^\varepsilon(m_t - \bar{F}_{\theta_H}^s(s^*(p_t, q_t)))}{p_{t|R_t} f^\varepsilon(m_t - \bar{F}_{\theta_H}^s(s^*(p_t, q_t))) + q_{t|R_t} f^\varepsilon(m_t - \bar{F}_{\theta_L + \bar{\xi}}^s(s^*(p_t, q_t))) + (1 - p_{t|R_t} - q_{t|R_t}) f^\varepsilon(m_t - \bar{F}_{\theta_L}^s(s^*(p_t, q_t)))}, \\ q_{t+1} &= \frac{q_{t|R_t} f^\varepsilon(m_t - m_e(p_{t|R_t}, q_{t|R_t}, \theta_L, \bar{\xi}))}{p_{t|R_t} f^\varepsilon(m_t - \bar{F}_{\theta_H}^s(s^*(p_t, q_t))) + q_{t|R_t} f^\varepsilon(m_t - \bar{F}_{\theta_L + \bar{\xi}}^s(s^*(p_t, q_t))) + (1 - p_{t|R_t} - q_{t|R_t}) f^\varepsilon(m_t - \bar{F}_{\theta_L}^s(s^*(p_t, q_t)))}. \end{aligned}$$

A.2 Propositions

Proposition 1. *There exists a unique equilibrium.*

Proof. The threshold p^* is uniquely determined by (8). The result is established recursively. Fix the fundamental (θ, ξ) and the realization of the shocks $\{u_0, \varepsilon_0, u_1, \varepsilon_1, \dots\}$. At any date t , given public beliefs Λ_t , (5) and (7) yield a unique distribution of private beliefs $\{\Lambda_{jt}\}_{j \in [0,1]}$ and $\{p_{jt}\}_{j \in [0,1]}$. Given these, under the tie-breaking rule that indifferent agents invest, there is a unique m_t^e derived from (10) and, therefore a unique m_t from (9). As a result, updating beliefs through (11) and (12) yields unique $\Lambda_{t|R}$ and Λ_{t+1} . We have shown that the updating of public beliefs yields a unique Λ_{t+1} from Λ_t and the realization of shocks $\{u_t, \varepsilon_t\}$. Starting from public beliefs Λ_0 , there is therefore a unique equilibrium path $\{\Lambda_0, \Lambda_1, \dots\}$ for any history of shocks, and all other quantities can be uniquely determined from it. \square

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

$$m^e(p_t, q_t, \theta, \xi) = \bar{F}_{\theta + \bar{\xi}}^s(s^*(p_t, q_t)).$$

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

$$p_j(p_t, q_t, s_j) = \frac{p_t}{p_t + q_t \frac{f_{\theta_L + \bar{\xi}}^s(s_j)}{f_{\theta_H}^s(s_j)} + (1 - p_t - q_t) \frac{f_{\theta_L}^s(s_j)}{f_{\theta_H}^s(s_j)}}.$$

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

$$m^e(p_t, q_t, \theta, \xi) = \int \mathbf{1}(p_j(p_t, q_t, s_j) \geq p^*) f_{\theta + \xi}^s(s_j) ds_j = \bar{F}_{\theta + \xi}^s(s^*(p_t, q_t)).$$

□

Proposition 2. *In the Gaussian case, i.e., $F^\xi \sim \mathcal{N}(0, \sigma_\xi^2)$, $F^s | \theta, \xi \sim \mathcal{N}(\theta + \xi, \sigma_s^2)$, $F^\varepsilon \sim \mathcal{N}(0, \sigma_\varepsilon^2)$, $F^u \sim \mathcal{N}(0, \sigma_u^2)$, 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 good state in date 1, p_1 , increases after observing m_0 . Since beliefs must converge to the truth in the long-run ($p_t \rightarrow 0$), due to the strictly positive flow of information, and the law of large numbers, this guarantees the existence of a boom-and-bust cycle in beliefs. 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 s^* . The probability that individual j puts on the good state is given by

$$p_j(p_0, s_j) = \frac{\int \Lambda_0(\theta_H, \xi) f_{\theta_H + \xi}^s(s_j) d\xi}{\int \Lambda_0(\theta_H, \xi) f_{\theta_H + \xi}^s(s_j) d\xi + \int \Lambda_0(\theta_L, \xi) f_{\theta_L + \xi}^s(s_j) d\xi}.$$

Since ξ is independent from θ , $\Lambda_0(\theta_H, \xi) = p_0 f^\xi(\xi)$ and $\Lambda_0(\theta_L, \xi) = (1 - p_0) f^\xi(\xi)$. Notice, then, that $\int f^\xi(\xi) f_{\theta + \xi}^s(s_j) d\xi$ is the pdf of s_j given θ , which is a normal, $s_j | \theta \sim \mathcal{N}(\theta, \sigma_\xi^2 + \sigma_s^2)$. Denote ϕ the pdf of a unit normal, we have:

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

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 good state at the beginning of period 1, p_1 , is given by

$$p_1 = \int \Lambda_1(\theta_H, \xi) d\xi = \frac{\int \Lambda_0(\theta_H, \xi) f^\varepsilon(m_0 - m^e(\Lambda_0, \theta_H, \xi)) d\xi}{\int \Lambda_0(\theta_H, \xi) f^\varepsilon(m_0 - m^e(\Lambda_0, \theta_H, \xi)) d\xi + \int \Lambda_0(\theta_L, \xi) f^\varepsilon(m_0 - m^e(\Lambda_0, \theta_L, \xi)) d\xi}.$$

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

$$p_1 = \frac{1}{1 + \frac{1-p_0}{p_0} \frac{\int f^\xi(\xi) f^\varepsilon(m_0 - \bar{F}_{\theta_L + \xi}^s(s_0^*)) d\xi}{\int f^\xi(\xi) f^\varepsilon(m_0 - \bar{F}_{\theta_H + \xi}^s(s_0^*)) d\xi}}.$$

Denoting ξ_0 the true shock, the impulse response in m_t yields $m_0 = \bar{F}_{\theta_L + \xi_0}^s(s_0^*)$, which goes to 1 as $\xi_0 \rightarrow \infty$. Because the MLRP property implies first-order stochastic dominance, we have $\bar{F}_{\theta_L + \xi}^s(s_0^*) < \bar{F}_{\theta_H + \xi}^s(s_0^*)$. Since $f^\varepsilon(\varepsilon)$

is decreasing for $\varepsilon \geq 0$, we have

$$f^\varepsilon (m_0 - \bar{F}_{\theta_L + \xi}^s (s_0^*)) < f^\varepsilon (m_0 - \bar{F}_{\theta_H + \xi}^s (s_0^*))$$

