

Speculation-Driven Business Cycles*

Saki Bigio

Department of Economics, UCLA
NBER

Eduardo Zilberman

Department of Economics, PUC-Rio
Research Department, Central Bank of Chile

January 2020

Abstract

We introduce speculation, in the spirit of Harrison and Kreps (1978), into a standard real business cycle model. Investors (speculators) hold heterogeneous beliefs about firm growth. Firm ownership, and thus, the firm's discount factor, varies with waves of optimism and leverage. These waves ripple into the firm's investment in hours. The firm's discount factor links the equity premium and labor volatility puzzles. We find that the amplification that can be generated by speculation for any model of beliefs, for a standard calibration, has an analytic upper bound of 1.5. When we specialize to a model of diagnostic beliefs, the volatility of hours is amplified by a factor of 1.15 and a bubble component of 20 percent in asset prices emerges.

Keywords: Heterogeneous Beliefs; Business Cycles; Asset Prices; Speculation; Bubbles.

JEL Classification: D84, E32, E44, E71, G41.

*Bigio: sbigio@ucla.edu. Zilberman (corresponding author): zilberman@econ.puc-rio.br. We are grateful for helpful conversations with Sebastian Di Tella, Fatih Guvenen, Fernando Mendo and Alp Simsek. We also thank participants at internal presentations in the Central Bank of Chile and PUC-Rio, as well as at the Santiago Macro Workshop 2019 for helpful comments. The views expressed in this paper do not necessarily reflect the position of the Central Bank of Chile. Bigio thanks the hospitality of the Federal Reserve Bank of San Francisco. All errors are ours.

“I’m happy to say I am a Harrison-Kreps-Keynesian.”

Thomas Sargent, New York Times, October 29, 2011.

1 Introduction

One of the most prominent views about deep economic downturns is that, in part, these are due to waves of speculation. This tradition can be traced to historical narratives of Hyman Minsky, Charles Kindleberger, and more recently Robert Shiller. These narratives share common elements. A wave of excess optimism about earnings-growth prospects prompts a wave of speculation. Speculation then triggers asset bubbles, which are fueled by increased leverage. Implicit in this narrative is a call for corrective measures, because once the bubble bursts on Wall Street, the blast wave impacts Main Street. These narratives are supported by evidence that large stock market turnover, high price-earnings (P/E) ratios, and expansionary credit measures increase the hazard rate of recessions (Shiller, 2000). We dub these narratives as the theory of “speculation-driven business cycles.”

Different from models of rational bubbles, the models of Harrison and Kreps (1978) and more recently of Scheinkman and Xiong (2003) formalize the notion of asset-price bubbles and crashes driven by waves of optimism and pessimism. The virtue of these models is that speculation is done in a way that only subtly departs from the discipline of rational expectations. Yet, in these models bubbles are speculative because investors believe they can resell assets to someone more optimistic in the future. These models lay out the necessary ingredients to generate speculative asset-price bubbles (Xiong, 2013) but remain silent about how bubbles spill over to the real economy. To close the loop toward a quantitative model of speculation-driven business cycles, we must embed speculative bubbles into a real business cycle backbone. This step, to the best of our knowledge, is missing. In fact, when asked about his view about the nature of business cycles, Thomas Sargent, a pioneer of the rational expectations movement, gave the following answer:¹

“[...] economists have been working hard to refine rational expectations theory. [...] An influential example of such work is the 1978 QJE paper by Harrison and Kreps. [...], for policymakers to know whether and how they can moderate bubbles, we need to have well-confirmed quantitative versions of such models up and running.”

In this response, Sargent proposes a departure from rational expectations, embraces the theory of “speculation-driven business cycles”, but then calls for a quantitative model that can scrutinize the theory.² However, Sargent’s challenge presents several immediate barriers. For one thing, it is not clear how to marry models of speculation with business cycle theory; for another, the curse of dimensionality adds analytic and computational complexity. This paper does two things. First, it presents a model that lays out the necessary elements to marry speculation with real business cycle (RBC) theory in a way that qualitatively fits the narrative. Second, it presents a calibration

¹Interview with Thomas Sargent, *The Region*, August 26, 2010. Available at <https://www.minneapolisfed.org/article/2010/interview-with-thomas-sargent>.

²Interestingly, Hyman Minsky was Thomas Sargent’s undergraduate thesis advisor. Whereas Sargent departs methodologically and calls for a quantitative approach to economic research, there is agreement in the nature of business cycles.

exercise that gives a sense of the quantitative impact that speculation can have on business cycles, with minimal frictions. In this quantitative exercise, to circumvent the curse of dimensionality, we remove some model ingredients, necessary to generate the asset-price dynamics inherent in the narratives, without altering the business cycle properties of the model. We see these tasks as a first step toward having “well-confirmed quantitative versions” of speculation-driven business cycle theory. These tasks are important. Policy should not care about speculation and bubbles *per se*, but only inasmuch as these are a source of generalized downturns.

In a nutshell, we minimally modify a standard RBC model to fit the narrative of speculation-driven business cycles. As in any business cycle model, there is a representative worker and representative firm. Different from an RBC model, shares of the firm are bought and sold by a set of investors who agree to disagree about the evolution of the growth rate of total factor productivity (TFP). To this core, we introduce three essential ingredients, each producing a different effect, as depicted in Figure 1. The first ingredient is time-to-build, which is enough to produce fluctuations driven exclusively by investor sentiments, even in the absence of speculative bubbles. However, to fit the narrative, we need two additional ingredients. As shown by Harrison and Kreps (1978), to produce speculative bubbles, the second ingredient is short-selling constraints coupled with alternating degrees of optimism. The third ingredient is market segmentation. This ingredient produces bubbles that manifest in high P/E ratios, and not in high interest rates.

We introduce time-to-build as in the original Kydland and Prescott (1982) framework, which translates into a model where labor is hired one period in advance. Thus, labor is a form of investment. This assumption is important to produce economic fluctuations from changes in asset prices. It is also important that the firm’s investment is in hours and not in physical capital because it is well understood that capital investment cannot drive the cycle (Chari et al., 2007).³ Along those lines, we present a formula that connects the firm’s discount factor across states to the excess returns in the stock market. When excess returns are high, the firm’s discount factor is also high. The connection between asset prices and the business cycles emerge because the firm’s discount factor determines the investment in labor. This mechanism is grounded on evidence by Lustig and Verdelhan (2012), which finds that risk-adjusted excess returns are high in recessions.⁴

In the model, the firm’s discount factor is a function of investor sentiments. Here, investors make portfolio decisions; they borrow and lend at a risk-free rate and invest in shares.⁵ When more optimistic investors increase their ownership of firms, they become more representative in the shareholder pool. Their wealth, leverage, and relative optimism influence the firm’s hiring decisions through the firm’s discount factor. In this environment, the equity premium and labor volatility puzzles are intimately linked.

³Investment is small relative to the capital stock in a business cycle model. Thus, fluctuations in investment do not meaningfully impact the production possibility frontier.

⁴As the authors emphasize, due to higher risk-adjusted costs of capital during recessions, even unconstrained firms should invest and hire less. Recent work by Hall (2017) puts a similar mechanism to work in a Diamond-Mortensen-Pissarides framework. When discounts (or risk premiums) are high, firms invest less in creating jobs and, thus, unemployment increases. By imposing fluctuations in premiums that match the data, he shows the model can account for the bulk of fluctuations in unemployment.

⁵Since Mehra and Prescott (1985), the field of asset-pricing has found environments that can help explain the equity premium puzzle (Cochrane, 2017). In particular, Bansal and Yaron (2004) show that to produce volatile stochastic discount factors we need recursive preferences, long-run risk, and time-varying volatility. We follow Epstein and Zin (1991) and Tallarini (2000), and endow speculators with non-separable recursive preferences. We also allow for “long-run” risks, shocks to firm TFP growth—rather than the level. Speculators disagree about the distribution of TFP growth rates. We let heterogeneous beliefs be the source of time-varying volatility.

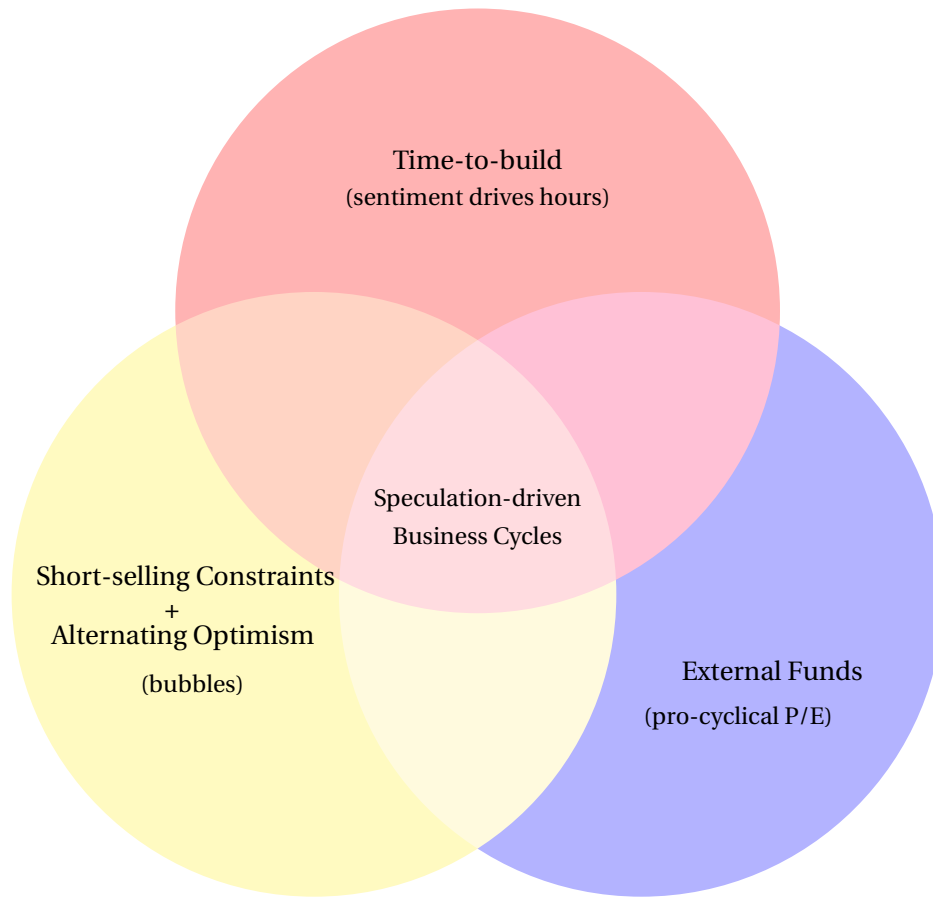


Figure 1: Diagram of model features.

With only time-to-build, belief heterogeneity can be enough to amplify fluctuations in hours. This amplification is due to counter cyclical movements in excess returns. However, on their own, these features do not produce asset-price dynamics that fit the narratives of speculation-driven business cycles. In principle, it is not obvious that fluctuations in excess returns will feed into fluctuations in P/E ratios; they can also be absorbed by the risk-free rate. For that, we show that a model needs a particular form of market segmentation.⁶ Namely, some agents (workers, in the present model) do not participate in the stock market, and their supply of funds must be sufficiently interest-rate elastic. This formulation of segmented markets is motivated by Guvenen (2009), for example.⁷ With this formulation of segmented markets, the pressure of speculation is reflected in movements in the P/E ratio, and not movements in the risk-free rate. This particular issue does not appear in Harrison and Kreps (1978) or Scheinkman and Xiong (2003) because the interest rate is exogenous in their models; thus, speculation shows in stock prices. In a business cycle model without segmentation, the interest rate absorbs the impact of speculation.

⁶We demonstrate this through a simple formula for the stock price, which holds for any specification of beliefs and degrees of risk-aversion, and for intertemporal elasticity of substitution approximately equal to one.

⁷Guvenen (2009) shows that this assumption, coupled with heterogeneous intertemporal elasticity of substitution among participants and non participants, renders the model consistent with several features of asset prices, including a high equity premium.

The final ingredient of the model is short-selling constraints coupled with alternating degrees of optimism. As highlighted by Harrison and Kreps (1978), this is needed to produce speculative asset-price bubbles. With this ingredient, the model is complete to fit the narratives. Section 3 lays out the environment and Section 4 showcases the effects of these ingredients and why they are needed.

To calibrate the model and make quantitative statements, one must model beliefs parsimoniously. Beliefs must be calibrated to generate an alternating optimism, and this must be done in a way that can be disciplined with data. Toward that end, we build on the work of Gennaioli and Shleifer (2010). Concretely, we assume that speculators are either rational—who hold the correct beliefs—or “diagnostic.” Gennaioli, Shleifer, and a series of co-authors argue that diagnostic expectations explain a myriad of social phenomena.⁸ Here, when TFP shocks are persistent, diagnostic beliefs produce extrapolative behavior. Investors expect shocks to be more persistent than they actually are. As a result, diagnostic beliefs produce the alternating pattern of optimism that is core to the narrative. In particular, compared with rational investors, diagnostic investors are over optimistic in high growth states, but over-pessimistic in adverse states. Beyond this appeal, diagnostic beliefs are a convenient formulation for quantitative purposes because these can be summarized with a single parameter. We discuss the calibration of the belief process in Section 5.

After calibrating the model, we complement the analysis by studying the qualitative predictions of the model, via simulations in Section 6. The framework reproduces the qualitative patterns of the speculation-driven business cycle narrative. Waves of optimism amplify the business cycle as the optimistic investors lever to buy shares. As the high TFP growth state persists, optimists accumulate wealth. This increasing wave of overoptimism leads to a greater willingness to bear risk, which induces lower excess returns and thus a higher P/E ratio. Since firm beliefs reflect the average stockholder’s stochastic discount factors, the firm adds more workers. The presence of diagnostic investors produces deeper recessions and large turnover after the economy transitions from high to low TFP growth. This is because diagnostic investors accumulate wealth during booms, but once a recession hits, they become pessimists. However, diagnostics remain relatively wealthy since they accumulated wealth in prior high growth states. Thus, during downturns, diagnostics remain representative in the shareholder pool and inject pessimism into the firm’s discount factor. This lowers the desire to hire workers. Ultimately, this mechanism produces asymmetric real business cycles: the longer the boom, the more severe the bust.

We also compare simulations between economies that feature short-selling constraints and economies that do not. This showcases the importance of speculative behavior. Short-selling constraints increase employment across states. Intuitively, overall willingness to bear risk by the marginal buyer increases in all states due to the reselling option. We also compare simulations with and without the external supply of funds by workers. Consistent with the analysis, without the external supply of funds, P/E ratios are stable, and excess returns are mostly driven by changes in the real interest rate. However, the business cycle dynamics are quantitatively similar in both cases, whether we include or exclude the external supply of funds.

We lever the latter result to produce quantitative statements in Section 7. As explained above, a full-blown

⁸Related to this paper, Bordalo et al. (2018) and Bordalo et al. (2019a) show that diagnostic expectations can explain how agents forecast stock returns, and Bordalo et al. (2018) argue that they can generate credit cycles.

quantitative exercise must overcome the curse of dimensionality. However, since the supply of external funds plays only a role in producing the cyclicalities of the P/E ratio but does not alter the business cycle properties, we can study how beliefs amplify the cycle in a version that shuts down the external fund to investors. This allows us to make progress toward quantitative questions because it eliminates one state variable from the model.

In the quantitative analysis, we show that the amplitude of the cycle of hours is amplified by speculation by a factor of at most 1.50. For a calibration of diagnostic beliefs, based on Bordalo et al. (2019a), the amplification factor is 1.15. That calibration is also consistent with a stock market bubble component of 20 percent relative to its fundamental value. We interpret this quantitative result as a sizable direct effect due to speculation, that would be amplified should other frictions, such as sticky prices or fire-sales externalities, were accounted for. We discuss how this paper fits in the literature in the next section, and then proceed to the main body.

2 Literature review

This paper is related to the literatures on real business cycles and the equity premium puzzle originated in Kydland and Prescott (1982) and Mehra and Prescott (1985), respectively. As in our paper, a recent strand links fluctuations in risk premiums to real business cycles. Di Tella and Hall (2019) stress the role of uninsurable idiosyncratic risk and precautionary savings, whereas Hall (2017), Borovička and Borovičková (2019) and Kehoe et al. (2019) study unemployment fluctuations in the context of the Diamond-Mortensen-Pissarides search model. Our paper also fits into the recent macro-finance literature that emphasizes the importance of the wealth share of special individuals (e.g., financial intermediaries) for the business cycle. For example, He and Krishnamurthy (2011), Brunnermeier and Sannikov (2014), Mendo (2018), and Silva (2019), among others. In our case, the wealth share of diagnostic investors is key. Relatedly, Caballero and Simsek (2019) show how financial trading between optimistic and pessimistic investors, by affecting the evolution of the distribution of wealth among them, amplifies a recession generated by a decline in risky asset valuations when output is determined by aggregate demand.

Our paper is also related to the natural selection literature, which asks whether those agents with incorrect beliefs eventually disappear. Blume and Easley (1992, 2006) and Sandroni (2000) argue that only those with more accurate beliefs survive in the long run in an environment with complete markets and separable preferences. However, this result is not robust to the market structure, as shown by Beker and Chattopadhyay (2010), Blume et al. (2018) and Cao (2018), and also not robust to preferences that are non-separable recursive even when markets are complete, as shown recently by Dindo (2019) and Borovička (n.d.). Closely related is Cao (2018), who works out the same investor problem as ours but does not link beliefs to TFP shocks in an RBC economy. In fact, the paper studies the natural selection hypothesis in an endowment economy with incomplete markets.⁹

Regarding the literature on heterogeneous beliefs and speculative behavior, we borrow the key ingredients from Harrison and Kreps (1978) and Scheinkman and Xiong (2003). The interaction with financial markets is explored

⁹Below we confirm the natural selection hypothesis in an example with separable preferences and without short-selling constraints. In addition, despite these recent contributions, in all simulations reported in the paper, rational investors eventually accumulate the entire stock of investors' wealth.

by Geanakoplos (2003, 2010), Fostel and Geanakoplos (2008), Simsek (2013), Iachan et al. (2019), among others.

Other papers have studied different transmission mechanisms of speculative behavior and bubbles to the real sector. In Gilchrist et al. (2005), monopolistic firms can overcome short-selling by issuing shares at a price above fundamental value, which lowers the cost of capital and enhances investment. Bolton et al. (2006) present an agency model in which over investment occurs during a bubble episode due to stock-based executive optimal compensation contracts that emphasize short-term stock performance. In contrast, Panageas (2005) shows that once investment subject to quadratic costs is introduced in a model with heterogeneous beliefs and a short-selling constraint, despite the speculative behavior of agents, the neoclassical q theory of investment remains valid. Related to our work, Buss et al. (2016) study policy implications in a quantitative framework in which agents trade for both risk-sharing and speculative reasons, and speculation reduces investment and welfare as it pushes the cost of capital up. As opposed to our work, these papers focus on investment in capital rather than hours, and not all of them feature models that are amenable to quantitative exercises.

By assuming that some agents hold diagnostic beliefs in the spirit of Gennaioli and Shleifer (2010), this paper is also related to the literature that explores how subjective beliefs affect the business cycle. See, for example, Eusepi and Preston (2011), Angeletos et al. (2018), Bordalo et al. (2018), and Bhandari et al. (2019), among others. Relatedly, Adam and Merkel (2019) show that (homogeneous) extrapolative beliefs can explain the stock price and business cycles altogether. Both cycles are connected as high stock prices signal profitable investment opportunities to capital producers. In contrast to our work, these papers abstract from speculative behavior.

Finally, a large literature studies other types of bubbles that emerge for reasons other than heterogeneous beliefs, such as the so-called “rational bubbles” (Blanchard and Watson, 1982; Santos and Woodford, 1997). Similar to our purpose, Martin and Ventura (2012) and Miao and Wang (2018) provide environments in which the collapse of rational bubbles leads to a recession. Other recent contributions emphasize the interaction of rational bubbles and policy, for example, Galí (2014), Hirano et al. (2015), Allen et al. (2018), and Asriyan et al. (2019). We leave the study of the role of policy in versions of our speculation-driven business cycle framework for future research.

3 Environment

Consider an infinite-horizon closed economy set in discrete time ($t = 0, 1, \dots$). We introduce investors into a standard RBC model with a representative worker and a representative firm. Investors (or potential speculators) differ in beliefs regarding the evolution of the growth of TFP and may hold (or issue) risk-free bonds and hold (or short-sell) risky shares of the firm. Workers do not hold stocks. The differences of beliefs induce the desire to lever and may introduce speculative portfolios, in the spirit of Harrison and Kreps (1978) and Scheinkman and Xiong (2003). In addition, we assume the firm hires labor one period in advance. This links portfolio decisions and labor fluctuations through the valuation of the firm.

Investors. The economy is populated by a finite number of infinite-lived investors, indexed by $i \in \{1, \dots, I\}$, and with corresponding masses μ_i . Investor i derives utility from the flow of consumption $c_{i,t}$. In particular, we

adopt Epstein-Zin recursive preferences:

$$U_{i,t} = c_{i,t}^{1-\beta} \left(\mathbb{E}_{i,t} \left[U_{i,t+1}^{1-\gamma} \right] \right)^{\frac{\beta}{1-\gamma}}, \quad (1)$$

where $\beta \in (0, 1)$ is the discount factor and $\gamma \geq 0$ is the risk-aversion parameter. The coefficient associated with the intertemporal elasticity of substitution (IES) in the Epstein-Zin formulation is set to 1 to obtain analytic expressions.¹⁰

Heterogeneity regards beliefs. Thus, expectations about future states, $\mathbb{E}_{i,t}$, are indexed by the agent identity. In particular, investor i forms beliefs $\{p_{ss'}^i\}$ regarding the TFP growth, g_t , which takes value in $\{\bar{g}_1, \dots, \bar{g}_S\}$. TFP growth may move from state s to s' , and is assumed to follow a Markov process with S states. Differences in beliefs regarding TFP growth of the representative firm translate into differences in beliefs about its future profits and stock returns, which creates a motive for trade in the financial market.

Investor i chooses consumption $c_{i,t}$, shares of the representative firm $n_{i,t+1}$ and risk-free bonds $b_{i,t+1}$ to maximize (1) subject to the borrowing constraint,

$$c_{i,t} + q_t n_{i,t+1} + b_{i,t+1} = (q_t + \pi_t) n_{i,t} + R_t b_{i,t}.$$

Here, q_t is the price per share,¹¹ and π_t are the profits of the representative firm. Investors can increase their leverage by issuing bonds. To get speculative behavior, we study versions of the model that differ in the extent of short-selling constraints. But in general, both $b_{i,t+1}$ and $n_{i,t+1}$ can take negative values. Finally, R_t is the risk-free rate that accrues to bonds bought in period $t - 1$ and carried over period t .

Within this framework, investors who are overly optimistic about future TFP prospects will buy more shares and tend to issue bonds. The more pessimistic tend to save by holding bonds that yield the risk-free rate. As we explain below, differences in beliefs determine ownership, which in turn defines the representative investor of the firm. This determines the firm's discount factor, which in turn produces labor-market fluctuations.

The worker. The representative worker derives utility from consumption $c_{w,t}$ but disutility from labor hours h_t . We assume preferences are GHH. In particular, the worker chooses $c_{w,t} \geq 0$ and $h_t \geq 0$ to maximize

$$\mathbb{E}_{w,0} \left[\sum_{t=0}^{\infty} \beta^t u \left(c_{w,t} - \xi A_{t-1} \frac{h_t^{1+\nu}}{1+\nu} \right) \right], \quad \text{subject to} \quad c_{w,t} + B_{t+1} = w_t h_t + R_t B_t,$$

where $u' > 0$ and $u'' \leq 0$. $\xi > 0$ is the scale factor on labor disutility, $\nu > 0$ is the inverse of the Frisch elasticity, and $\beta \in (0, 1)$ is the discount factor. Finally, w_t is the wage rate, and B_{t+1} is the worker's savings.

¹⁰To obtain these preferences, take the limit $\rho \rightarrow 1$ of the more standard formula in Epstein and Zin (1991),

$$U_{i,t} = \left\{ (1 - \beta) c_{i,t}^{1-1/\rho} + \beta \mathbb{E}_{i,t} \left(\left[U_{i,t+1}^{1-\gamma} \right] \right)^{\frac{1-1/\rho}{1-\gamma}} \right\}^{\frac{1}{1-1/\rho}},$$

where ρ is the IES if the model is deterministic.

¹¹The total amount of shares are normalized to 1.

The worker's budget constraint features an implicit assumption. Asset markets are segmented. In particular, the representative worker cannot hold shares.¹² This assumption implies that the worker's savings B_{t+1} is a variable that determines an important feature of the model. This variable, as we explain below, determines how the impact of speculation is absorbed by the cost of capital. In particular, it determines whether speculation shows up in q_t or R_{t+1} . To make the arguments as simple as possible, we do not take a stance on how the worker makes consumption-saving decisions. In principle, next-period savings (or the supply of funds schedule) can be any function of the state variables, as long as it follows a balanced growth path. Hence, although our numerical simulations assume rule-of-thumb behavior, all analytical results derived in what follows are consistent with optimal behavior.

The disutility of labor supply also merits some discussion. The scaling factor A_{t-1} is necessary to guarantee the existence of a balanced growth path as in Jaimovich and Rebelo (2009). It can be interpreted as a long-run wealth effect. Also, note that since the firm hires labor one period in advance, then h_t is the labor hired at $t - 1$ but only supplied at t , when the worker experiences the disutility of working. In principle, the worker's perceived stochastic process of TFP growth, $\{p_{ss'}^w\}$, used to form expectations $\mathbb{E}_{w,t}$, can be different from the true process and from those perceived by the other agents in the economy.

The firm. The representative competitive firm hires labor h_{t+1} that will be employed at $t + 1$ one period in advance (before the realization of the shock) to maximize expected profits,

$$\mathbb{E}_{f,t} [A_t g_{t+1} h_{t+1}^\alpha - w_{t+1} h_{t+1}],$$

where w_{t+1} is the wage rate paid at $t + 1$. Given an initial level A_0 , TFP evolves according to $A_t = A_{t-1} g_t$, where $g_t \in \{\bar{g}_1 < \bar{g}_2 < \dots < \bar{g}_S\}$ follows an S -state Markov process with transition probabilities $\{p_{ss'}\}$. The current state s and future state s' take values in $\{1, \dots, S\}$. Qualitative results do not change if we assume disagreement regarding the evolution of the TFP level rather than its trend but assume the latter improves the behavior of risk premiums (Bansal and Yaron, 2004). We assume the firm uses its own beliefs $\{p_{t,ss'}^f\}$ regarding the evolution of TFP growth to form expectations $\mathbb{E}_{f,t}$. The firm's beliefs reflect ownership. Given that labor is chosen one period in advance, probabilities are adjusted to account for the stockholders' stochastic discount factors. Akin to Hall (2017), fluctuations in risk premiums generate fluctuations in investment in hours. We spell out the precise formula for the firm's beliefs below, after we introduce the recursive version of the model.

3.1 Market clearing

Given the consistent notation for labor h_t in both the firm's and worker's optimization problems, we already impose market clearing in the labor market. By taking the first order conditions (FOCs) with respect to h_t in both problems,

¹²We could relax this assumption by assuming a less extreme form of market segmentation. The worker, for instance, could hold some shares, as long as some costs prevent the adjustment of the portfolio immediately.

and equalizing supply and demand in the labor market,¹³ one obtains,

$$h_{t+1} = \left(\frac{\alpha \mathbb{E}_{f,t}[g_{t+1}]}{\xi} \right)^{\frac{1}{1+\nu-\alpha}} \quad \text{and} \quad \frac{w_{t+1}}{A_t} = \xi \left(\frac{\alpha \mathbb{E}_{f,t}[g_{t+1}]}{\xi} \right)^{\frac{\nu}{1+\nu-\alpha}}. \quad (2)$$

Hence, realized profits at $t + 1$ are given by

$$\frac{\pi_{t+1}}{A_t} = \left(\frac{1}{\xi} \right)^{\frac{\alpha}{1+\nu-\alpha}} (\mathbb{E}_{f,t}[g_{t+1}])^{\frac{1+\nu}{1+\nu-\alpha}} \left\{ \frac{g_{t+1}}{\mathbb{E}_{f,t}[g_{t+1}]} \alpha^{\frac{\alpha}{1+\nu-\alpha}} - \alpha^{\frac{1+\nu}{1+\nu-\alpha}} \right\}.$$

To close the model, we specify the remaining market clearing conditions. Market clearing for goods requires that

$$\sum_{i=1}^I \mu_i c_{i,t} + c_{w,t} = A_{t-1} g_t h_t^\alpha,$$

whereas market clearing for bonds and shares requires that

$$B_{t+1} = - \sum_{i=1}^I \mu_i b_{i,t+1} \quad \text{and} \quad \sum_{i=1}^I \mu_i n_{i,t+1} = 1,$$

respectively. Note that the total amount of shares are normalized to 1. The definition of the equilibrium is standard. The model features a balanced growth path, with all variables growing at the same rate except labor hours h_t , the risk-free interest rate R_{t+1} and share holdings $n_{i,t+1}$. Let $\hat{x}_t = x_t/A_{t-1}$ denote a generic detrended variable.

4 Analysis

The model features one exogenous state variable $s \in \{1, \dots, S\}$, which indexes the growth in TFP. Denote the current aggregate endogenous state variables by X , to be defined below, and the future ones by X' . Let the law of motion of X be given by a transition function ψ , such that $X' = \psi(X, s, s')$. Finally, the recursive formulation for the firm's detrended profits is the following,

$$\hat{\pi}(X, s, s') = \frac{1}{\xi^{\frac{\alpha}{1+\nu-\alpha}}} \left\{ \bar{g}_{s'} \left(\alpha \mathbb{E}_{X,s}^f[g] \right)^{\frac{\alpha}{1+\nu-\alpha}} - \left(\alpha \mathbb{E}_{X,s}^f[g] \right)^{\frac{1+\nu}{1+\nu-\alpha}} \right\},$$

which depends on X only indirectly through the firm's beliefs, which reflect the firm's ownership as well as the shareholders' stochastic discount factors, to be specified below. The next subsections present the recursive representation of the model and some additional results.

¹³The labor supply schedule and the labor demand schedule are $w_{t+1} = \xi A_t h_{t+1}^\nu$ and $w_{t+1} = \alpha A_t \mathbb{E}_{f,t}[g_{t+1}]/h_{t+1}^{1-\alpha}$, respectively.

4.1 Investor problem

Recall that detrended variables are expressed as $\hat{x}_t = x_t/A_{t-1}$. To solve the investors' problem, we perform a change of variables. First, define the individual state to be

$$\hat{a}_{i,t} = (\hat{q}_t + \hat{\pi}_t)n_{i,t} + R_t \hat{b}_{i,t},$$

which is the investor i 's detrended wealth in the current period. Hence, the budget constraint can be rewritten as

$$\hat{c}_{i,t} + \hat{q}_t n_{i,t+1} + g_t \hat{b}_{i,t+1} = \hat{a}_{i,t}.$$

Second, after omitting the subscript t and using the prime for next-period variables, define $\tilde{c}_i = \hat{c}_i/\hat{a}_i = c_i/a_i$ as the consumption share of wealth; $\tilde{n}_i = \hat{q}n'_i/\hat{a}_i(1 - \tilde{c}_i)$ as the share of invested wealth, i.e. after deducting consumption, that is invested in risky shares; and, analogously, $\tilde{b}_i = g\hat{b}'_i/\hat{a}_i(1 - \tilde{c}_i)$ is the share of invested wealth that goes to risk-free bonds. Hence, the budget constraint can be written as $\tilde{n}_i + \tilde{b}_i = 1$.