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

$$\begin{aligned} & \int f^\xi (\xi) f^\varepsilon (m_0 - \bar{F}_{\theta_H + \xi}^s (s_0^*)) d\xi - \int f^\xi (\xi) f^\varepsilon (m_0 - \bar{F}_{\theta_L + \xi}^s (s_0^*)) d\xi \\ & \xrightarrow{\xi_0 \rightarrow \infty} \int_{-\infty}^{\infty} f^\xi (\xi) [f^\varepsilon (1 - \bar{F}_{\theta_H + \xi}^s (s_0^*)) - f^\varepsilon (1 - \bar{F}_{\theta_L + \xi}^s (s_0^*))] d\xi > 0 \end{aligned}$$

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

$$\int f^\xi (\xi) f^\varepsilon (\bar{F}_{\theta_L + \xi_0}^s (s_0^*) - \bar{F}_{\theta_L + \xi}^s (s_0^*)) d\xi < \int f^\xi (\xi) f^\varepsilon (\bar{F}_{\theta_L + \xi_0}^s (s_0^*) - \bar{F}_{\theta_H + \xi}^s (s_0^*)) d\xi$$

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


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

$$\tau^* = (E_{\theta, \xi} [f_{\theta + \xi}^p (p^*) | \mathcal{I}])^{-1} \beta \frac{\partial E_{\theta, \xi} [V(\mathcal{I}') | \mathcal{I}]}{\partial p^*}, \quad ((15))$$

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, (8) shows that the marginal investor p^* is such that $p^* \theta_H + (1 - p^*) \theta_L = c + \tau$. Combining with (14) and reorganizing yields (15). \square

A.3 Additional figures



Notes: The simulation was performed with parameters: $\theta_H = 1$, $\theta_L = 0.5$, $\bar{\xi} = 0.4$, $c = 0.79$. 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 11: Impulse response in the case of a true positive



Notes: The simulation was performed with parameters: $\theta_H = 1$, $\theta_L = 0.5$, $\bar{\xi} = 0.4$, $c = 0.79$. 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 12: Impulse response in the case of a true negative

A.4 Intensive margin of investment

In this appendix, we consider an extension of the 3-state version of the simplified learning model from section 2 in which firms invest along the intensive margin. As a result, investment i_t can take on any values $i_{jt} \geq 0$ instead of being limited to $i_{jt} \in \{0, 1\}$. We will show that the boom-bust dynamics and the information cascades explored in Section 3.1 survive to that change in investment technology.

Most of the equations of the model remain the same. The main change is that the return equation (2) now becomes

$$y_{jt} = R_t i_{jt} - c(i_{jt}), \quad (27)$$

where, as before, R_t is (1) and where c is now a smooth investment cost function that we assume to be strictly increasing and strictly convex with $c(0) = 0$ and $c'(0) = 0$. To generate information cascades it is often convenient to assume a large amount of curvature in the cost function c for higher values of investment. One simple way to capture that feature is to assume that there is a threshold \bar{i} such that $i_{jt} \leq \bar{i}$. In a richer model, this threshold could come from a budget constraint, a financial constraint, decreasing returns in production, etc.³⁰

Under the return technology (27), agents pick i_{jt} to maximize

$$(p_{jt}\theta_h + (1 - p_{jt})\theta_l) i_{jt} - c(i_{jt}),$$

and the first-order condition implies that the optimal investment decision is given by

$$i_{jt} = \begin{cases} \bar{i} & \text{if } (c')^{-1}(p_{jt}\theta_h + (1 - p_{jt})\theta_l) \geq \bar{i} \\ (c')^{-1}(p_{jt}\theta_h + (1 - p_{jt})\theta_l) & \text{if } p_{jt}\theta_h + (1 - p_{jt})\theta_l \geq 0 \\ 0 & \text{otherwise,} \end{cases} \quad (28)$$

where $(c')^{-1}$ denotes the inverse of the marginal cost function. To simplify the notation, we define the investment function $i(\Lambda_t, s_j)$ as the optimal investment, given by (28), of an agent with a private signal s_j when the public beliefs are Λ_t .

We also need to adapt the endogenous public signal to the new investment technology. We assume that firms

³⁰This threshold is neither necessary to generate booms-and-busts, nor information cascades. We introduce it to generate sharper information cascades.

observe the total amount of investment in the economy, such that

$$m_t = m^e(\Lambda_t, \theta, \xi) + \varepsilon_t \quad (29)$$

$$\text{where } m^e(\Lambda_t, \theta, \xi) = \int i(\Lambda_t, s_j) f_{\theta+\xi}^s(s_j) ds_j. \quad (30)$$


Under these assumptions, the updating of beliefs follows the same equations as Section 2.6, such that (11) and (12) remain valid.

Simulations

It is straightforward to simulate this model numerically. For that purpose, we assume that $c(i) = \frac{i^2}{2\alpha}$ so that

$$i = \min \{ \alpha (p_{jt}\theta_h + (1 - p_{jt})\theta_l), \bar{i} \}$$

when investment takes place. Figure 13 shows the impulse response of total investment m and the public beliefs in the false-positive state. As we can see, this model can also generate boom-bust cycles, with aggregate investment rising for 64 periods before starting to decline. The logic is the same as in Section 2: agents rationally interpret the large aggregate investment as evidence that the fundamental is $\theta = \theta_h$, which pushes for further investment. After several periods, as more and more information is released, agents realize that the false-positive state is the most likely and aggregate investment declines. The example also shows that slow booms and rapid busts are compatible with an intensive investment margin.




Notes: The simulation was performed with parameters: $\theta_H = 1$, $\theta_L = 0.5$, $\bar{\xi} = 0.48$, $\alpha = 12$, $\bar{i} = 7.5$. The priors are set to $p_0 = 0.05$ and $q_0 = 0.04$. All the distributions are Gaussian: $F_{\theta+\xi}^s \sim \mathcal{N}(\theta + \xi, \sigma_s)$, $F^\varepsilon \sim \mathcal{N}(0, \sigma_\varepsilon)$ and $F^u \sim \mathcal{N}(0, \sigma_u)$ with standard deviations $\sigma_s = 0.2$, $\sigma_\varepsilon = 2.5$, $\sigma_u = 2.5$.

Figure 13: Impulse response in the false-positive state with an intensive investment margin

The model with an intensive investment margin can also generate information cascades. Figure 14 provides an example. We see that aggregate investment first goes through a rapid increase to a plateau in which firms all invest at levels close to \bar{i} . As in Figure 4 in the main text, an information cascade follows, in which little information is transmitted through the observation of aggregate investment (agents still learn from the exogenous public signal). As a result, the good and the bad states remain a likely possibility for an extended period of time, and agents have to wait until period $t = 50$ for the false-positive state to become the most likely outcome. As we can see from this example, information cascades remain even when agents are allowed to invest along an intensive margin. Note, finally, that this type of information cascade due to binding constraints is reminiscent of the mechanism present in

Straub and Ulbricht (2017; 2019).




Notes: The simulation was performed with parameters: $\theta_H = 1$, $\theta_L = 0.5$, $\bar{\xi} = 0.45$, $\alpha = 10$, $\bar{i} = 7.5$. The priors are set to $p_0 = 0.25$ and $q_0 = 0.05$. All the distributions are Gaussian: $F_{\theta+\xi}^s \sim \mathcal{N}(\theta + \xi, \sigma_s)$, $F^\varepsilon \sim \mathcal{N}(0, \sigma_\varepsilon)$ and $F^u \sim \mathcal{N}(0, \sigma_u)$ with standard deviations $\sigma_s = 0.25$, $\sigma_\varepsilon = 0.8$, $\sigma_u = 2$.