Let the returns on shares be denoted by

$$R_n(X, s, s') = \left[\frac{\hat{q}(\psi(X, s, s'), s') + \hat{\pi}(X, s, s')}{\hat{q}(X, s)} \right] \bar{g}_s,$$

and the risk-free rate on bonds bought today and carried until tomorrow by $R_b(X, s)$. Next period detrended wealth, after substituting $\tilde{n}_i = 1 - \tilde{b}_i$, is given by

$$\hat{a}'_i = \frac{1}{\bar{g}_s} \left[R_n(X, s, s')(1 - \tilde{b}_i) + R_b(X, s)\tilde{b}_i \right] \hat{a}_i(1 - \tilde{c}_i). \quad (3)$$

Let the subscript -1 denote last-period variables. The optimization problem of investors can be written as

$$U_i(\hat{a}_i, X, s; A_{-1}) = \max_{\tilde{c}_i, \tilde{b}_i} \left\{ (\tilde{c}_i \hat{a}_i A_{-1})^{1-\beta} \left(\mathbb{E}_{i,s} [U_i(\hat{a}'_i, X', s'; A)^{1-\gamma}] \right)^{\frac{\beta}{1-\gamma}} \right\},$$

subject to (3). For now, we consider the version without short-selling constraints. We follow a guess-and-verify method to characterize the solution. Conjecture that $U_i(\hat{a}_i, X, s; A_{-1}) = V_i(X, s)\hat{a}_i A_{-1}$. We state the following lemma, whose proof is presented in Appendix A.1.

Lemma 1. *Given the law of motion $X' = \psi(X, s, s')$ and prices $R_b(X, s)$ and $R_n(X, s, s')$ for all X, s, s' , the optimal consumption share for investor i is $\tilde{c}_i(X, s) = 1 - \beta$, and the optimal portfolio weight $\tilde{b}_i(X, s)$ is defined implicitly by*

$$\mathbb{E}_{i,s} \left[\frac{(V_i(X', s'))^{1-\gamma} [R_b(X, s) - R_n(X, s, s')]}{\left[R_n(X, s, s')(1 - \tilde{b}_i(X, s)) + R_b(X, s)\tilde{b}_i(X, s) \right]^\gamma} \right] = 0. \quad (4)$$

In addition, V_i satisfies the following recursion,

$$\ln V_i(X, s) = (1 - \beta) \ln \tilde{c}_i(X, s) + \beta \ln(1 - \tilde{c}_i(X, s)) + \frac{\beta}{1 - \gamma} \times \\ \ln \left(\mathbb{E}_{i,s} \left[\left(V_i(X', s') \left[R_n(X, s, s')(1 - \tilde{b}_i(X, s)) + R_b(X, s)\tilde{b}_i(X, s) \right] \right)^{1-\gamma} \right] \right).$$

Firm beliefs. Given the characterization of the investor problem, we are now in a position to define the firm's beliefs, $\{p_{ss'}^f(X)\}$, which reflect ownership and stochastic discount factors. To spare notation, let $v_i(X, s, s') = V_i(\psi(X, s, s'), s')$. An inspection of FOC (4) reveals that the stochastic discount factor of the speculator reads

$$SDF_i(X, s, s') = \frac{(v_i(X, s, s'))^{1-\gamma}}{\left[R_n(X, s, s')(1 - \tilde{b}_i(X, s)) + R_b(X, s)\tilde{b}_i(X, s) \right]^\gamma}. \quad (5)$$

Invoking equation (3) and the fact that $\hat{c}_i = (1 - \beta)\hat{a}_i$, the stochastic discount factor reads $v_{i,t+1}^{1-\gamma} \beta^\gamma \left(\frac{g_t \hat{c}_{i,t+1}}{\hat{c}_{i,t}} \right)^{-\gamma}$, a more familiar representation.

We let the firm's beliefs be given by investors' risk-adjusted beliefs averaged across them. In particular,

$$p_{ss'}^f(X) = \sum_i \omega_i(X, s) \left[\frac{SDF_i(X, s, s') p_{ss'}^i}{\sum_{s'} SDF_i(X, s, s') p_{ss'}^i} \right], \quad (6)$$

where $\left\{ \frac{SDF_i(X, s, s') p_{ss'}^i}{\sum_{s'} SDF_i(X, s, s') p_{ss'}^i} \right\}_{s'}$ are the risk-neutral transition probabilities (or risk-adjusted beliefs) of investor i regarding all possible future states, and $\{\omega_i(X, s)\}_i$ are weights (summing one) that reflect ownership of the firm by investors.

Risk-neutral probabilities are simply beliefs adjusted for stochastic discount factors. The idea is that the firm maximizes on behalf of its shareholders, and internalizes that employment decisions induce profit risk due to the assumption that they must be made one period in advance. Since the interaction of profit risk and portfolio decisions affects investors' well-being, the firm's beliefs must be adjusted in a way that maximizes shareholders' preferences. Risk-neutral probabilities accomplish this goal. Importantly, portfolio and employment decisions become closely connected, tying fluctuations in stochastic discount factors to fluctuations in labor hours.

Risk-adjusted beliefs across investors are aggregated using weights $\{\omega_i(X, s)\}_i$. For now, we do not take a stance on the weights given to investors. One possible example is to assume that weights are given by the proportions of shares. In this case,

$$\omega_i(X, s) = \frac{\mathbf{I}_{\{\tilde{n}_i(X, s) > 0\}} \eta_i \tilde{n}_i(X, s)}{\sum_i \mathbf{I}_{\{\tilde{n}_i(X, s) > 0\}} \eta_i \tilde{n}_i(X, s)},$$

where $\tilde{n}_i(X, s) = 1 - \tilde{b}_i(X, s)$ is the proportion of wealth after deducting consumption that is invested in risky shares, and $\eta_i = \frac{\mu_i \hat{a}_i}{\sum_i \mu_i \hat{a}_i}$ is the wealth share of speculator i .

Specialization: $S = 2$ and $\gamma = 1$. In Appendix B, we work out FOCs (4) and SDFs (5) when $S = 2$, say $s \in \{L, H\}$ with $\bar{g}_L < \bar{g}_H$. $S = 2$ is a natural benchmark. Absent any bounds on portfolio choices as assumed in this subsection,

it renders market completeness due to the presence of two assets (risk-free bonds and risky shares) that can be used to transfer consumption through time and across states.

For $\gamma \neq 1$, the solution makes explicit the dependence of $\tilde{b}_i(X, s)$ and $SDF_i(X, s, s')$ on $v_i(X, s, s') = V_i(\psi(X, s, s'), s')$, for which we do not have a closed form solution. When $\gamma = 1$, this dependence on $v_i(X, s, s')$ is eliminated, which enhances tractability. The quasi-closed forms in the following lemma are obtained.

Lemma 2. *Let $\gamma = 1$ and $S = 2$ with $s \in \{L, H\}$, then the optimal portfolio decision is*

$$\tilde{b}_i(X, s) = p_{sL}^i \frac{R_n(X, s, H)}{R_n(X, s, H) - R_b(X, s)} - p_{sH}^i \frac{R_n(X, s, L)}{R_b(X, s) - R_n(X, s, L)}.$$

In addition, individual wealth evolves according to

$$\frac{\bar{g}_s \hat{a}'_i(X, s, s')}{\beta \hat{a}_i} = \frac{1}{SDF_i(X, s, s')} = p_{ss'}^i R_b(X, s) \left[\frac{R_n(X, s, H) - R_n(X, s, L)}{|R_b(X, s) - R_n(X, s, -s')|} \right].$$

Hence, for any configuration of weights $\omega_i(X, s)$, the firm's probabilities adjusted for investors' stochastic discount factors are given by

$$p_{ss'}^f(X) = \frac{|R_b(X, s) - R(X, s, -s')|}{R_n(X, s, H) - R_n(X, s, L)}. \quad (7)$$

The result is immediate once we set $\gamma = 1$ in equations (12) and (13) in Appendix B, where we characterize the evolution of wealth and the investor's portfolio for the case when $\gamma = 1$ (log preferences) and $S = 2$. In that specific case, the model is solved almost analytically. By multiplying SDF_i and $p_{ss'}^i$ to compute shareholder i 's risk-adjusted beliefs in (6), $p_{ss'}^i$ in the numerator and in the denominator cancel out. This reveals that risk-adjusted beliefs are the same for all investors and are an explicit function of excess returns. In turn, excess returns reflect the degree of belief heterogeneity in the economy.

This result has two implications. First, we do not need to take a stance on how adjusted beliefs are aggregated to form the firm's beliefs. As long as weights $\{\omega_i(X, s)\}$ sum to one, they can be anything.

Second, the simple formula in (7) connects two branches of modern macroeconomics. To the extent that the firm's beliefs depend on excess returns, it ties the equity premium puzzle to the real business cycle, in particular labor market fluctuations. Intuitively, when excess returns are low (good state), and therefore investors are willing to bear risk, probabilities are distorted to favor the good state, and more labor hours are employed. Analogously, when excess returns are high, the firm employs less. In that sense, the inability of a standard real business cycle model to generate both large fluctuations in hours (without relying on a large labor supply elasticity) and large volatility in the equity premium can be interpreted as a single puzzle.

Given that the FOCs with respect to the portfolio decisions hold with equality, subjective valuations of (perceived) flow of dividends of all speculators must coincide and, hence, must equal the price per share. Indeed, if someone's subjective valuation differs from the price per share, there would be an arbitrage opportunity. As Miller (1977) and Harrison and Kreps (1978) emphasize, the presence of short-selling constraints, by preventing those willing to short-sell from arbitraging the market, generates subjective valuations that differ from the price

per share. This leads to the possibility of speculative bubbles, something we explore in this paper.

Toward that end, we study what occurs when we impose a margin constraint on shares, $\tilde{b}_i \leq 1 + \kappa$. The coefficient $\kappa \geq 0$ is introduced to avoid excessive short-selling. In the extreme case of no short-selling, the constraint $n_{i,t+1} \geq 0$ must be satisfied, which is equivalent to $\tilde{b}_i \leq 1$. When $\kappa > 0$, then $q_t n_{i,t+1} \geq -\kappa \beta a_{i,t}$ is equivalent to $\tilde{b}_i \leq 1 + \kappa$, meaning that wealthier investors can short-sell proportionally more. If this upper bound is binding, $\tilde{b}_i(X, s) = 1 + \kappa$, the evolution of individual detrended wealth in (3) reads

$$\hat{a}'_i(X, s, s') = \frac{\beta}{g_s} [(1 + \kappa)R_b(X, s) - \kappa R_n(X, s, s')] \hat{a}_i.$$

4.2 State variables

Next we introduce the endogenous aggregate state variables that determine X . One natural candidate is the whole distribution of wealth: $R\hat{B}$ for the worker,¹⁴ and $\{\hat{a}_i\}_{i=1}^I$ for investors. We can encode the wealth distribution among investors by keeping track of aggregate wealth $\sum_{i=1}^I \mu_i \hat{a}_i$ coupled with the wealth shares for $I - 1$ agents, i.e. $\{\eta_i\}_{i=1}^{I-1}$, where $\eta_i = \mu_i \hat{a}_i / \sum_{i=1}^I \mu_i \hat{a}_i$.

Nonetheless, although R is predetermined, $\sum_{i=1}^I \mu_i \hat{a}_i$ depends on price \hat{q} which is an equilibrium object. To circumvent this problem, and get a predetermined state variable, we use the definition of individual wealth, the market clearing conditions for bonds and shares, and the government budget constraint, to get

$$\sum_i \mu_i \hat{a}_i = \hat{q}(X, s) + \underbrace{\bar{g}_s h^\alpha - \hat{w}h}_{=\hat{E}} - R\hat{B}. \quad (8)$$

Note that \hat{E} and $R\hat{B}$ are predetermined. Hence, $X \equiv \{R\hat{B}, \hat{E}, \{\eta_i\}_{i=1}^{I-1}\}$ is the set of aggregate endogenous state variables. Both \hat{E} and $R\hat{B}$ have simple interpretations. The former is detrended realized profits (dividends received by shareholders), whereas the latter is the aggregate previous detrended debt the speculators need to honor or the worker's current wealth.

4.3 Equilibrium

The next proposition derives an equilibrium relationship between the price per share, current profits, and “net liquidity” available to speculators. Fix $\hat{B}'(X, s)$ which is the worker's detrended aggregate savings carried out from today to tomorrow. We take it as a partial equilibrium object, albeit a useful one.

Proposition 1. *The price per share is given by*

$$\hat{q}(X, s) = \frac{1}{1 - \beta} \left[\beta \hat{E} + \bar{g}_s \hat{B}'(X, s) - \beta R\hat{B} \right]. \quad (9)$$

¹⁴As opposed to R and B , which are the risk-free interest rate R on bonds B issued yesterday, we coupled the subscript b and the superscript prime to denote the risk-free interest rate $R_b(X, s)$ on bonds $B'(X, s)$ issued today, respectively. Finally, by detrending the worker's budget constraint, one obtains $\hat{c}_{w,t} + g_t \hat{B}_{t+1} = \hat{w}_t h_t + R_t \hat{B}_t$.

The proof is found in Appendix A.2. Equation (9) is a quasi-closed form for the price per share. It is derived using all equilibrium conditions in the model, except the FOC (4) with respect to the portfolio choices of the investors. This equation is fairly general, holding for any value of the risk aversion parameter γ , any configuration of beliefs $\{p_{ss'}^i\}$ including homogeneous beliefs, and any supply of funds schedule $\hat{B}(X, s)$ including the optimal one. It also holds if bounds, such as debt and/or short-selling constraints, are imposed on portfolio decisions. Hence, it is consistent with speculative behavior. Critical in the equation, however, is that the IES equal to one anchors the marginal propensity to consume out of wealth to $1 - \beta$. In more general formulations, the marginal propensity to consume is time varying but remarkably stable—hence the equity premium puzzle. We lever this result to argue that equation (9) captures the main quantitative forces that drive stock price dynamics.

In particular, equation (9) directly links the price per share $\hat{q}(X, s)$ to discounted profits, $\beta \hat{E}$, and a measure of flow of funds (or “net” liquidity available) to investors, $\bar{g}_s \hat{B}'(X, s) - \beta R \hat{B}$. Both expressed in present value, given the $1 - \beta$ in the denominator. If the worker is hand-to-mouth such that $\hat{B}'(X, s) = \hat{B} = 0$, then the P/E ratio, $\frac{\hat{q}(X, s)}{\hat{E}} = \frac{\beta}{1 - \beta}$, is constant. This is true even in the presence of the ingredients that make the economy prone to speculative bubbles. Hence, aside from the dynamic effects on the stock price through the evolution of profits, the interest rate $R_b(X, s)$ absorbs the effects of speculation. A bubble component, meaning that current stock price $q(X, s)$ is above all subjective valuations (based on the investors’ perceived flow of dividends), may arise because future dividends are discounted at a higher interest rate.

To obtain a procyclical P/E ratio, the flow of funds to investors relative to current profits, $\frac{\bar{g}_s \hat{B}'(X, s) - \beta R \hat{B}}{\hat{E}}$, must be higher at the good states, which of course, depends on how $\hat{B}'(X, s)$ is specified, although there is a force through growth \bar{g}_s pushing funds up at good states. As a final extreme example, if liquidity supply is unlimited at a given interest rate R , speculation only affects $\hat{q}(X, s)$. This emphasizes the importance of market clearing conditions, as well as the net liquidity supply available to investors. To sum up, given (X, s) , if “appetite” for shares (in fixed supply) due to speculation is high, the extent to which $\hat{q}(X, s)$ or $R_b(X, s)$ reflects such “appetite” depends on the flow of funds to speculators.¹⁵

Bubble component. In the presence of a bidding short-selling constraint, and alternating optimism and pessimism among speculators, the economy is prone to bubbles. Following Scheinkman and Xiong (2003), we define the bubble component encoded in the price per share as the log-difference of the price per share and the largest subjective valuation,

$$\log \hat{q}(X, s) - \log(\max_i \hat{q}_i(X, s)),$$

where the subjective valuation $\hat{q}_i(X, s)$ satisfies the following recursion:

$$\hat{q}_i(X, s) = \sum_{s'} p_{ss'}^i SDF_i(X, s, s') (\hat{q}_i(\psi(X, s, s'), s') + \hat{\pi}(X, s, s')) \bar{g}_s.$$

In words, the bubble component is precisely the resale option value of the marginal buyer of shares.

¹⁵In Appendix C, we describe the remaining equations that characterize the equilibrium, and outline an algorithm to solve the model numerically.