Figure 14: Impulse response in the false-positive state with a bounded intensive investment margin

A.5 Permanent information cascades

In this appendix, we show that a slight modification of the environment considered in Figure 4 in the main text can lead to a permanent information case in which no information is provided by the observation of the mass of investors. We depart from that specification by imposing a lower bound $\underline{s} = 0$ on the distribution of private signal and by making the exogenous public signal uninformative ($\sigma_u = \infty$). Figure 15 shows the behavior of the economy in the false-positive state under this new parametrization. As we can see, the economy quickly converges to a permanent plateau at $m = 1$. The public beliefs are so optimistic at that point that even the most pessimistic agent ($s_j = 0$) invests. Since all agents invest regardless of the true state of the world, observing the mass of investors provides no information and the economy remains in a permanently elevated state of investment.



Notes: The simulation was performed with parameters: $\theta_H = 1$, $\theta_L = 0.5$, $\bar{\xi} = 0.4$, $c = 0.79$. The priors are set to $p_0 = 0.25$ and $q_0 = 0.05$. The distributions of all signals are Gaussian with standard deviations $\sigma_s = 0.5$, $\sigma_\varepsilon = 0.2$, $\sigma_u = \infty$. The private signal distribution is bounded below by zero.

Figure 15: Impulse response in the false-positive state with permanent information cascades

B Appendix of Section 5

This appendix contains additional details about the full quantitative model, our resolution method, the calibration as well as exercises to test the robustness of the simulations of Figure 9.

B.1 Equilibrium definition

In what follows, to lighten notation, we denote by Ω_t^A the aggregate state space at the beginning of a period upon making the technology adoption decision (stage A of the timing) and Ω_t^B the aggregate state space at the end of the period when consumption, production and market clearing take place (stage B). Before maturation of the technology, the state space that describes the information set of the household and any participant that is not an entrepreneur is $\Omega_t^A = (K_t^{IT}, K_t^T, \Lambda_t)$. The individual state space that describes the information set of entrepreneurs at stage A is (Ω_t^A, s_j) . Finally, the aggregate state space that describes the information set at stage B is $\Omega_t^B = (K_t^{IT}, K_t^T, \Lambda_t, m_t, s_t^u)$. After maturation, all the uncertainty related to the technology is resolved, so the aggregate state spaces can be described by $\Omega_t^A = (K_t^{IT}, K_t^T, \theta)$ and $\Omega_t^B = (K_t^{IT}, K_t^T, \theta, m_t)$.

We are now ready to define an equilibrium for the full quantitative model. All the model's equations that the definition refers to have been compiled in subsection B.2.

Definition 2. An equilibrium is a set of a) entrepreneur j 's technology adoption decision $i_{jt} (\Omega_t^A, s_j)$, b) a measure of new technology adopters $m_t (\Omega_t^A, \theta, \xi, \varepsilon_t)$, c) household decisions $\{C_t, L_t, I_t^{IT}, I_t^T, K_{t+1}^{IT}, K_{t+1}^T, B_t\}$ which are functions of Ω_t^B , d) firm decisions $\{Y_t, Y_t^w, Y_t^n, Y_t^o, Y_{jt}^r, L_t^n, L_t^o, K_t^n, K_t^o, K_t^{n,IT}, K_t^{n,T}, K_t^{o,IT}, K_t^{o,T}, \Pi_t^{IT}, \Pi_t^T\}$ which are functions of Ω_t^B , e) prices $\{P_t, P_t^w, P_t^n, P_t^o, w_t, R_t, \pi_t, z_t^n, z_t^o, z_t^{IT}, z_t^T, \mu_t\}$ which are functions of Ω_t^B for all t such that

1. The entrepreneurs' technology adoption decision and the measure of new technology adopters solve (53)-(54);
2. The household decisions satisfy the household's optimality conditions and constraints (37)-(43);
3. The firm's decisions satisfy the its optimality conditions (44)-(52) and the production functions (31)-(36);
4. Prices satisfy (60)-(60), the Phillips curve (66), the (18) and the market clearing constraints (55)-(59).

B.2 Model equations

We provide below all the 37 nonlinear equations that characterize the full business cycle model (36 unknowns).³¹

1. Production functions and capital bundles

$$Y_t^n = A_t^n (K_t^n)^\alpha (L_t^n)^{1-\alpha} \quad (\text{individual production new technology}) \quad (31)$$

$$Y_t^o = A_t^o (K_t^o)^\alpha (L_t^o)^{1-\alpha} \quad (\text{individual production old technology}) \quad (32)$$

$$Y_t^w = \left(\int_0^1 m_t (Y_t^n)^{\frac{\sigma-1}{\sigma}} + (1-m_t) (Y_t^o)^{\frac{\sigma-1}{\sigma}} \right)^{\frac{\sigma}{\sigma-1}} \quad (\text{aggregate production wholesale sector}) \quad (33)$$

$$Y_t = \left(\int_0^1 (Y_{jt}^r)^{\frac{\sigma-1}{\sigma}} dj \right)^{\frac{\sigma}{\sigma-1}} \quad (\text{aggregate production final good}) \quad (34)$$

$$K_t^n = \kappa_n \left((1-\omega_n) (K_t^{n,T})^{\frac{\zeta-1}{\zeta}} + \omega_n (K_t^{n,IT})^{\frac{\zeta-1}{\zeta}} \right)^{\frac{\zeta}{\zeta-1}} \quad (\text{new tech. capital bundle}) \quad (35)$$

$$K_t^o = \kappa_o \left((1-\omega_o) (K_t^{o,T})^{\frac{\zeta-1}{\zeta}} + \omega_o (K_t^{o,IT})^{\frac{\zeta-1}{\zeta}} \right)^{\frac{\zeta}{\zeta-1}} \quad (\text{old tech. capital bundle}) \quad (36)$$

³¹One equation is redundant because of Walras' Law, but we keep it in the equilibrium definition for the sake of completeness.