Specialization: $S = 2$ and $\gamma = 1$. Recall that, absent any bounds on portfolio choices, $S = 2$ implies complete markets. In addition, $\gamma = 1$ makes the model highly tractable. If $S = 2$, $\gamma = 1$, and there are no short-selling constraints ($\kappa = \infty$), we get fairly simple formulas for portfolio shares and the evolution of wealth, as Lemma 2 highlights. In addition, we also get a simple expression for the evolution of wealth shares $\{\eta_i\}$, with a precise implication for its limit distribution.

Lemma 3. *If $\gamma = 1$, $S = 2$ with $s \in \{L, H\}$, and $\kappa = \infty$, then the evolution of the individual investor j 's wealth share is governed by*

$$\eta_j'(X, s, s') = \frac{p_{ss'}^j \eta_j}{\sum_{i=1}^I p_{ss'}^i \eta_i}, \quad \text{which implies} \quad \frac{\eta_j'(X, s, s')}{\eta_k'(X, s, s')} = \frac{p_{ss'}^j \eta_j}{p_{ss'}^k \eta_k} \quad \text{for all } j, k. \quad (10)$$

In addition, if one investor is rational, say speculator $i = 1$ such that $p_{ss'}^1 = p_{ss'}$ for all s, s' , then $\eta_1 \rightarrow 1$ almost surely.

A proof is found in Appendix A.3. Regarding the evolution of the wealth distribution, upon the realization of s' , the larger the ratio $p_{ss'}^j/p_{ss'}^k$, the higher the increase in relative wealth, η_j/η_k . Intuitively, investors who believed that the realized state s' was more likely to realize ex-ante also chose portfolios that perform better in such a state.

Regarding the limit result, only rational investors (i.e., those with correct beliefs) survive. In other words, they eventually acquire the whole investors' stock of wealth. This is a well-known result in a complete markets context with nonrecursive preferences. See, for example, Sandroni (2000) and Blume and Easley (2006). In the next section, we assume a rational investor exists. But if none of the investors are rational, one can extend this result by following the steps in Blume and Easley (2006), and show that investors with the closest beliefs to the truth survive.

Finally, the limit result is not general to other values (other than one) of the risk aversion parameter γ , as shown by the recent contribution of Borovička (n.d.). Also, the result does not generalize to any incomplete markets structure in the presence of a debt-limit constraint, even when $\gamma = 1$, as Beker and Chattopadhyay (2010) and others show in related contexts.

4.4 Discussion: model assumptions and properties

The model we have discussed so far is tailored to speak to many features of the speculation-driven business cycles. This can be easily seen for the specialization in which $\gamma = 1$, $S = 2$ with $s \in \{L, H\}$, and $\kappa = \infty$. As the good state $s = H$ persists, the law of motion for wealth shares, equation (10), implies that optimistic investors accumulate wealth on average. Hence, asset prices reflect this increasing overoptimistic view. As in a standard macro-finance model, this wave of overoptimism leads to more willingness to bear risk, and this induces lower excess returns. The transmission from lower excess returns to higher hours worked can be seen through adjusted firm's beliefs in equation (7). A great willingness to bear risk not only affects excess returns, but also distorts the firm's beliefs in favor of good states. Whether lower excess returns reflect a higher risk-free interest rate or higher stock prices (i.e., lower stock returns) depends on the supply of liquidity to investors. As equation (9) highlights, if the supply is high relative to profits during good times, the economy displays high P/E ratio, credit expansion, and increased leverage, commonly associated with speculative episodes. The counterpart of this argument as the bad states persist also

holds. As pessimistic investors accumulate wealth, the economy displays higher excess returns and lower labor hours. Aiming to generate speculation and bubbles in a parsimonious way, the next section takes a stand on the heterogeneity of beliefs as a final step to make the model speak comprehensively regarding the speculation-driven business cycle view.

When presenting the paper, we prioritize parsimony and computational simplicity over realism and computational complexity. At this point, it is worth discussing some assumptions and model properties, and the extent to which alternative modeling assumptions would modify the qualitative or quantitative results.

TFP growth. By contrast to most RBC models, technological shocks affect growth rates rather than levels. This is a proxy for growth prospects in the economy. Nonetheless, the same conclusions would emerge if we had assumed TFP shocks at the level in an economy without trend growth. Indeed, by reinterpreting g_t as levels, it is enough to substitute \bar{g}_s for one in all equations above. We opt to keep growth shocks for two reasons. First, they are important to make stochastic discount factors more volatile and thus, generate somewhat larger risk premiums. Second, as equation (9) highlights, \bar{g}_s pushes the net flow of funds to speculators up at good states, something the model requires to generate a procyclical P/E ratio.

Degenerate ergodic distribution. In the model with $\gamma = 1$, $S = 2$ and $\kappa = \infty$, the share of rational wealth converges asymptotically to one. The same pattern holds in all quantitative simulations below. Hence, by construction, the model does not generate speculation-driven business cycles in its ergodic distribution. For that reason, in what follows, we focus our quantitative analysis along transition paths. A common easy-to-implement technical assumption is to allow exogenous exit/entry—death/birth—among speculators. This additional assumption is known to introduce mixing into the Markov chain that governs the dynamics of the wealth share. With this mixing, the Markov chain of the wealth share can feature a non-degenerate ergodic distribution. Given that we find a slow speed of convergence toward the steady-state without mixing, we do not expect our amplification results to change in a quantitatively meaningful way if we were to have obtained them from a steady-state version of the model with mixing. For parsimony, we do not introduce deaths/births into the model.

Risk-aversion, number of shocks, and the equity premium puzzle. We presented the model in the full generality of Epstein-Zin preferences and for any arbitrary number of shocks, but later specialized to a risk-aversion of one and only two shocks to obtain analytic expressions. Whereas a higher value for risk-aversion, or a richer shock structure—that could admit disaster risk, for example—would help quantitatively in producing a greater risk-premium, the upper bound of the amplification that can be produced by speculation depends only on actual, and not risk-adjusted probabilities—this is something we show later. Thus, the upper bound on amplification holds for any configuration of beliefs and risk aversion, and thus, is independent of the equity premium in the model. In turn, a greater equity premium could help increase the extent to which a calibrated model of beliefs can reach that upper bound. For the calibration we present later, where the risk-aversion coefficient is set to one, the model goes a long way toward the upper bound on amplification. We decided to work with a risk-aversion of one to see how far the model goes without a high level of risk aversion, while keeping computations simple. We acknowledge, however, that hitting a reasonable equity premium would increase amplification. In a nutshell, our

quantitative results below on amplification should be seen as a lower bound relative to versions of the model that produce an equity premium closer to the one observed in the data.

Exogenous supply of funds. We adopted an exogenous supply of funds by workers. In reality, this can represent supply of funds from agents outside of the economy. Importantly, all theoretical results are consistent with an endogenous supply of funds obtained from the worker's dynamic optimization problem. Although conceptually, it is not difficult to introduce such endogenous supply of funds, computationally, this adds an additional state variable. For computational reasons, we shut down external funds in the quantitative section. Later, we argue that although shutting down external savings affects the extent to which the stock price and the interest rate would absorb the pressure from speculation, the overall quantitative conclusions regarding amplification should not change much.

5 Modeling beliefs

So far, we have been agnostic about beliefs, but in order to progress toward quantitative statements, we need a model of beliefs. Hence, the last step to formalize the speculation-driven business cycle view is to impose some discipline on the heterogeneity of beliefs. We take a stance on the number of participants and assume $I = 2$ types of investors. We let $i = 1$ represent the rational investor who holds correct beliefs, $p_{ss'}^1 = p_{ss'}$ for all s, s' . We call $i = 2$ the diagnostic investor.

Diagnostic expectations are formalized by Gennaioli and Shleifer (2010, 2018) based on prior work by Daniel Kahneman and Amos Tversky, and applied by them and co-authors to explain a wide range of social phenomena. Suppose an agent wants to form beliefs regarding the distribution of types (next-period shocks in our case) in a given group (current shock s in our case). Then a specific type (say s') is diagnostic or representative of this group if its true probability of realization ($p_{ss'}$) is large relative to its true probability of realization in some reference group (for example, $p_{ks'}$, such that the reference group is a shock k other than s). Diagnostic expectations attribute more weights to diagnostic types.

We follow Bordalo et al. (2018), who formalize diagnostic expectations in the context of an AR(1) process. As a reference group, they consider past conditions as if no news were received in the meantime. If past conditions mean one period with no news, the probability of realization of s' in such reference group is $\sum_k p_{s_{-1}k} p_{ks'}$, which comes at a cost of tracking one more state variable, the previous shock s_{-1} . This extra state variable would be avoided if the reference group were described by conditions in the long past, such that the probability of s' sense is $\sum_k \bar{p}_k p_{ks'}$, where $\{\bar{p}_s\}$ is the invariant distribution associated with $\{p_{ss'}\}$. In particular, $i = 2$'s diagnostic beliefs are given by:

$$p_{ss'}^2 = p_{ss'} \left(\frac{p_{ss'}}{\sum_k \bar{p}_k p_{ks'}} \right)^\theta Z_s, \quad (11)$$

where Z_s is a constant that guarantees that $\sum_{s'} p_{ss'}^2 = 1$, and $\theta > 0$ measures the extent to which beliefs are distorted.

Hence, given that $\theta > 0$, beliefs are distorted by attributing more probability towards types (shocks) that are

more diagnostic. Psychologically, diagnostic types are oversampled from limited and selective memory. Important for our purposes is the relevance of diagnostic expectations to explain how individuals forecast stock market returns (Gennaioli et al., 2016; Bordalo et al., 2019a).

If $S = 2$ with $s \in \{L, H\}$, a simple algebra reveals that $p_{LL}^2 > p_{LL}$ and $p_{HH}^2 > p_{HH}$ if and only if $p_{LL}p_{HH} > (1 - p_{LL})(1 - p_{HH})$. In words, whenever shocks are persistent, diagnostic investors believe that states are more persistent than they really are. This is the case for TFP shocks as observed in the data.

This simple and intuitive result has two implications for the speculation-driven business cycle view we evaluate in this paper. First, the diagnostic investor is optimistic at the good state but pessimistic at the bad state. Under the presence of a binding short-selling constraint, this observation allows the possibility of bubbles, as noted by Harrison and Kreps (1978). Indeed, at the good state, the price per share not only reflects the diagnostic's valuation but also the option to resell the share to the rational agent if the bad state is realized. Similarly, at the bad state, a bubble may arise because the relatively optimistic rational agent is willing to pay a price higher than the valuation due to the resell option. As becomes clear during speculation episodes, with the presence of diagnostic speculators and short-selling constraints, the turnover of shares is amplified, and a bubble component arises.

Second, the diagnostic agent accumulates proportionally more wealth as the states (whether good or bad) persist, whereas the rational agent accumulates wealth when the states change. This observation implies that the longer the boom within a cycle, the larger is the drop in employment once the bust arrives. Indeed, as the good state persists, the optimistic diagnostic investor accumulates an increasing amount of wealth, which also means an increasing amount of hours worked due to the aforementioned transmission mechanism. Once the bad state happens, the diagnostic becomes pessimistic. The larger the wealth previously accumulated, the larger is the fall in hours employed. Indeed, a wealthier pessimistic investor at the bad state implies fewer worker hours. We illustrate this and other implications of the model in the next section.

6 Speculation-driven business cycles

This section illustrates the mechanics of the model through simulation exercises. In a benchmark case, we assume there are no external funds to speculators, $\hat{B}'(X, s) = \hat{B} = 0$. Then, we show how the supply of funds changes the financial moments implied by the model, but does not alter the transmission mechanism and the amplification of shocks. Also, to highlight the role of speculation and bubbles, in both instances we present results with and without short-selling constraints.¹⁶

No external funds to investors. We begin with a description of the results when there is no flow of funds to investors, $\hat{B}'(X, s) = \hat{B} = 0$.¹⁷ Figures 2 and 3 compare simulations for three economies: the first economy features homogeneous investors and the other two, heterogeneous investors; in one economy investors are allowed to

¹⁶As discussed above, with the short-selling constraint, \tilde{b}_i is bounded above by $1 + \kappa$, where $\kappa = 0$ implies no short-selling at all. Except for Proposition 1, the analytical expressions obtained are not generalized when short-selling constraints bind for some speculators. However, as discussed in Appendix C, we can easily implement these constraints numerically.

¹⁷In Appendix C.1, we discuss how this assumption coupled with $\gamma = 1$, $S = 2$, and $I = 2$ simplifies the solution and the numerical implementation.

short-sell ($\kappa = \infty$) and, thus, bubbles cannot arise in this case; in the others short-selling constraints are present ($\kappa = 0$), which generates a bubble. In all the exercises that follow, the economy is initiated at a steady state with only rational investors, $\eta_1 = 1$ —which is an absorbing state. Also, TFP growth is initiated in the good state. For the economies with belief heterogeneity, at $t = 1$, a mass of agents that hold 20 percent of the wealth unexpectedly becomes diagnostic. We simulate these economies for eight periods of expansions, followed by four consecutive periods of recession. After that the economy recovers fully.

Figure 2 plots the evolution of the rational’s wealth share and of (log) hours worked. Figure 3 shows the evolution of the financial market outcomes (excess returns, firm’s adjusted probabilities, share’s price, risk-free rate and the bubble component of the stock price).

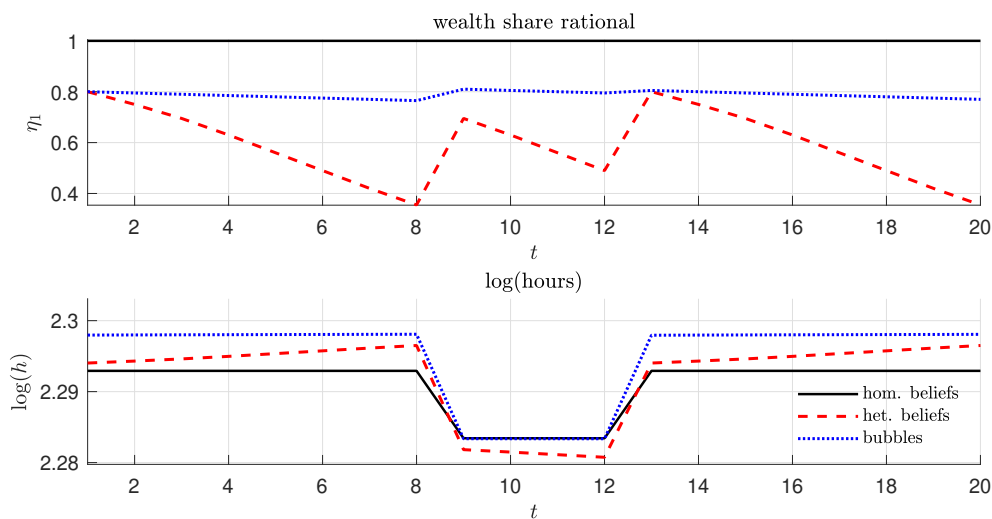


Figure 2: Typical business cycle.

Note: The figure is computed using a symmetric Markov chain, with $p_{ss'} = 0.70$ if $s' = s$, and $\theta = 2$, such that $p_{ss'}^2 = 0.927$ if $s' = s$. In addition, $\beta = 0.99$, $\nu = 0.5$, $\xi = 1$, $\gamma = 1$, $\alpha = 2/3$, $\bar{g}_L = 1$ and $\bar{g}_H = 1.02$.

First, compare the economy with homogeneous beliefs (solid line) with the economy with heterogeneous beliefs without short-selling constraints (dashed line). It is clear that the economy with heterogeneous beliefs features a cycle with greater persistence and amplitude. The explanation is that, as the expansion persists, increasing overoptimism dominates the economy. With a greater overall willingness to bear risk, excess returns decline during the expansion. As a consequence, the firm’s risk-adjusted probability of a transition to a recession falls and, thus, employment increases throughout the cycle. Once the recession is realized, the diagnostic investors become pessimistic. Because diagnostics are relatively pessimist during recessions and because they hold a non-negligible share of total wealth, employment falls to a level below its counterpart in an economy with only rational investors. As the recession persists, the diagnostic investors further accumulate wealth, implying an increasing pessimism of the representative investor. This is reflected in a path of increasing excess returns and declining hours.