2. Household optimality and constraints

$$C_t^{-1} = \beta E_t \left[C_{t+1}^{-1} \frac{1 + R_t}{1 + \pi_{t+1}} \right] \quad (\text{Euler eq. nominal bonds}) \quad (37)$$

$$C_t^{-1} = \beta E_t \left[C_{t+1}^{-1} (z_{t+1}^{IT} + 1 - \delta) \right] \quad (\text{Euler eq. IT capital}) \quad (38)$$

$$C_t^{-1} = \beta E_t \left[C_{t+1}^{-1} (z_{t+1}^T + 1 - \delta) \right] \quad (\text{Euler eq. T capital}) \quad (39)$$

$$L_t^\psi = w_t \quad (\text{labor supply}) \quad (40)$$

$$K_{t+1}^{IT} = (1 - \delta) K_t^{IT} + I_t^{IT} \quad (\text{law of motion IT capital}) \quad (41)$$

$$K_{t+1}^T = (1 - \delta) K_t^T + I_t^T \quad (\text{law of motion T capital}) \quad (42)$$

$$C_t + \sum_{i=T,IT} I_t^i + \frac{B_t}{P_t} = w_t L_t + \sum_{i=T,IT} z_t^i K_t^i + \frac{1 + R_{t-1}}{1 + \pi_t} \frac{B_{t-1}}{P_{t-1}} + m_t \Pi_t^{IT} + (1 - m_t) \Pi_t^T \quad (\text{budget constraint}) \quad (43)$$

3. Firm optimality

$$\left(A_t^n (K_t^n)^\alpha (L_t^n)^{1-\alpha} \right)^{\frac{\sigma-1}{\sigma}} / L_t^n = \left(A_t^o (K_t^o)^\alpha (L_t^o)^{1-\alpha} \right)^{\frac{\sigma-1}{\sigma}} / L_t^o \quad (\text{wage equalization}) \quad (44)$$

$$\frac{\alpha}{1 - \alpha} = \frac{z_t^n K_t^n}{w_t L_t^n} \quad (\text{cost minimization new tech.}) \quad (45)$$

$$\frac{\alpha}{1 - \alpha} = \frac{z_t^o K_t^o}{w_t L_t^o} \quad (\text{cost minimization old tech.}) \quad (46)$$

$$K_t^{n,IT} = \kappa_n^{\zeta-1} \omega_n^\zeta \left(\frac{z_t^{IT}}{z_t^n} \right)^{-\zeta} K_t^n \quad (\text{new tech. demand for IT capital}) \quad (47)$$

$$K_t^{n,T} = \kappa_n^{\zeta-1} (1 - \omega_n)^\zeta \left(\frac{z_t^T}{z_t^n} \right)^{-\zeta} K_t^n \quad (\text{new tech. demand for T capital}) \quad (48)$$

$$K_t^{o,IT} = \kappa_o^{\zeta-1} \omega_o^\zeta \left(\frac{z_t^{IT}}{z_t^o} \right)^{-\zeta} K_t^o \quad (\text{old tech. demand for IT capital}) \quad (49)$$

$$K_t^{o,T} = \kappa_o^{\zeta-1} (1 - \omega_o)^\zeta \left(\frac{z_t^T}{z_t^o} \right)^{-\zeta} K_t^o \quad (\text{old tech. demand for T capital}) \quad (50)$$

$$\Pi_t^{IT} = \frac{1}{\sigma - 1} \left(\frac{\sigma}{\sigma - 1} \right)^{-\sigma} P_t^{1-\sigma} (P_t^w)^\sigma Y_t^w \left(\frac{1}{A_t^n} \left(\frac{z_t^n}{\alpha} \right)^\alpha \left(\frac{w_t}{1 - \alpha} \right)^{1-\alpha} \right)^{1-\sigma} \quad (\text{new tech. profits}) \quad (51)$$

$$\Pi_t^T = \frac{1}{\sigma - 1} \left(\frac{\sigma}{\sigma - 1} \right)^{-\sigma} P_t^{1-\sigma} (P_t^w)^\sigma Y_t^w \left(\frac{1}{A_t^o} \left(\frac{z_t^o}{\alpha} \right)^\alpha \left(\frac{w_t}{1 - \alpha} \right)^{1-\alpha} \right)^{1-\sigma} \quad (\text{old tech. profits}) \quad (52)$$

4. Technology adoption decisions

$$i_{jt} = \operatorname{argmax}_{i \in [0,1]} i E \left[\Pi_t^n | \Omega_t^A, s_j \right] + (1 - i) E \left[\Pi_t^o | \Omega_t^A, s_j \right] \quad (\text{optimal technology choice}) \quad (53)$$

$$m_t = \int_0^{1-\mu} i_{jt} dj + \mu \varepsilon_t \quad (\text{measure new tech. adopters}) \quad (54)$$

5. Market clearing constraints

$$Y_t = C_t + I_t^{IT} + I_t^T \quad (\text{final good}) \quad (55)$$

$$Y_t^w = \int Y_{jt}^r dj \quad (\text{wholesale good}) \quad (56)$$

$$L_t = m_t L_t^n + (1 - m_t) L_t^o \quad (\text{labor}) \quad (57)$$

$$K_t^{IT} = m_t K_t^{n,IT} + (1 - m_t) K_t^{o,IT} \quad (\text{IT capital}) \quad (58)$$

$$K_t^T = m_t K_t^{n,T} + (1 - m_t) K_t^{o,T} \quad (\text{T capital}) \quad (59)$$

6. Prices

$$z_t^n = \frac{1}{\kappa_n} \left((1 - \omega_n)^\zeta (z_t^T)^{1-\zeta} + \omega_n^\zeta (z_t^{IT})^{1-\zeta} \right)^{\frac{1}{1-\zeta}} \quad (\text{price new tech. capital bundle}) \quad (60)$$

$$z_t^o = \frac{1}{\kappa_o} \left((1 - \omega_o)^\zeta (z_t^T)^{1-\zeta} + \omega_o^\zeta (z_t^{IT})^{1-\zeta} \right)^{\frac{1}{1-\zeta}} \quad (\text{price old tech. capital bundle}) \quad (61)$$

$$P_{nt}/P_t = \frac{\sigma}{\sigma - 1} \frac{1}{A_{nt}} \left(\frac{z_t^n}{\alpha} \right)^\alpha \left(\frac{w_t}{1 - \alpha} \right)^{1-\alpha} \quad (\text{price new tech.}) \quad (62)$$

$$P_{ot}/P_t = \frac{\sigma}{\sigma - 1} \frac{1}{A_o} \left(\frac{z_t^o}{\alpha} \right)^\alpha \left(\frac{w_t}{1 - \alpha} \right)^{1-\alpha} \quad (\text{price old tech.}) \quad (63)$$

$$P_t^w = \left(m_t (P_t^n)^{1-\sigma} + (1 - m_t) (P_t^o)^{1-\sigma} \right)^{\frac{1}{1-\sigma}} \quad (\text{price wholesale}) \quad (64)$$

$$\mu_t = P_t/P_t^w \quad (\text{markup retail sector}) \quad (65)$$

7. New Keynesian block (log-linearized)

$$\hat{\pi}_t = \beta E_t \hat{\pi}_{t+1} - \frac{(1 - \chi)(1 - \beta\chi)}{\chi} \hat{\mu}_t \quad (\text{Phillips curve}) \quad (66)$$

$$\hat{R}_t = \phi_\pi \hat{\pi}_t + \phi_y (\hat{Y}_t - \bar{Y}) \quad (\text{Taylor rule}) \quad (67)$$

B.3 Resolution method

Our resolution approach relies on linearizing the equations of the economic model while keeping the full nonlinear equations that govern the evolution of beliefs. Besides providing greater tractability, this approach allows us to focus on the nonlinearities implied by our learning model while putting aside those of the DSGE model, which are usually weak.