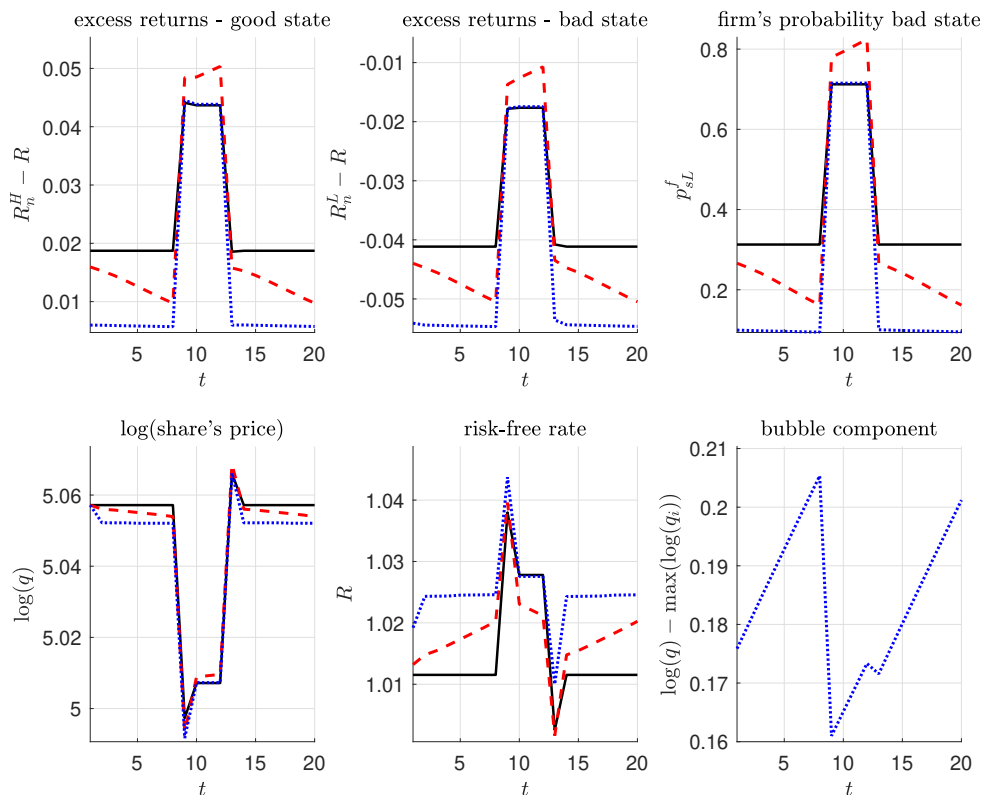


Figure 3: Financial variables along a typical cycle.

Note: Parameter values are the same as those used in Figure 2.

When comparing the results between the economy with a short-selling constraint (dotted lines) and the economy without (dashed line), two conclusions emerge. First, the rational's wealth becomes a slower moving variable and fluctuates less once the constraint is present. Thus, conditional on being in an expansion, risk premiums, adjusted probabilities, and hours employed move more slowly. Second, with short-selling constraints, the firm employs more hours irrespective of the state. Intuitively, once the short-selling constraint is imposed, the overall willingness to bear risk by the marginal buyer increases in both states—due to the bubble component that captures the resell option. In particular, in recessions, when the rational speculator is the marginal buyer, outcomes become closer to the economy with only the rational investor. Similarly, in expansions, outcomes reflect the more optimistic beliefs of the diagnostic speculator, the marginal buyer in booms. Hence, excess returns are depressed across states in the economy with short-selling constraints. Lower excess returns translate into higher firm risk-adjusted probabilities of future high TFP growth and, thus, more hours employed.

Regarding the financial variables, excess returns are qualitatively aligned with the data. They decline in expansions and increase in recessions. Also, the model captures the boom-bust narrative of bubbles. However, without an elastic supply of funds by workers, the P/E ratio is constant due to general equilibrium forces—see equation (9). Also, as in a standard RBC model, we find a counter-cyclical risk-free rate, although the effect is mitigated by het-

erogeneous beliefs and short-selling constraints. As emphasized by Harrison and Kreps (1978) and Scheinkman and Xiong (2003), the presence of short-selling constraints produces an asset-price bubble. Yet, without external funds it is the risk-free rate that mostly absorbs the pressure from speculation. A bubble component arises not because the P/E ratio increases, but because subjective valuations (of the perceived flow of dividends) fall due to higher discounting.

As we emphasize throughout the paper, the supply of funds to speculators is crucial to get the full narrative right. In the next subsection, we simulate an economy with external funds that produces the right cyclicity of the P/E ratio. Nonetheless, it is important to reiterate that excess returns are the relevant financial moment for the pass-through to real activity. Despite missing the P/E ratio dynamics, as long as the supply of funds does not affect much excess returns, the business cycles are quantitatively similar with or without external funds.

Finally, to clarify the extent of amplification as a function of the length of a boom, Figure 4 reproduces two cycles for the economy with heterogeneous beliefs, but without short-selling constraints. The comparison is now between two expansions, one that lasts eight periods (dot-dashed line) and the original four-period expansion (dashed line). There is a noticeable pattern. The longer the boom, the more severe the bust. This is an immediate implication of the extra additional time for diagnostic investors to accumulate wealth. The takeaway is that cycles are asymmetric, although we calibrate the Markov chain symmetrically.

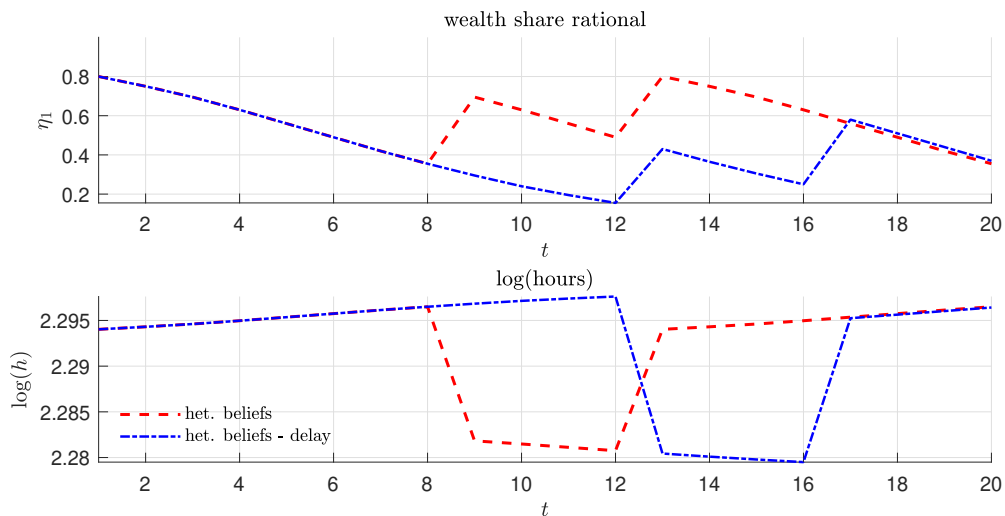


Figure 4: Responses for cycles of different length.

Note: Parameter values are the same as those used in Figure 2.

Flow of funds to investors. To illustrate how the liquidity flow to investors fixes the financial moments, we consider a rule-of-thumb supply of funds from the worker,

$$\hat{B}'(X, s) = \phi_s(\hat{w}h + R\hat{B}),$$

where ϕ_s is the state-dependent propensity to accumulate out of current income plus current wealth, $\hat{w}h + R\hat{B}$.¹⁸ Recall that $\hat{w}_t = \xi h_t^\nu$, and note that both the wage bill $\xi h_t^{1+\nu}$ and labor h_t are predetermined variables that can be easily recovered from predetermined profits \hat{E} . By combining the equation above with equation (9), one obtains

$$\hat{q}(X, s) = \frac{1}{1-\beta} \left[\beta \hat{E} + \phi_s \xi h_t^{1+\nu} + (\phi_s - \beta) R \hat{B} \right].$$

This expression showcases that as the boom persists higher wages and higher flow of funds to the speculators boost the stock price. In addition, $\phi_H > \phi_L$ enhances the chances of getting a procyclical P/E ratio.¹⁹ Note that the supply of funds we employ here does not depend on the risk-free interest rate $R_b(X, s)$, which is countercyclical as in standard RBC models. Hence, if savings were very sensitive to $R_b(X, s)$, there would be a counteracting force depressing the stock price during booms. As long as this substitution effect is not too strong, the rule-of-thumb supply of funds should approximate fairly well a properly derived counterpart. We do not pursue a full derivation of the worker supply of funds due to tractability and for ease of exposition.

Figures 5 and 6 are analogues of Figures 2 and 3 but for economies with external supply of funds. In this case, we report the evolution of the P/E ratio instead of the bubble component.

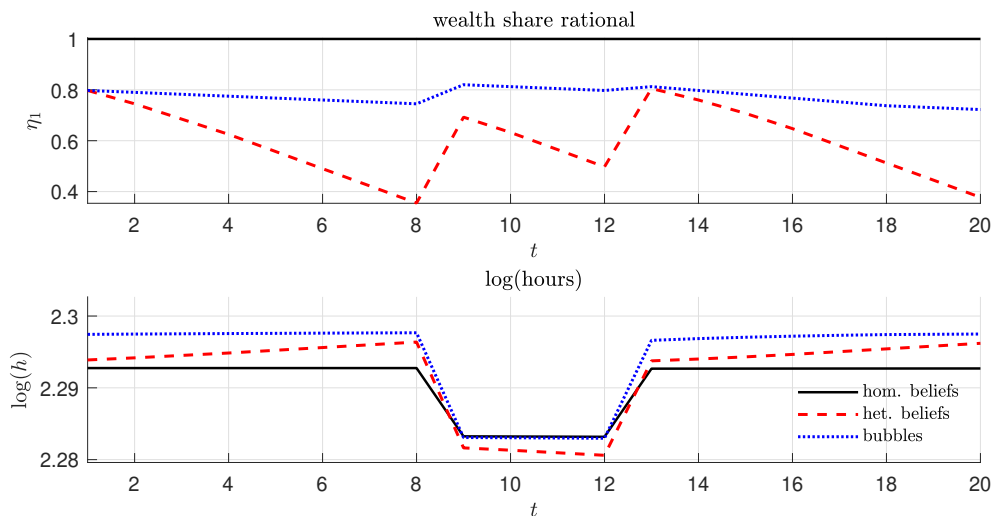


Figure 5: Typical business cycle with elastic supply of funds.

Note: Parameter values are the same as those used in Figure 2. Additionally, we set $\phi_H = \beta = 0.99$ and $\phi_L = 0.988$. These figures imply that the worker holds 23.7 percent of the total wealth in the steady state with $\eta_1 = 1$.

The figures show how an external supply of funds can fix, at least qualitatively, the financial moments. In this case, the price per share increases (declines) in expansions (recessions). Not only does the P/E ratio become procyclical, but heterogeneous beliefs and speculation increase its amplitude. Nonetheless, the business cycle properties remain very similar, regardless of the presence of an external supply of funds. We build on this last

¹⁸In Appendix C.2, we discuss how to solve the model numerically for this case.

¹⁹This formulation of the supply of funds can be justified by a precautionary motive as the worker faces incomplete markets, or by consumption smoothness as in Guvenen (2009), who assumes that non-stockholders have a lower IES than stockholders.

result to study the role of speculation on business cycles in the version without external funds. We know this version misses some of the financial aspects of the theory, but it is likely to produce business cycle properties similar to the version with external funds.

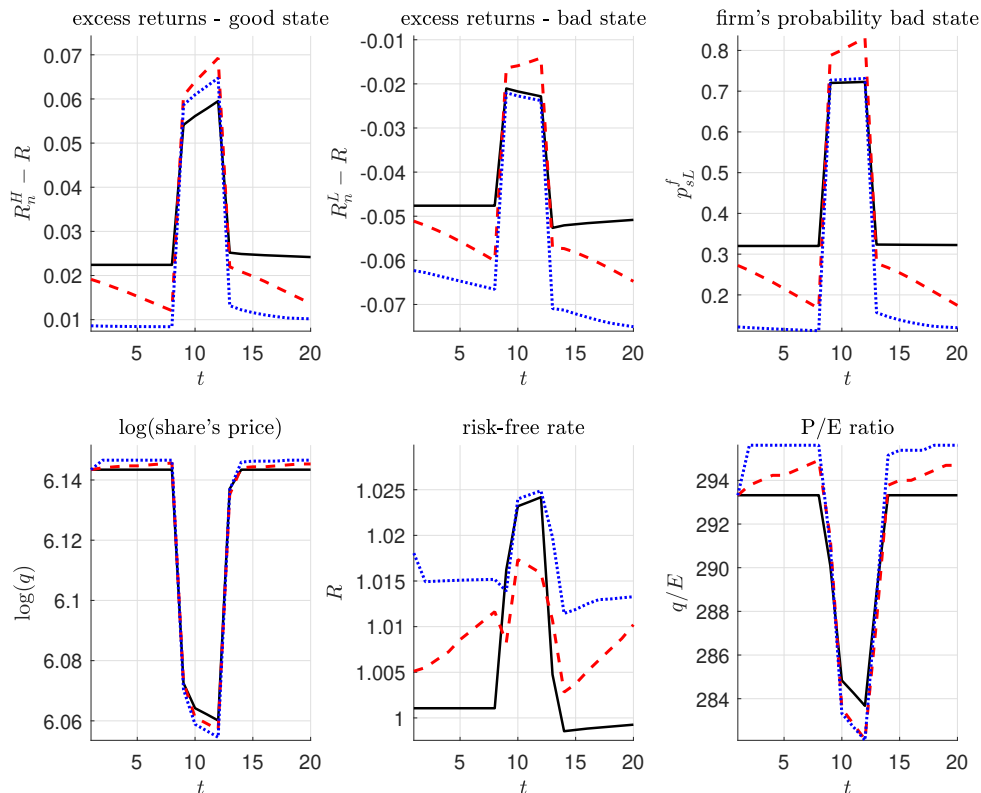


Figure 6: Financial market variables with elastic supply of funds.

Note: Parameter values are the same as those used in Figure 2.

7 Quantitative exploration

In this section we conduct a quantitative analysis of the potential that speculation can have to magnify business cycle fluctuations, along the lines of the narratives of Minsky, Kindleberger, and Shiller. Because we abstract from financial frictions and pecuniary externalities as in Dávila and Korinek (2017), or sticky prices and aggregate demand externalities as in Caballero and Simsek (2019), if anything, this quantitative exploration should be seen as a lower bound. In this section, $\hat{B}'(X, s) = \hat{B} = 0$, so as anticipated, we miss the procyclicality of the P/E ratio, but arguably obtain similar quantitative business cycle implications.

Calibration. The calibration is standard. We set the model period to quarters, so $\beta = 0.991$. Regarding the production function, we assume a labor share of $\alpha = 2/3$. The good state is associated with quarterly growth of 1.2 percent, $g_H = 1.012$, where the bad state with a recession of -0.4 percent, $g_L = 0.996$. The average duration

of a recession (expansion) is four (ten) quarters, implying that $p_{LL} = 0.75$ ($p_{HH} = 0.90$). These figures are in congruence with rough calculations using the NBER recession dates, or with more elaborated estimations such as Hamilton (1989). The Frisch elasticity is set to 2, meaning that $\nu = 0.5$. We set the relative risk aversion to $\gamma = 1$ (log preferences). ξ is just a scale parameter, normalized to one. These parameters are fixed throughout all specifications, and summarized in Table 1.

Parameter	Meaning	Value	Reference
β	discount factor	0.991	average real rate of 4%
α	decreasing returns	2/3	labor share of 2/3
g_H	high-state growth	1.012	quarterly growth of 1.2% in booms
g_L	low-state growth	0.996	quarterly growth of -0.4% in recessions
p_{HH}	high-to-high transition prob.	0.90	Hamilton (1989)
p_{LL}	low-to-low transition prob.	0.75	Hamilton (1989)
ν	inverse labor-supply elasticity	0.5	standard RBC
ξ	labor scaling factor	1	normalization
γ	risk-aversion	1	tractability

Table 1: Calibration table.

An upper bound for amplification. Consider the baseline model with time-to-build in which beliefs are homogeneous and rational. For the aforementioned calibration, firm's adjusted probabilities are $p_{LL}^f = 0.759$ and $p_{HH}^f = 0.896$, implying that the standard deviation of (log) hours is 0.57 percent, and hours decline by 1.25 percent upon a transition from high to low TFP growth states. In this section we aim to compute how much these figures are amplified due to heterogeneous beliefs and speculation.