Exploiting the martingale property of beliefs

As explained in Section 5.1, we exploit the property that beliefs follow a martingale under rational learning. Because of this property coupled with the linearization, the linear system of equations that characterize the economy can be solved at time t to a first-order approximation as if beliefs were held constant going forward, i.e., as if $\Lambda_{t+s}(\theta, \xi) = \Lambda_t(\theta, \xi)$, $\forall s \geq 0$. In other words, the beliefs Λ_t are no longer a state variable when solving the economic model forward, avoiding the need to keep track of a potentially large object in the state space (infinite dimensional in the case of continuous ξ).

Linearization around moving steady states

The interaction between technology adoption and beliefs is of a highly nonlinear nature and the economy can settle on very different steady states in the long run. We have found that linearizing the economy around a fixed non-stochastic

steady state produces poor results and that linearizing instead around a moving steady state that takes into account these nonlinearities yields more accurate predictions.

Our preferred approach thus involves linearizing the model every period around the long-run non-stochastic steady state in which all random variables are set equal to their means according to current beliefs.³² In particular, the productivity of the new technology is equal to the current average public expectation at time t with $\bar{A}^n(\Lambda_t) = E[A_{t+s}^n | \mathcal{I}_t] = (1 - \lambda) A^o + \lambda (p_t \theta_H + (1 - p_t) \theta_L)$, where $p_t = \int \Lambda_t(\theta_H, \xi) d\xi$. We refer to this steady state as $S.S(\Lambda_t)$.

In a non-stochastic steady state, the rental rates on the two types of capital are equalized and the threshold $s^{*,S.S(\Lambda_t)}$ which characterizes the marginal technology adopter solves

$$E \left[\theta | \Lambda_t, s^{*,S.S(\Lambda_t)} \right] = A^o.$$

The corresponding expected measure of new technology adopters is given by

$$m^{S.S(\Lambda_t)} = E \left[(1 - \mu) \bar{F}_{\theta+\xi}^s \left(s^{*,S.S(\Lambda_t)} \right) + \mu \varepsilon_t \mid \Lambda_t \right].$$

Every period, the model is solved forward around the long-run steady state associated to current beliefs but the variables are then translated back to a reference steady state which stays the same throughout the simulation in order to construct the impulse responses.³³

Technology adoption decision with GE effects

In each period, given the current capital stocks and public beliefs $\Omega_t^A = (K_t^{IT}, K_t^T, \Lambda_t)$, we solve nonlinearly for the marginal technology adopter threshold $s_t^* = s^*(\Omega_t, \Lambda_t)$ ³⁴ which solves

$$E \left[\hat{A}_t^n - \alpha \hat{z}_t^n \mid \Lambda_t, s_t^* \right] = E \left[\underbrace{\hat{A}_t^o}_{=0} - \alpha \hat{z}_t^o \mid \Lambda_t, s_t^* \right],$$

where the hatted variables indicate log-deviations from the current steady state around which the economy is linearized.

Solving for the threshold thus requires solving for the relative rental rates of capital. For a given stock of capital (K_t^{IT}, K_t^T) , m_t and productivity realization A_t^n , the relative rental rate of capital $\hat{z}_t^n - \hat{z}_t^o = \widehat{z_t^n / z_t^o}$ can be computed by solving a static system of equations. We first combine the firms' first-order conditions over capital and labor to express the relative demands for capital

$$\frac{k_{nt}}{k_{ot}} = \left(\frac{A_t^n}{A_o} \right)^{\sigma-1} \left(\frac{z_t^n}{z_t^o} \right)^{-\sigma(1-(1-\alpha)\frac{\sigma-1}{\sigma})}. \quad (68)$$

Using the IT and traditional capital market clearing constraints

$$K_t^{IT} = m_t \kappa_n^{\zeta-1} \omega_n^\zeta \left(\frac{z_t^{IT}}{z_t^n} \right)^{-\zeta} k_t^n + (1 - m_t) \kappa_o^{\zeta-1} \omega_o^\zeta \left(\frac{z_t^{IT}}{z_t^o} \right)^{-\zeta} k_t^o \quad (69)$$

$$K_t^T = m_t \kappa_n^{\zeta-1} (1 - \omega_n)^\zeta \left(\frac{z_t^T}{z_t^n} \right)^{-\zeta} k_t^n + (1 - m_t) \kappa_o^{\zeta-1} (1 - \omega_o)^\zeta \left(\frac{z_t^T}{z_t^o} \right)^{-\zeta} k_t^o, \quad (70)$$

³²This includes not only the productivity but also the signal noise and the mass of noise entrepreneurs.

³³Note that the output gap in the Taylor rule (18), expressed in terms of the reference steady state, must be translated to the current steady state in every period.

³⁴Due to the presence of general equilibrium effects, even under the assumption of MLRP for $\{F_x^s\}$, the technology adoption decision is not guaranteed to take the form of a cutoff rule. We check numerically that the necessary monotonicity condition holds to ensure the existence of a cutoff s_t^* .

and definitions of the capital bundle prices

$$z_t^n = \frac{1}{\kappa_n} \left((1 - \omega_n)^\zeta (z_t^T)^{1-\zeta} + \omega_n^\zeta (z_t^{IT})^{1-\zeta} \right)^{\frac{1}{1-\zeta}} \quad (71)$$

$$z_t^o = \frac{1}{\kappa_o} \left((1 - \omega_o)^\zeta (z_t^T)^{1-\zeta} + \omega_o^\zeta (z_t^{IT})^{1-\zeta} \right)^{\frac{1}{1-\zeta}}, \quad (72)$$

we obtain a system of 5 equations (68)-(72) in 5 unknowns $(k_t^n, k_t^o, z_t^{IT}, z_t^T, z_t^n/z_t^o)$ which can be solved using a nonlinear solver.

Our procedure to solve for m_t in a given period with state variables $(K_t^{IT}, K_t^T, \Lambda_t)$ involves a simple bisection algorithm which we describe as follows:

1. Guess s_t^* ;
2. Using certainty equivalence, set $A_t^n \equiv E[A_t^n | \Lambda_t, s_t^*]$ and $m_t \equiv E_t[m_t | \Lambda_t, s_t^*]$ and solve for $\hat{z}_t^n - \hat{z}_t^o$ using the system (68)-(72).³⁵;
3. Check if $\left| E \left[\hat{A}_t^n - \alpha (\hat{z}_t^n - \hat{z}_t^o) | \Lambda_t, s_t^* \right] \right| < \varepsilon$ with ε small. If yes, stop, otherwise return to (1) with a lower s_t^* if the expectation strictly positive, or a higher s_t^* if negative.

We denote the resulting technology adoption threshold by $s^*(K_t^{IT}, K_t^T, \Lambda_t)$.

Dynamics of m_t

The previous paragraph explains how s_t^* can be computed in a given period t for a certain state space $(K_t^{IT}, K_t^T, \Lambda_t)$. One difficulty remains in the fact that agents in the economic model need to form expectations about future m_{t+s} , $s > 0$, and that one equation in our linearized system should govern the dynamics of m_t . We do so by numerically linearizing the expectation of m_{t+s} as a function of $(K_{t+s}^{IT}, K_{t+s}^T, \Lambda_t)$.

To fix notation, denote m^{exp} the public expectation of m that corresponds to the technology adoption threshold $s^*(K^{IT}, K^T, \Lambda_t)$:

$$m^{exp}(K^{IT}, K^T, \Lambda_t) = E \left[(1 - \mu) \bar{F}_{\theta+\xi}^s \left(s^*(K^{IT}, K^T, \Lambda_t) \right) + \mu \varepsilon_t \mid \Lambda_t \right].$$

In the long-run steady state around which the economy is linearized, the expected measure of new technology adopters $m^{S.S(\Lambda_t)}$ coincides with m^{exp} when the stocks of capital have reached their steady states $(K^{IT,S.S(\Lambda_t)}, K^{T,S.S(\Lambda_t)})$:

$$m^{S.S(\Lambda_t)} = m^{exp}(K^{IT,S.S(\Lambda_t)}, K^{T,S.S(\Lambda_t)}, \Lambda_t).$$

The linearized equation that governs the expectation of m is given by

$$\begin{aligned} E[m_{t+s} | \Lambda_t] &= m^{exp}(K_{t+s}^{IT}, K_{t+s}^T, \Lambda_t) = m^{S.S(\Lambda_t)} + \frac{\partial m^{exp}}{\partial K^{IT}}(K^{IT,S.S(\Lambda_t)}, K^{T,S.S(\Lambda_t)}, \Lambda_t) (K_{t+s}^{IT} - K^{IT,S.S(\Lambda_t)}) \\ &+ \frac{\partial m^{exp}}{\partial K^T}(K^{IT,S.S(\Lambda_t)}, K^{T,S.S(\Lambda_t)}, \Lambda_t) (K_{t+s}^T - K^{T,S.S(\Lambda_t)}). \end{aligned} \quad (73)$$

We use numerical differentiation to evaluate the two derivatives $\frac{\partial m^{exp}}{\partial K^{IT}}$ and $\frac{\partial m^{exp}}{\partial K^T}$ and feed equation (73) to our linear rational expectation model solver.

Resolution method step by step

We start by selecting a reference steady state (indicated with the upperscript *ref*) with respect to which all the impulse responses are expressed. We choose the long-run nonstochastic steady state associated with the initial public

³⁵Since the economy is eventually linearized, the use of certainty equivalence has no bearing on the results, which remain correct to the first order.

beliefs Λ_0 at the time 0 with the corresponding technology adoption $m^{S.S(\Lambda_0)}$.

In what follows, hatted variables \hat{x}^{ref} or $\hat{x}^{S.S(\Lambda_t)}$ denote (log) deviations from the particular steady state it refers to (*ref* or *S.S* (Λ_t)). The various steps of the resolution method can be described as follows:

1. At a given date t with state variables $(\hat{K}_t^{IT,ref}, \hat{K}_t^{T,ref}, \Lambda_t)$, compute the long-run steady state associated with current beliefs *S.S* (Λ_t): $(Y^{S.S(\Lambda_t)}, m^{S.S(\Lambda_t)}, K^{IT,S.S(\Lambda_t)}, K^{T,S.S(\Lambda_t)}, \dots)$;
2. Translate the state variables $(\hat{K}_t^{IT,ref}, \hat{K}_t^{T,ref})$ to $(\hat{K}_t^{IT,S.S(\Lambda_t)}, \hat{K}_t^{T,S.S(\Lambda_t)})$;
3. Compute the threshold $s^*(K_t^{IT}, K_t^T, \Lambda_t)$ and the realization of $m_t = (1 - \mu) \bar{F}_{\theta+\xi}^s(s^*(K_t^{IT}, K_t^T, \Lambda_t)) + \mu \varepsilon_t$; express it in deviation from the current steady state $\hat{m}_t^{S.S(\Lambda_t)} = m_t - m^{S.S(\Lambda_t)}$;
4. Use numerical differentiation to evaluate the law of $m^{exp}(K_{t+s}^{IT}, K_{t+s}^T, \Lambda_t)$ and express it with respect to the current steady state associated to Λ_t ;
5. Using a linear rational expectation solver, solve for the economy linearized around the current steady state *S.S* (Λ_t), keeping beliefs constant;
6. Evaluate the impulse response associated with the chosen set of shocks from t onward;
7. Collect the time- t variables $(\hat{Y}_t^{S.S(\Lambda_t)}, \hat{K}_t^{IT,S.S(\Lambda_t)}, \hat{K}_t^{T,S.S(\Lambda_t)}, \dots)$ and translate them back to the reference steady state $(\hat{Y}_t^{ref}, \hat{K}_t^{IT,ref}, \hat{K}_t^{T,ref}, \dots)$. Store their values as the time- t realization of the desired impulse response;
8. Update beliefs using the particular realization of the signals and m_t ;
9. Set $t := t + 1$ and go back to (1).

B.4 Additional details about the calibration

We use data from the “nowcast” of real GDP in the Survey of Professional Forecasters as they capture beliefs about net changes in productivity, net of changes in factors shares, because factors don’t have time to adjust over this short time horizon. The average real GDP growth between between 1991 and 1998 (before the boom-bust cycle started) is 2.4%, and we assume that we are on a balanced growth path with A_{ot} , θ_{Ht} and θ_{Lt} growing at rate $\gamma = 1 + 0.024/4 = 1.006$ per quarter.

To compute θ_H , we use the highest expectation of real GDP growth over our period, which is 4.19% in 2000Q2 in annualized term. We interpret this most optimistic scenario as a situation in which everybody is convinced that the state of the new technology is good and so $m = 1$ and $p = 1$. In this case, we can write

$$\frac{\gamma(\lambda\theta_H + (1 - \lambda)A_o)}{A_o} = 1 + 0.0419/4$$

where recall $\lambda = 1/22$ is the likelihood that the new technology becomes operational. Since A_o is normalized to one we find that $\theta_H = 1.099$.

We perform a similar computation to find θ_L but in this case relying on the most pessimistic ($m = 1$ and $p = 1$) forecast (0.8% annualized in 2001Q1). We find

$$\frac{\gamma(\lambda\theta_L + (1 - \lambda)A_o)}{A_o} = 1 + 0.008/4$$

which yields $\theta_L = 0.912$.

B.5 Mass of noise entrepreneurs μ

Figure 16 shows the benchmark simulation of Section 5 with a smaller ($\mu = 0.1$) and larger ($\mu = 0.2$) mass of noise entrepreneurs. We see that when there are fewer noise entrepreneurs, the public signals that agents receive are more precise which leads to a stronger boom that happens sooner. In contrast, if the public signals are more noisy

($\mu = 0.2$), agents learn little by observing m and the economy does not go through a boom-bust cycle. This last simulation emphasizes the importance of social learning for the mechanism.