Initially, for any model of beliefs, we compute the maximal amplifications of the standard deviation of (log) hours and the decline in hours once the recession hits. These upper bounds are given by a version of the model without time-to-build. To see why, notice that equation (2) reflects that labor is employed taking into consideration *expected* TFP growth one period ahead, and not the *realized* TFP growth. Since TFP growth is a mean-reverting process, expected TFP growth is lower than actual TFP growth in booms and higher in recessions. Thus, time-to-build *per se* limits the amplitude of the cycle. The fluctuations in a model without time-to-build are an upper bound to fluctuations in a model with time-to-build for any configuration of beliefs. Absent time-to-build, hours decline by 1.91 percent upon a transition from high to low TFP growth states. The overall standard deviation of (log) hours is 0.86 percent.²⁰ Heterogeneous beliefs and speculation cannot amplify business cycles beyond those levels.

Taking these observations into account, for any model of beliefs, the maximal amplifications are given by a factor of 1.51 in the standard deviation (0.86 in the model without time-to-build divided by 0.57 in the model

²⁰This quantity is in line with the findings of early real business cycle models, e.g. Cooley and Prescott (1995).

with time-to-build and homogeneous-rational beliefs), and a factor of 1.53 in the amplitude of hours decline (1.91 divided by 1.25). Indeed, the maximal amplifications are obtained when beliefs distort expectations and make perceived TFP growth as persistent as possible. This is because, the more persistent the beliefs, the closer the model with time-to-build is to a model without time-to-build, which is in turn an upper bound. Hence, the maximal cycle amplification due to heterogeneous beliefs and speculation is obtained when investor beliefs always reflect in firm risk-adjusted probabilities of $p_{LL}^f = p_{HH}^f = 1$, yielding the same business cycle properties of the model without time-to-build.

Amplification in the calibration. To gauge the role of heterogeneous beliefs in amplifying the cycle, we next discuss the amplification we obtain for various values of θ . As above, we present results with and without short-selling, $\kappa = \infty$ and $\kappa = 0$, respectively.²¹

Figure 7 depicts how p_{LL}^2 and p_{HH}^2 vary with θ . As explained above, these two distorted probabilities govern the amplitude of hours worked in that they approximate the model with belief heterogeneity to the upper-bound obtained in a model without time-to-build. Bordo et al. (2019a), Bordo et al. (2018) and Bordo et al. (2019b) estimate θ to be around 0.9, 0.9, and 0.6 in the context of forecasting stock returns, explaining credit cycles, and forecasting macroeconomic variables, respectively. Hence, moderate values of θ can distort substantially subjective probabilities of remaining in a recession. For instance, if $\theta = 0.4$, diagnostic investors believe that the economy will remain in a recession with a probability of 87 percent as opposed to the true probability of 75 percent. As θ further increases, such probability actually surpasses the subjective probability of remaining in a boom, even though the high TFP-growth state is more persistent. In words, conditional on being in the bad state, the bad state is very diagnostic, and thus subject to lots of disagreement given the presence of rational investors.

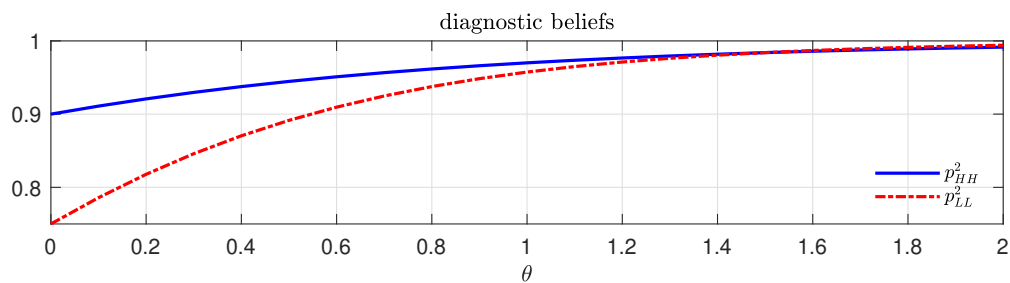


Figure 7: Diagnostic beliefs.

Note: Diagnostic beliefs for $\theta \in [0, 2]$ are computed using equation (11) in Section 5.

The top panels of Figure 8 showcase the amplification of the decline in hours when the economy moves from an expansion to a recession, for different values of θ . In particular, for each value of θ in the x axis, each of the top panels reports the effects of the amplification, relative to a model with homogeneous beliefs, when a recession occurs after two-and-a-half years of expansions—the average cycle encoded in the Markov chain of TFP growth. As before, each top panel reports the amplification, for different values of η_1 in the initial period, capturing differ-

²¹In all simulations when $\kappa = 0$, the economy always converges to the steady state in which $\eta_1 = 1$. Hence, short-selling constraints are not binding in the steady state.

ent waves of optimism.²² Dashed lines depict the economy without short-selling constraints ($\kappa = \infty$), and thus, without bubbles; in the case of dotted lines, the economy does not feature short-selling ($\kappa = 0$).

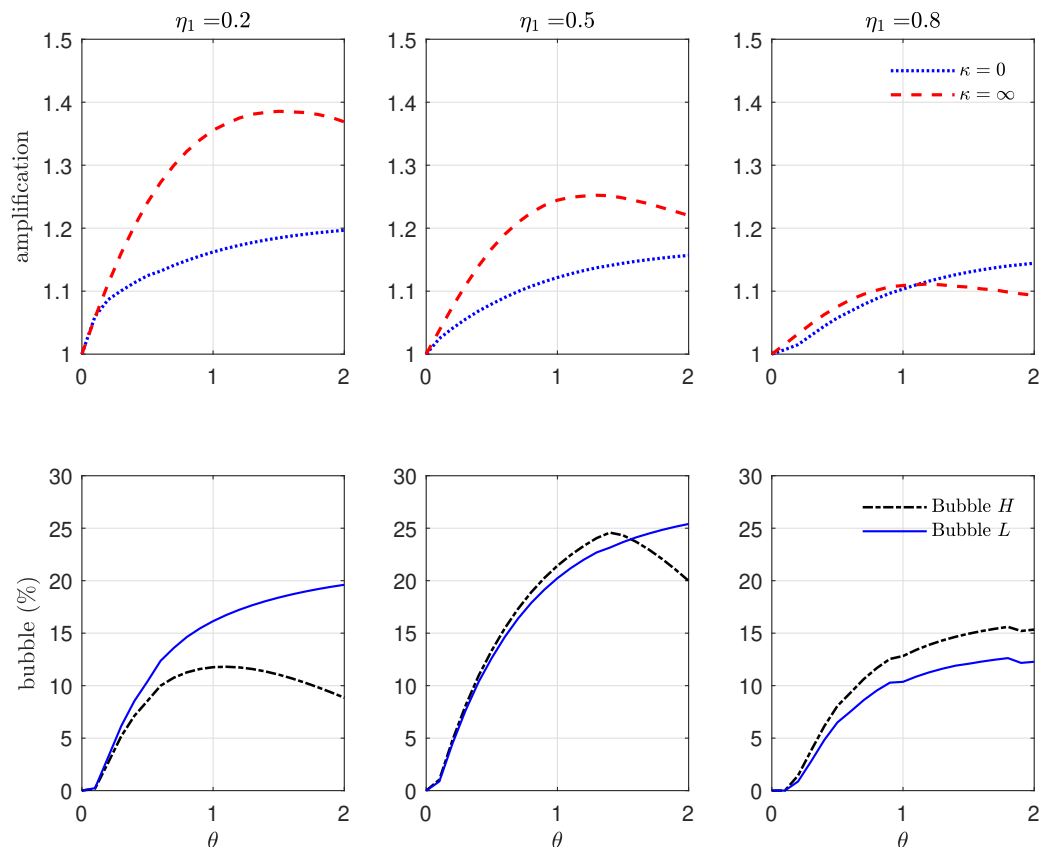


Figure 8: Amplification in labor responses and bubble components.

Note: Simulations start from the steady state reached during a boom when $\eta_1 = 1$. Columns correspond to belief shocks at $t = 1$ that bring the rational wealth share to 0.2, 0.5, and 0.8. A recession hits after two-and-a-half years of expansions. For $\theta \in [0, 2]$, and for $\kappa \in \{0, \infty\}$, top panels plot the amplification in labor decline relative to a model with homogeneous-rational beliefs. For $\theta \in [0, 2]$, when $\kappa = 0$, bottom panels plot bubble components right before the recession (Bubble H), and in its first year (Bubble L). The calibration is reported in Table 1.

The figure shows that when $\theta = 1$ and $\kappa = \infty$, for example, the decline in labor is amplified by 36, 24, and 11 percent when the economy is initiated at η_1 equal to 0.2, 0.5 and 0.8, respectively.²³ Given that the upper bound for amplification is about 50 percent for any configuration of beliefs, the calibration to diagnostic beliefs studied here goes a long way toward that upper bound.

A salient feature is that the role of short-selling constraints ($\kappa = 0$, dotted lines) is nonmonotone. Once tight short-selling constraints are imposed, the amplification may be higher or lower than the one in the economy with

²²For each path, we start from the steady state reached during a boom when $\eta_1 = 1$, and then shock the economy by making a fraction of agents diagnostic. The columns correspond to belief shocks that bring the rational wealth share to 0.2, 0.5, and 0.8, respectively. One can interpret each path as waves of optimism inherent to the speculation-driven business cycle narratives.

²³We consider $\theta = 1$ as the benchmark value for the diagnostic parameter. As the model does not feature a representative agent, we increase it above the upper end of the estimations in Bordalo et al. (2019a), Bordalo et al. (2018) and Bordalo et al. (2019b), due to the presence of rational agents in the model that might countervail the aggregate effects of diagnostic expectations.

loose constraints. Intuitively, bubbles depress excess returns and boost employment in both states. Hence, the overall effect is ambiguous because they may attenuate or exacerbate the labor decline once the recession hits.

The bottom panels of Figure 8 depict the bubble components before and after the bust—for the case of short-selling constraints, the only case that produces a bubble. In particular, the dashed lines represent the bubble prior to the bust, whereas the solid lines represent the bubble component in the first period of the recession. The x axis shows different values of θ . For $\theta = 1$, for instance, the bubble component ranges from 10 percent to 20 percent, depending on the initial value for the share of rational wealth, η_1 . Interestingly, depending on this initial value, the bubble component in recessions can actually exceed the bubble component in booms. This does not mean that the stock price is higher but only that the price is above its fundamental value. Thus, the model fits the narrative of “bursting bubbles” in recessions only for specific degrees of optimism and initial wealth. For example, for an initial wealth of $\eta_1 = 0.8$ and $\theta = 1$, the burst of the bubble is moderate, taking the bubble component only from 12.8 to 10.4 percent. In fact, for lower values of initial η_1 , say for $\eta_1 = 0.2$, beliefs disagreement actually mitigates the decline in the stock price due to the bubble component, as shown in the bottom left panel.

We also study the degree of amplification contributed by short-selling constraints. At the same time that short-selling constraints may mitigate the decline in hours in the first cycle after a wave in optimism, they can also generate greater unconditional volatility of hours. These results seem to contradict each other, but they are an artifact of a countervailing force. Short-selling constraints limit the extent of losses by diagnostic agents, prolonging their survival: diagnostic wealth becomes a slower moving variable in the presence of constraints.

To illustrate this point, we set $\theta = 1$ and simulate economies. We compute (a) the evolution of the average 20-year rolling-window standard deviation of (log) hours, and (b) the path of the average wealth share of rational agents. Results are reported in Figure 9. Each column of panels represents simulations initiated at η_1 equal to 0.2, 0.5, and 0.8. Dotted lines represent economies without short-selling ($\kappa = 0$), whereas dashed lines correspond to economies that allow for short-selling ($\kappa = \infty$). The top panels present the amplification in standard deviation of (log) hours, relative to an economy with homogeneous-rational beliefs, and bottom panels show the respective evolution of the rational wealth share.

The conclusions that emerge from these simulations are that the volatility of hours increases in both economies, with and without short-selling constraints, relative to the economy with homogeneous beliefs. Nonetheless, the long-run amplification is stronger in the presence of short-selling constraints. Again, this is because short-selling constraints make diagnostic agents survive longer, and this sustains a longer amplification of labor responses. In terms of unconditional moments, this translates into more volatile business cycles when short-selling constraints are allowed. The effects are sizable considering that there is an upper bound of a nearly 50 percent increase in hours volatility. Indeed, in the initial periods, the standard deviation of hours is amplified by 15, 12, to 8 percent for initial values of rational wealth share of 0.2, 0.5, and 0.8, respectively.

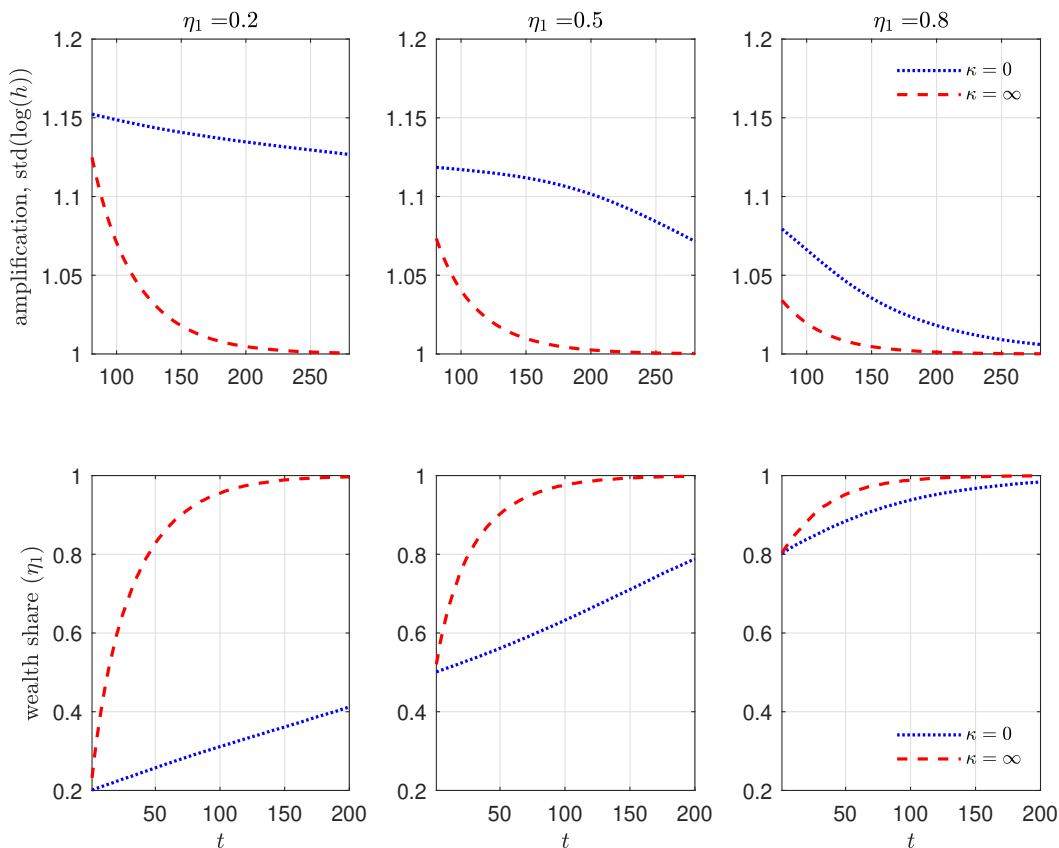


Figure 9: Amplification of the volatility of hours and wealth evolution.

Note: We simulate 10,000 economies that start from the steady state when $\eta_1 = 1$. Columns correspond to belief shocks at $t = 1$ that bring the rational wealth share to 0.2, 0.5, and 0.8. Top panels plot the evolution of the average (across economies) 20-year rolling-window standard deviation of (log) hours. Bottom panels plot the evolution of the average (across economies) wealth share of rational agents. We consider economies with $\kappa \in \{0, \infty\}$ and $\theta = 1$. The remaining calibration is reported in Table 1.

8 Conclusion

This paper adapts a standard real business cycle model to fit the speculation-driven business cycle narrative. We spell out some key ingredients that are sufficient to account for that narrative and present a quantitative evaluation of the narrative.