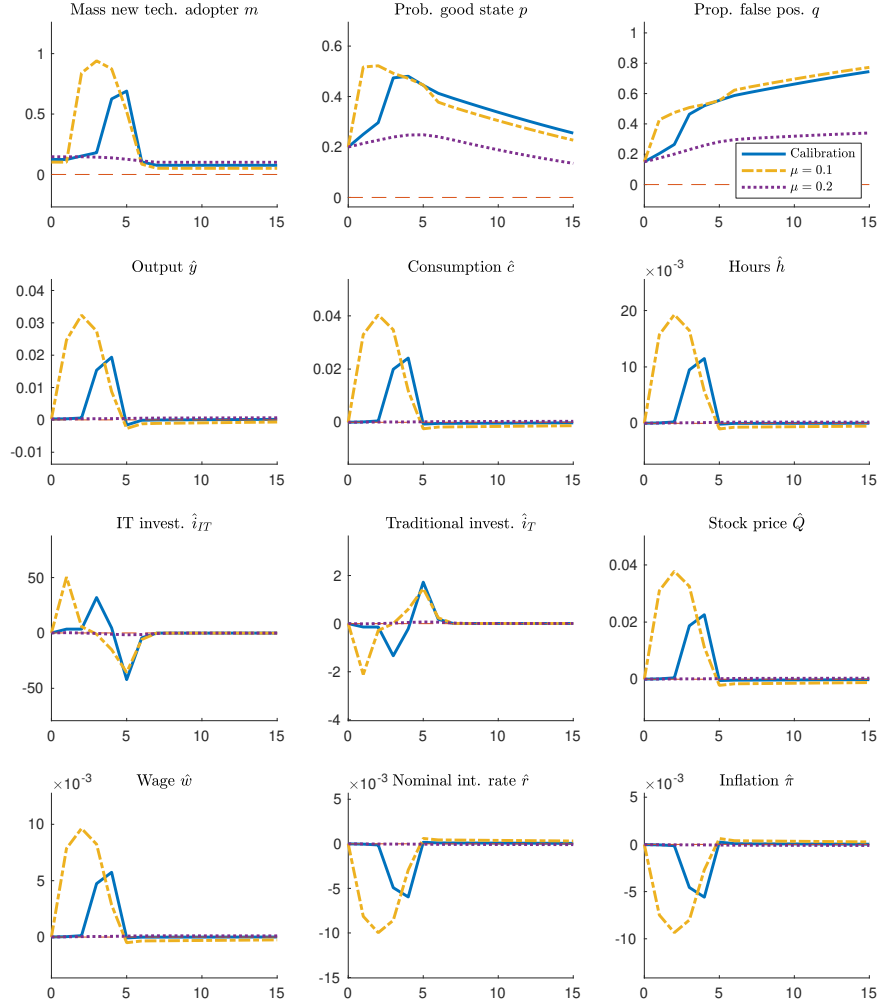


Figure 16: Simulations under different mass of noise traders μ

B.6 Size of common noise signal $\bar{\xi}$

Figure 17 shows the same simulation as Section 5 but with various sizes for the common noise shock $\bar{\xi}$. Unsurprisingly, the size of the boom is correlated with $\bar{\xi}$. When the common noise essentially mimics θ_H the increase in output reaches 3 percent. In contrast, when $\bar{\xi} = 0.9 \times (\theta_H - \theta_L)$ aggregate output goes up by about 1.5 percent during the boom.

B.7 Prior p_0 associated with the good state: $\theta = \theta_H$ and $\xi = 0$

Figure 18 shows the simulations of Section 5 under various values for the prior p_0 associated with the good state. We see that the boom is much larger when entrepreneurs expect the good state to be more likely ($p_0 = 25$). In that case, the economy goes in the information cascades zone as in Figure 4. In contrast, when agents expect that the good state is unlikely ($p_0 = 0.15$) their pessimism is enough to prevent the economy from going through a boom-bust cycle.

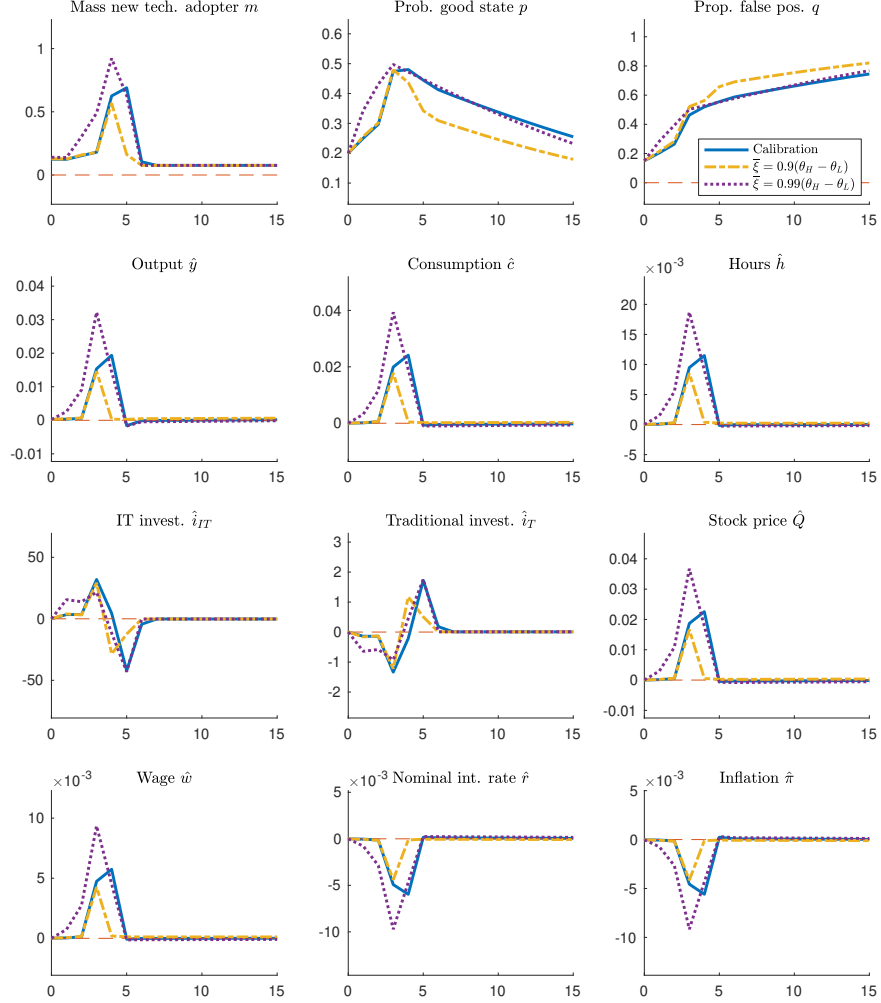


Figure 17: Simulations under different size of the common noise signal $\bar{\xi}$

B.8 Prior q_0 associated with the false positive state: $\theta = \theta_L$ and $\xi = \bar{\xi}$

Figure 19 shows what happens to the benchmark simulations of Section 5 when the prior on the false positive state changes. When that state is less likely ($q_0 = 0.1$) the boom-bust cycle dynamics is stronger and the economy enters the information cascades state, as in Figure 4. If instead, the false positive state is more likely ($q_0 = 0.2$), agents are particularly skeptical of good signal and the economy does go through a boom.