The paper presents an equilibrium equation that relates the price per share to current profits and the flow of funds to investors. Given that markets are segmented, the price per share should reflect fluctuations in the net supply of funds to the financial sector. This expression is valid for any configuration of beliefs, any value for risk aversion, and independent of the presence of short-selling (or borrowing) constraints. The price that absorbs speculation, whether the price per share or the risk-free interest rate, depends on the amount of funds supplied at good and bad states. In particular, a bubble arises even without a surge in the stock price.

This result speaks to the leaning against the wind debate—see Gourio et al. (2018) for a recent contribution.

Suppose the government can control the risk-free rate—this can be accomplished through fiscal, monetary, or macro-prudential policy. If heterogeneous beliefs and short-selling constraints are the driving forces behind bubbles, this result suggests that the government cannot fully “burst” a bubble, but it can decide to what extent the stock price or the risk-free rate should absorb speculation. Hence, policy must trade-off the stabilization of one price against the cost of making the other one fluctuate more. In a world with aggregate demand externalities or pecuniary externalities, the policymaker faces a trade-off.

With respect to the qualitative amplification of the business cycle, a few conclusions are worth emphasizing. First, as diagnostic investors accumulate wealth during booms and recessions, real business cycles are amplified. Second, the longer the boom period, the more severe is the bust. Although the modeling of workers is stylized, these results hold for any supply of funds schedule—including one appropriately derived from optimal behavior. Finally, if short-selling constraints are imposed, such that the economy becomes prone to bubbles, simulations suggest we get further amplification in booms but an attenuation during busts.

A quantitative exercise suggests that the effects of heterogeneous beliefs and speculation can be sizable relative to a homogeneous beliefs benchmark. The bubble component of asset prices can also be sizable. By construction, cycles cannot be amplified more than 50 percent in our quantitative exploration. A concrete calibration shows that the volatility of hours can be amplified up to 15 percent, and bubbles can reach up to 20 percent of the asset’s fundamental value. We see these quantitative results as sizable because they respond directly to heterogeneous beliefs and speculation. Arguably these could be amplified by account for other frictions, such as sticky prices or fire-sales externalities, accounted for.

References

- Adam, Klaus and Sebastian Merkel**, “Stock Price Cycles and Business Cycles,” 2019. Working Paper.
- Allen, Franklin, Gadi Barlevy, and Douglas Gale**, “A Theoretical Model of Leaning Against the Wind,” 2018. Working Paper.
- Angeletos, George-Marios, Fabrice Collard, and Harris Dellas**, “Quantifying Confidence,” *Econometrica*, 2018, 86 (5), 1689–1726.
- Asriyan, Vladimir, Luca Fornaro, Alberto Martin, and Jaume Ventura**, “Monetary Policy for a Bubbly World,” 2019. Working Paper.
- Bansal, Ravi and Amir Yaron**, “Risks for the Long Run: A Potential Resolution of Asset Pricing Puzzles,” *Journal of Finance*, 2004, 59 (4), 1481–1509.
- Beker, Pablo and Subir Chattopadhyay**, “Consumption Dynamics in General Equilibrium: A Characterisation when Markets Are Incomplete,” *Journal of Economic Theory*, 2010, 145 (6), 2133–2185.
- Bhandari, Anmol, Borovička, and Paul Ho**, “Survey Data and Subjective Beliefs in Business Cycle Models,” 2019. Working Paper.
- Blanchard, Olivier J. and Mark W. Watson**, “Bubbles, Rational Expectations and Financial Markets,” in Paul Wachtel, ed., *Crises in the Economic and Financial Structure*, Lexington, MA: D.C. Heath and Company, 1982, pp. 295–316.
- Blume, Lawrence and David Easley**, “Evolution and Market Behavior,” *Journal of Economic Theory*, 1992, 58 (1), 9–40.
- and —, “If You’re So Smart, Why Aren’t You Rich? Belief Selection in Complete and Incomplete Markets,” *Econometrica*, 2006, 74 (4), 929–966.
- Blume, Lawrence E., Timothy Cogley, David A. Easley, Thomas J. Sargent, and Viktor Tsyrennikov**, “A Case for Incomplete Markets,” *Journal of Economic Theory*, 2018, 178, 191–221.
- Bolton, Patrick, José Scheinkman, and Wei Xiong**, “Executive Compensation and Short-Termist Behaviour in Speculative Markets,” *Review of Economic Studies*, 07 2006, 73 (3), 577–610.
- Bordalo, Pedro, Nicola Gennaioli, and Andrei Shleifer**, “Diagnostic Expectations and Credit Cycles,” *Journal of Finance*, 2018, 73 (1), 199–227.
- , —, **Rafael La Porta, and Andrei Shleifer**, “Diagnostic Expectations and Stock Returns,” *Journal of Finance*, 2019, 74 (6), 2839–2874.

- , —, **Yueran Ma, and Andrei Shleifer**, “Over-reaction in Macroeconomic Expectations,” 2019. Working Paper.
- Borovička, Jaroslav**, “Survival and Long-Run Dynamics with Heterogeneous Beliefs under Recursive Preferences,” *Journal of Political Economy*, 128 (1), 206–251.
- and **Katarína Borovičková**, “Risk Premia and Unemployment Fluctuations,” 2019. Working Paper.
- Brunnermeier, Markus K. and Yuliy Sannikov**, “A Macroeconomic Model with a Financial Sector,” *American Economic Review*, 2014, 104 (2), 379–421.
- Buss, Adrian, Bernard Dumas, Raman Uppal, and Grigory Vilkov**, “The Intended and Unintended Consequences of Financial-market Regulations: A General-equilibrium Analysis,” *Journal of Monetary Economics*, 2016, 81, 25–43.
- Caballero, Ricardo J. and Alp Simsek**, “A Risk-centric Model of Demand Recessions and Speculation,” 2019. Working Paper.
- Cao, Dan**, “Speculation and Financial Wealth Distribution Under Belief Heterogeneity,” *Economic Journal*, 2018, 128 (614), 2258–2281.
- Chari, V. V., Patrick J. Kehoe, and Ellen R. McGrattan**, “Business Cycle Accounting,” *Econometrica*, 2007, 75 (3), 781–836.
- Cochrane, John H.**, “Macro-Finance,” *Review of Finance*, 2017, 21 (3), 945–985.
- Cooley, Thomas F. and Edward C. Prescott**, “Economic Growth and Business Cycles,” in Thomas F. Cooley, ed., *Frontiers of Business Cycle Research*, Princeton University Press Princeton, 1995, pp. 1–38.
- Dávila, Eduardo and Anton Korinek**, “Pecuniary Externalities in Economies with Financial Frictions,” *Review of Economic Studies*, 02 2017, 85 (1), 352–395.
- Di Tella, Sebastian and Robert E. Hall**, “Risk Premium Shocks Can Create Inefficient Recessions,” 2019. Working Paper.
- Dindo, Pietro**, “Survival in Speculative Markets,” *Journal of Economic Theory*, 2019, 181, 1–43.
- Epstein, Larry G. and Stanley E. Zin**, “Substitution, Risk Aversion, and the Temporal Behavior of Consumption and Asset Returns: An Empirical Analysis,” *Journal of Political Economy*, 1991, 99 (2), 263–286.
- Eusepi, Stefano and Bruce Preston**, “Expectations, Learning, and Business Cycle Fluctuations,” *American Economic Review*, 2011, 101 (6), 2844–72.
- Fostel, Ana and John Geanakoplos**, “Leverage Cycles and the Anxious Economy,” *American Economic Review*, 2008, 98 (4), 1211–44.

- Galí, Jordi**, “Monetary Policy and Rational Asset Price Bubbles,” *American Economic Review*, 2014, 104 (3).
- Geanakoplos, John**, “Liquidity, Default, and Crashes: Endogenous Contracts in General Equilibrium,” in Mathias Dewatripont, Lars Peter Hansen, and Stephen J. Turnovsky, eds., *Advances in Economics and Econometrics: Theory and Applications, Eighth World Congress*, Vol. 2 of *Econometric Society Monographs*, Cambridge University Press, 2003, pp. 170–205.
- , “The Leverage Cycle,” in Daron Acemoglu, Kenneth Rogoff, and Michael Woodford, eds., *NBER Macroeconomics Annual 2009, Volume 24*, University of Chicago Press, 2010, pp. 1–65.
- Gennaioli, Nicola and Andrei Shleifer**, “What Comes to Mind,” *Quarterly Journal of Economics*, 2010, 125 (4), 1399–1433.
- and —, *A Crisis of Beliefs: Investor Psychology and Financial Fragility*, Princeton, NJ: Princeton University Press, 2018.
- , **Yueran Ma, and Andrei Shleifer**, “Expectations and Investment,” in Martin Eichenbaum and Jonathan A. Parker, eds., *NBER Macroeconomics Annual 2015, Volume 30*, University of Chicago Press, 2016, pp. 379–431.
- Gilchrist, Simon, Charles P. Himmelberg, and Gur Huberman**, “Do Stock Price Bubbles Influence Corporate Investment?,” *Journal of Monetary Economics*, 2005, 52 (4), 805–827.
- Gourio, François, Anil K. Kashyap, and Jae W. Sim**, “The Trade offs in Leaning Against the Wind,” *IMF Economic Review*, 2018, 66 (1), 70–115.
- Guvenen, Fatih**, “A Parsimonious Macroeconomic Model for Asset Pricing,” *Econometrica*, 2009, 77 (6), 1711–1750.
- Hall, Robert E.**, “High Discounts and High Unemployment,” *American Economic Review*, 2017, 107 (2), 305–30.
- Hamilton, James D.**, “A New Approach to the Economic Analysis of Nonstationary Time Series and the Business Cycle,” *Econometrica*, 1989, 57 (2), 357–384.
- Harrison, J. Michael and David M. Kreps**, “Speculative Investor Behavior in a Stock Market with Heterogeneous Expectations,” *Quarterly Journal of Economics*, 1978, 92 (2), 323–336.
- He, Zhigu and Arvind Krishnamurthy**, “A Model of Capital and Crises,” *The Review of Economic Studies*, 2011, 79 (2), 735–777.
- Hirano, Tomohiro, Masaru Inaba, and Noriyuki Yanagawa**, “Asset Bubbles and Bailouts,” *Journal of Monetary Economics*, 2015, 76, S71–S89.
- Iachan, Felipe S., Plamen T. Nenov, and Alp Simsek**, “The Choice Channel of Financial Innovation,” 2019. Working Paper.

- Jaimovich, Nir and Sergio Rebelo**, “Can News about the Future Drive the Business Cycle?,” *American Economic Review*, 2009, 99 (4), 1097–1118.
- Kehoe, Patrick J., Pierlauro Lopez, Virgiliu Midrigan, and Elena Pastorino**, “Asset Prices and Unemployment Fluctuations,” 2019. Working Paper.
- Kydland, Finn E. and Edward C. Prescott**, “Time to Build and Aggregate Fluctuations,” *Econometrica*, 1982, 50 (6), 1345–1370.
- Lustig, Hanno and Adrien Verdelhan**, “Business Cycle Variation in the Risk-return Trade-off,” *Journal of Monetary Economics*, 2012, 59, S35–S49.
- Martin, Alberto and Jaume Ventura**, “Economic Growth with Bubbles,” *American Economic Review*, 2012, 102 (6), 3033–3058.
- Mehra, Rajnish and Edward C. Prescott**, “The Equity Premium: A Puzzle,” *Journal of Monetary Economics*, 1985, 15 (2), 145–161.
- Mendo, Fernando**, “Risk to Control Risk,” 2018. Working Paper.
- Miao, Jianjun and Pengfei Wang**, “Asset Bubbles and Credit Constraints,” *American Economic Review*, 2018, 108 (9), 2590–2628.
- Miller, Edward M.**, “Risk, Uncertainty, and Divergence of Opinion,” *Journal of Finance*, 1977, 32 (4), 1151–1168.
- Panageas, Stavros**, “The Neoclassical Theory of Investment in Speculative Markets,” 2005. Working Paper.
- Sandroni, Alvaro**, “Do Markets Favor Agents Able to Make Accurate Predictions?,” *Econometrica*, 2000, 68 (6), 1303–1341.
- Santos, Manuel S. and Michael Woodford**, “Rational Asset Pricing Bubbles,” *Econometrica*, 1997, 65 (1), 19–57.
- Scheinkman, José A. and Wei Xiong**, “Overconfidence and Speculative Bubbles,” *Journal of Political Economy*, 2003, 111 (6), 1183–1220.
- Shiller, Robert**, *Irrational Exuberance*, Princeton, NJ: Princeton University Press, 2000.
- Silva, Dejanir H.**, “The Risk Channel of Unconventional Monetary Policy,” 2019. Working Paper.
- Simsek, Alp**, “Belief Disagreements and Collateral Constraints,” *Econometrica*, 2013, 81 (1), 1–53.
- Tallarini, Thomas D.**, “Risk-sensitive Real Business Cycles,” *Journal of Monetary Economics*, 2000, 45 (3), 507–532.
- Xiong, Wei**, “Bubbles, Crises, and Heterogeneous Beliefs,” in Jean-Pierre Fouque and Joseph A. Langsam, eds., *Handbook on Systemic Risk*, Cambridge University Press, 2013, pp. 663–713.

Appendices

A Proofs

A.1 Proof of Lemma 1

Proof. By conjecturing that $U_i(\hat{a}_i, X, s; ; A_{-1}) = V_i(X, s)\hat{a}_i A_{-1}$, after a monotone transformation by taking logs, one obtains,

$$\ln V_i(X, s) + \ln \hat{a}_i = \max_{\tilde{c}_i, \tilde{b}_i} \left\{ (1 - \beta)(\ln \tilde{c}_i + \ln \hat{a}_i) + \frac{\beta}{1 - \gamma} \ln \left(\mathbb{E}_{i,s} \left[(V_i(X', s') \hat{a}_i)^{1-\gamma} \right] \right) \right\}.$$

By plugging (3) into the equation above, after some algebra, one verifies that $\ln \hat{a}_i$ cancels on both sides, and obtains,

$$\begin{aligned} \ln V_i(X, s) &= \max_{\tilde{c}_i} \{ (1 - \beta) \ln \tilde{c}_i + \beta \ln(1 - \tilde{c}_i) \} + \\ &+ \frac{\beta}{1 - \gamma} \ln \left(\max_{\tilde{b}_i} \left\{ \mathbb{E}_{i,s} \left[\left(V_i(X', s') \left[R_n(X, s, s')(1 - \tilde{b}_i) + R_b(X, s)\tilde{b}_i \right] \right)^{1-\gamma} \right] \right\} \right), \end{aligned}$$

which separates consumption from portfolio decisions. In particular, the FOC with respect to consumption yields $\tilde{c}_i(X, s) = 1 - \beta$ for all i and (X, s) , and the FOC with respect the portfolio weight yields equation (4). \square

A.2 Proof of Proposition 1

Proof. We further develop equation (8) by using the market clearing condition for goods, the optimal consumption plan $\hat{c}_i = (1 - \beta)\hat{a}_i$ from the investors' optimization problem, the worker's budget constraint, and the equilibrium labor market outcomes in equation (2). Hence, the following equation is obtained,

$$\beta \sum_i \mu_i \hat{a}_i = \hat{q}(X, s) - \bar{g}_s \hat{B}'(X, s).$$

By plugging this expression back into equation (8), one obtains equation (9). \square

A.3 Proof of Lemma 3

Proof. The first part is a direct implication of Lemma 2. To prove the second part, note that there are I absorbing states, each with $\eta_j = 1$ for some j , and $\eta_i = 0$ for $i \neq j$. Indeed, suppose an absorbing state features $\eta_j \in (0, 1)$ for some js . Hence, $\sum_{i=1}^I p_{ss'}^i \eta_i = p_{ss'}^j$ for all j such that $\eta_j \in (0, 1)$, which contradicts heterogeneous beliefs. With an abuse of notation, upon reaching state s today, denote the continuation history m periods ahead by $s^m|s = \{s, s_1, s_2, \dots, s_m\}$. Investor j attaches probability $\mathbf{p}^j(s^m|s) = p_{s,s_1}^j \prod_{t=1}^{m-1} p_{s_t, s_{t+1}}^j$ that such history will occur. In

addition, given that investor $i = 1$ has the correct beliefs, $\frac{p^j(s^m|s)}{p^1(s^m|s)}$ is a non-negative martingale and, thus, converges almost surely. Hence, $\frac{\eta_j}{\eta_1}$ also converges almost surely, meaning that $\eta_1 \rightarrow 1$ almost surely. \square