B.9 Standard deviation of the exogenous public signal σ_u

Figure 20 shows the benchmark simulation under different values for σ_u . When the exogenous public signal is less information ($\sigma_u = 4 \times \sigma_s$), agents have a harder time figuring out that the economy is in the false positive state and the boom is much larger, with output reaching a high of about 3 percent above trend. In contrast, when σ_u is small, the exogenous public signal is very informative and agents learn rapidly that the technology is bad ($\theta = \theta_L$). In this case, there is no room for our social learning mechanism to create a boom-bust cycle.




Figure 18: Simulations under different priors p_0

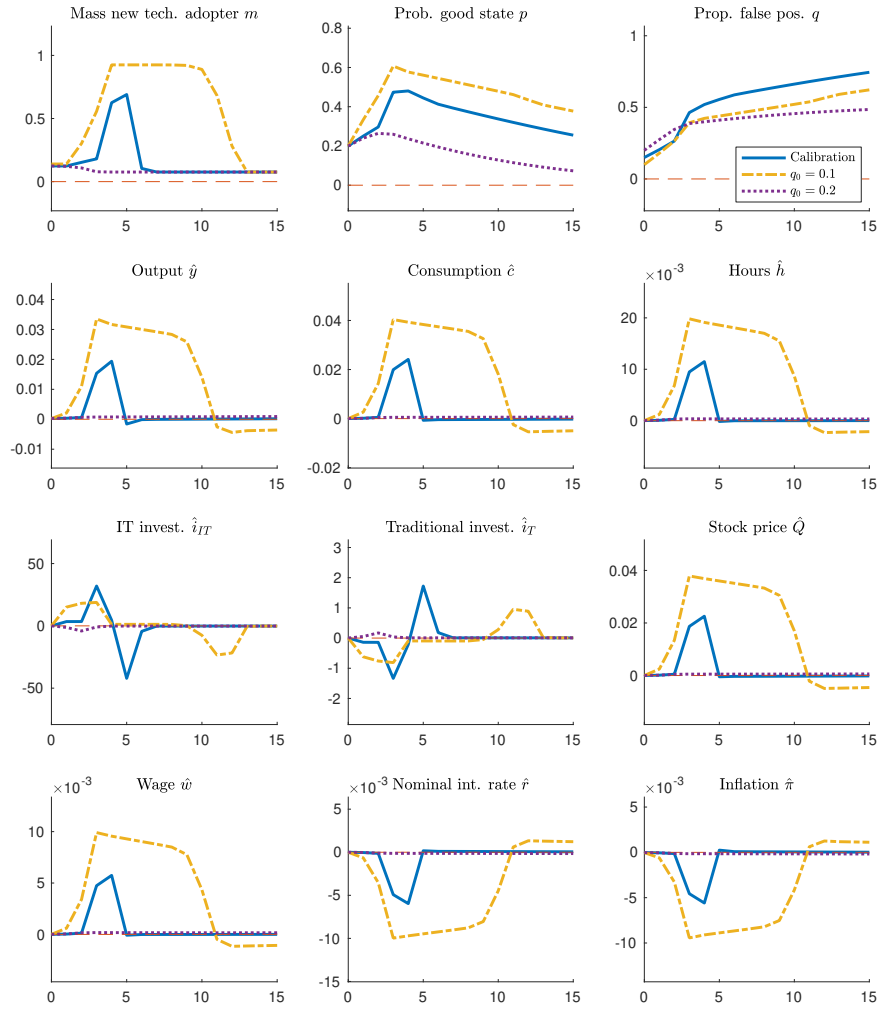


Figure 19: Simulation under different priors q_0

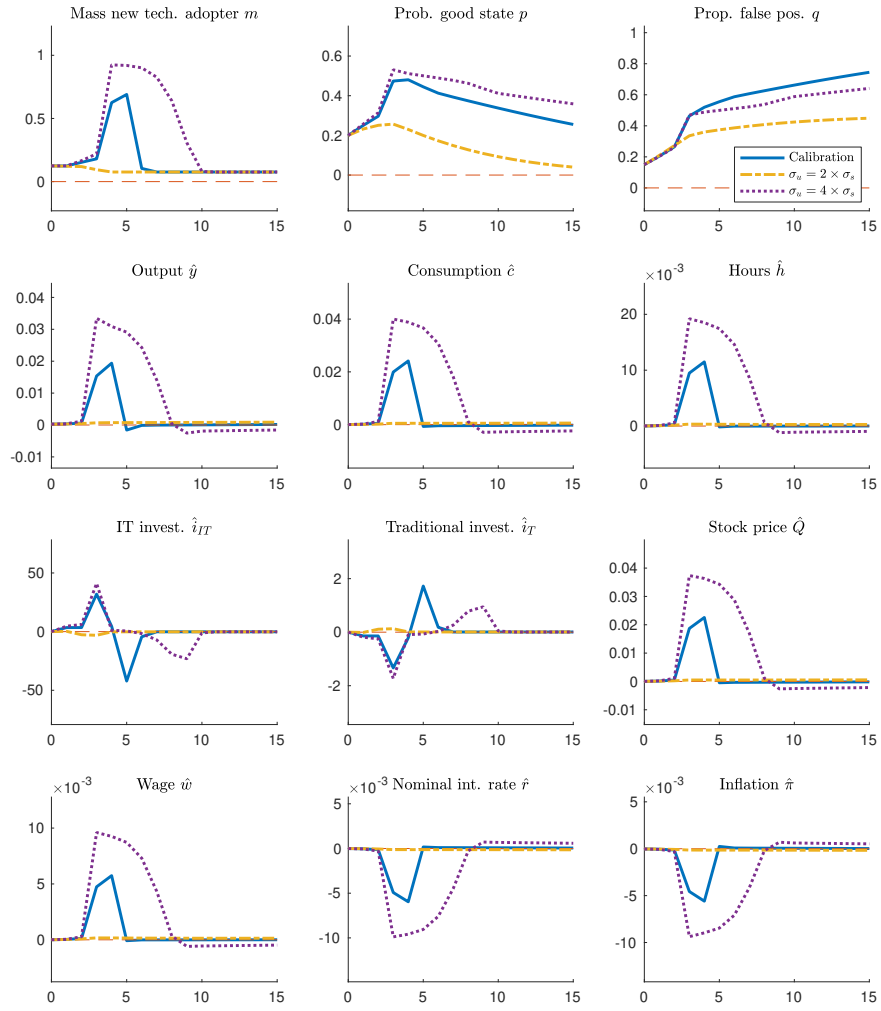


Figure 20: Simulation under different informativeness for the exogenous public signal σ_u