B Wealth evolution (3) and FOC (4) for S=2

In this appendix, given $q(X, s)$, $R_b(X, s)$, $R_n(X, s)$, $\psi(X, s, s')$ and $v_i(X, s, s') = V_i(\psi(X, s, s'), s')$, we derive quasi-closed form solutions for the optimal portfolio choices when $S = 2$. Equation (4) can be rewritten as

$$\mathbb{E}_{i,s} \left[\frac{(v_i(X, s, s'))^{1-\gamma} [R_b(X, s) - R_n(X, s, s')]}{\left[R_n(X, s, s')(1 - \tilde{b}_i(X, s)) + R_b(X, s)\tilde{b}_i(X, s) \right]^\gamma} \right] = 0.$$

Under the assumption of two shocks, say $s \in \{L, H\}$,

$$p_{s,L}^i \left[\frac{(v_i(X, s, L))^{1-\gamma} [R_b(X, s) - R_n(X, s, L)]}{\left[R_n(X, s, L)(1 - \tilde{b}_i(X, s)) + R_b(X, s)\tilde{b}_i(X, s) \right]^\gamma} \right] = -p_{s,H}^i \left[\frac{(v_i(X, s, H))^{1-\gamma} [R_b(X, s) - R_n(X, s, H)]}{\left[R_n(X, s, H)(1 - \tilde{b}_i(X, s)) + R_b(X, s)\tilde{b}_i(X, s) \right]^\gamma} \right],$$

which can be rewritten as

$$\frac{\{p_{s,L}^i(v_i(X, s, L))^{1-\gamma} [R_b(X, s) - R_n(X, s, L)]\}^{\frac{1}{\gamma}}}{R_n(X, s, L) + [R_b(X, s) - R_n(X, s, L)]\tilde{b}_i(X, s)} = \frac{\{p_{s,H}^i(v_i(X, s, H))^{1-\gamma} [R_n(X, s, H) - R_b(X, s)]\}^{\frac{1}{\gamma}}}{R_n(X, s, H) - [R_n(X, s, H) - R_b(X, s)]\tilde{b}_i(X, s)}.$$

By further developing the expression above,

$$\begin{aligned} & \left\{ p_{s,L}^i(v_i(X, s, L))^{1-\gamma} [R_b(X, s) - R_n(X, s, L)] \right\}^{\frac{1}{\gamma}} \{ R_n(X, s, H) - [R_n(X, s, H) - R_b(X, s)]\tilde{b}_i(X, s) \} = \\ & = \left\{ p_{s,H}^i(v_i(X, s, H))^{1-\gamma} [R_n(X, s, H) - R_b(X, s)] \right\}^{\frac{1}{\gamma}} \{ R_n(X, s, L) + [R_b(X, s) - R_n(X, s, L)]\tilde{b}_i(X, s) \}. \end{aligned}$$

Hence, after collecting terms,

$$\tilde{b}_i(X, s) = \frac{\{p_{s,L}^i(v_i(X, s, L))^{1-\gamma} [R_b(X, s) - R_n(X, s, L)]\}^{\frac{1}{\gamma}} R_n(X, s, H) - \{p_{s,H}^i(v_i(X, s, H))^{1-\gamma} [R_n(X, s, H) - R_b(X, s)]\}^{\frac{1}{\gamma}} R_n(X, s, L)}{\{p_{s,H}^i(v_i(X, s, H))^{1-\gamma} [R_n(X, s, H) - R_b(X, s)]\}^{\frac{1}{\gamma}} [R_b(X, s) - R_n(X, s, L)] + \{p_{s,L}^i(v_i(X, s, L))^{1-\gamma} [R_b(X, s) - R_n(X, s, L)]\}^{\frac{1}{\gamma}} [R_n(X, s, H) - R_b(X, s)]}. \quad (12)$$

Now, by plugging the equation above in equation (3), i.e. $\frac{\bar{g}_s a'(X, s, s')}{\beta a} = R_n(X, s, s') + [R_b(X, s) - R_n(X, s, s')]\tilde{b}_i(X, s)$, and working it out, one obtains that

$$\begin{aligned} & R_n(X, s, s') + [R_b(X, s) - R_n(X, s, s')] \times \\ & \times \frac{\{p_{s,L}^i(v_i(X, s, L))^{1-\gamma} [R_b(X, s) - R_n(X, s, L)]\}^{\frac{1}{\gamma}} R_n(X, s, H) - \{p_{s,H}^i(v_i(X, s, H))^{1-\gamma} [R_n(X, s, H) - R_b(X, s)]\}^{\frac{1}{\gamma}} R_n(X, s, L)}{\{p_{s,H}^i(v_i(X, s, H))^{1-\gamma} [R_n(X, s, H) - R_b(X, s)]\}^{\frac{1}{\gamma}} [R_b(X, s) - R_n(X, s, L)] + \{p_{s,L}^i(v_i(X, s, L))^{1-\gamma} [R_b(X, s) - R_n(X, s, L)]\}^{\frac{1}{\gamma}} [R_n(X, s, H) - R_b(X, s)]} = \\ & = \frac{\{p_{s,s'}^i(v_i(X, s, s'))^{1-\gamma} [R_b(X, s) - R_n(X, s, s')]\}^{\frac{1}{\gamma}} [R_b(X, s) - R_n(X, s, L)]R_n(X, s, H) + \{p_{s,s'}^i(v_i(X, s, s'))^{1-\gamma} [R_b(X, s) - R_n(X, s, s')]\}^{\frac{1}{\gamma}} [R_n(X, s, H) - R_b(X, s)]R_n(X, s, L)}{\{p_{s,H}^i(v_i(X, s, H))^{1-\gamma} [R_n(X, s, H) - R_b(X, s)]\}^{\frac{1}{\gamma}} [R_b(X, s) - R_n(X, s, L)] + \{p_{s,L}^i(v_i(X, s, L))^{1-\gamma} [R_b(X, s) - R_n(X, s, L)]\}^{\frac{1}{\gamma}} [R_n(X, s, H) - R_b(X, s)]} = \end{aligned}$$

$$\begin{aligned}
&= \frac{\left\{ p_{s,s'}^i (v_i(X, s, s'))^{1-\gamma} [|R_b(X, s) - R_n(X, s, s')|] \right\}^{\frac{1}{\gamma}} \{ [R_b(X, s) - R_n(X, s, L)] R_n(X, s, H) + [R_n(X, s, H) - R_b(X, s)] R_n(X, s, L) \}}{\left\{ p_{s,H}^i (v_i(X, s, H))^{1-\gamma} [R_n(X, s, H) - R_b(X, s)] \right\}^{\frac{1}{\gamma}} [R_b(X, s) - R_n(X, s, L)] + \left\{ p_{s,L}^i (v_i(X, s, L))^{1-\gamma} [R_b(X, s) - R_n(X, s, L)] \right\}^{\frac{1}{\gamma}} [R_n(X, s, H) - R_b(X, s)]} = \\
&= \frac{\left\{ p_{s,s'}^i (v_i(X, s, s'))^{1-\gamma} [|R_b(X, s) - R_n(X, s, s')|] \right\}^{\frac{1}{\gamma}} R_b(X, s) [R_n(X, s, H) - R_n(X, s, L)]}{\left\{ p_{s,H}^i (v_i(X, s, H))^{1-\gamma} [R_n(X, s, H) - R_b(X, s)] \right\}^{\frac{1}{\gamma}} [R_b(X, s) - R_n(X, s, L)] + \left\{ p_{s,L}^i (v_i(X, s, L))^{1-\gamma} [R_b(X, s) - R_n(X, s, L)] \right\}^{\frac{1}{\gamma}} [R_n(X, s, H) - R_b(X, s)]}. \tag{13}
\end{aligned}$$

To obtain the expressions reported in Lemma 2, just set $\gamma = 1$ into equations (12) and (13) above.

C Equilibrium equations and numerical algorithm

Throughout the main text, we consider $X = \{\hat{E}, \{\eta_i\}_{i=1}^{I-1}, R\hat{B}\}$ because it facilitates exposition and intuition. However, to implement the model numerically, we propose a change of the state space that renders the algorithm more efficient.

First, instead of using $R\hat{B}$ as a state variable, we consider $\eta_w = \frac{R\hat{B}}{\sum_i \mu_i \hat{a}_i + R\hat{B}}$. Recall that $\sum_i \mu_i \hat{a}_i + R\hat{B} = \hat{q}(X, s) + \hat{E}$. Hence by substituting $R\hat{B}$ for $\eta_w(\hat{q}(X, s) + \hat{E})$ in the model, we can characterize the equilibrium in terms of η_w , which is bounded in between zero and one, rather than $R\hat{B}$. Note that we do not change the definition of $\eta_i = \frac{\mu_i \hat{a}_i}{\sum_i \mu_i \hat{a}_i}$, which represents wealth shares among investors only.

Second, instead of using the current de-trended profits \hat{E} , we consider the **previous** expected value for TFP computed using firm's adjusted beliefs. Call it ε . Hence, with slight abuse of notation, current de-trended profits are given by

$$\hat{\pi}(X, s) = \frac{1}{\xi^{\frac{\alpha}{1+\nu-\alpha}}} \left\{ \bar{g}_s(\alpha\varepsilon)^{\frac{\alpha}{1+\nu-\alpha}} - (\alpha\varepsilon)^{\frac{1+\nu}{1+\nu-\alpha}} \right\}.$$

Note that by knowing s and ε , one can compute $\hat{E} = \hat{\pi}(X, s)$. The advantage of using ε instead of \hat{E} is twofold. First, it is bounded by the lowest and the highest TFP growth shocks. Second, it renders a more efficient numerical algorithm as the updating rule for ε , $\varepsilon'(X, s) = \mathbb{E}_{X,s}^f[g]$, does not depend on the future exogenous state s' .

Again, with a slight abuse of notation, let $X = \{\varepsilon, \{\eta_i\}_{i=1}^{I-1}, \eta_w\}$. Given this change of variables, a simple algebra reveals that the key equation (9) in the main text now reads

$$\hat{q}(X, s) = \frac{\beta(1 - \eta_w)\hat{\pi}(X, s) + \bar{g}_s \hat{B}'(X, s)}{1 - \beta(1 - \eta_w)}. \quad (14)$$

Under the assumption that $S = 2$, equation (12) in Appendix B solves for the optimal interior portfolio weights $\tilde{b}_i(X, s)$ as functions of $\hat{q}(X, s)$, $R_b(X, s)$, $R_n(X, s)$, $\psi(X, s, s')$ and $v_i(X, s, s') = V_i(\psi(X, s, s'), s')$. This is the only part of the solution method that the assumption of $S = 2$ kicks in.

The market clearing condition for bonds reads:

$$\bar{g}_s \hat{B}'(X, s) = - \frac{\beta(\hat{\pi}(X, s) + \bar{g}_s \hat{B}'(X, s) - \eta_w(\hat{q}(X, s) + \hat{\pi}(X, s)))}{1 - \beta} \sum_i \tilde{b}_i(X, s) \eta_i.$$

If $\gamma \neq 1$, then the value function $V_i(X, s)$ satisfies the following recursion:

$$\begin{aligned} \ln V_i(X, s) &= (1 - \beta) \ln(1 - \beta) + \beta \ln \beta + \\ &+ \frac{\beta}{1 - \gamma} \ln \left(\left\{ \mathbb{E}_{i,s} \left[\left(V_i(\psi(X, s, s'), s') \left[R_n(X, s, s')(1 - \tilde{b}_i(X, s)) + R_b(X, s) \tilde{b}_i(X, s) \right] \right)^{1-\gamma} \right] \right\} \right). \end{aligned}$$

If $\gamma = 1$ the solution simplifies as portfolios weights, $\tilde{b}_i(X, s)$, cease to depend on $v_i(X, s, s')$, and thus, the recursion above when $\gamma \rightarrow 1$ is immaterial for the solution.

The law of motion $X' = \psi(X, s, s')$ is implicitly defined by

$$\begin{aligned}\varepsilon'(X, s) &= \mathbb{E}_{X,s}^f[g]; \\ \eta'_i(X, s, s') &= \frac{\left[R_n(X, s, s')(1 - \tilde{b}_i(X, s)) + R_b(X, s)\tilde{b}_i(X, s) \right] \eta_i}{\sum_{i=1}^I \left[R_n(X, s, s')(1 - \tilde{b}_i(X, s)) + R_b(X, s)\tilde{b}_i(X, s) \right] \eta_i}; \\ \eta'_w(X, s, s') &= \frac{R_b(X, s)\hat{B}'(X, s)}{q(\psi(X, s, s'), s') + \hat{\pi}(\psi(X, s, s'), s')}.\end{aligned}$$

Recall that stock returns are given by

$$R_n(X, s, s') = \frac{\hat{q}(\psi(X, s, s'), s') + \hat{\pi}(\psi(X, s, s'), s')}{\hat{q}(X, s)},$$

whereas the stochastic discount factors are given by

$$SDF_i(X, s, s') = \frac{(v_i(X, s, s'))^{1-\gamma}}{\left[R_n(X, s, s')(1 - \tilde{b}_i(X, s)) + R_b(X, s)\tilde{b}_i(X, s) \right]^\gamma}.$$

Finally, firm's adjusted beliefs are given by

$$p_{ss'}^f(X) = \sum_i \omega_i(X, s) \left[\frac{SDF_i(X, s, s') p_{ss'}^i}{\sum_{s'} SDF_i(X, s, s') p_{ss'}^i} \right],$$

where we assume that weights are given by

$$\omega_i(X, s) = \frac{\mathbf{I}_{\{\tilde{b}_i(X, s) < 1\}} \eta_i (1 - \tilde{b}_i(X, s))}{\sum_i \mathbf{I}_{\{\tilde{b}_i(X, s) < 1\}} \eta_i (1 - \tilde{b}_i(X, s))}.$$

After assuming a functional form for the supply of funds to the investors, $B'(X, s)$, one can solve numerically the model in the computer.

We propose the following numerical algorithm to compute globally the equilibrium. Discretize ε , $\{\eta_i\}_{i=1}^{I-1}$ and η_w , fix $\hat{B}'(X, s)$, and conjecture the law of motion $X' = \psi(X, s, s')$. The idea is to iterate over $\psi(X, s, s')$ until convergence is reached. With $\hat{B}'(X, s)$ and a guess for $\psi(X, s, s')$ at hand, one can compute $\hat{q}(X, s)$ and $R_n(X, s, s')$. Hence, within each iteration, use the bisection method to find $R_b(X, s)$ that clears the bonds market, keeping in mind that if $\gamma \neq 1$, one also needs to iterate over $V_i(X, s)$ inside this inner loop to compute $\tilde{b}_i(X, s)$. With $R_n(X, s, s')$, $R_b(X, s)$ and $\tilde{b}_i(X, s)$ at hand, one can update $\eta'_i(X, s, s')$, $\eta'_w(X, s, s')$ and $\varepsilon'(X, s)$ to obtain a new guess for $\psi(X, s, s')$ in the next iteration.

This numerical algorithm is easily malleable if we impose short-selling constraints: $\tilde{b}_i(X, s) \leq 1 + \kappa$, with $\kappa \in [0, \infty)$. All we need is to replace $\tilde{b}_i(X, s)$ obtained in equation (12) after working out the FOCs by $1 + \kappa$ if the optimal

interior solution is outside this bound.

C.1 Specialization: $\gamma = 1, S = 2, \kappa = \infty, \hat{B}'(X, s) = \hat{B} = 0, I = 2$

As argued above, if $\gamma = 1$ and $S = 2$ with $s \in \{L, H\}$, the solution simplifies as $\tilde{b}_i(X, s)$ and $SDF_i(X, s, s')$ cease to depend on $v_i(X, s, s')$ (see expressions obtained in Lemma 2 in the main text), avoiding an inner loop to compute $v_i(X, s, s')$ recursively. If in addition $\kappa = \infty$, as discussed in the main text, we obtain the following closed form for the evolution of wealth shares,

$$\eta'_i(X, s, s') = \frac{p_{ss'}^i \eta_i}{\sum_{i=1}^I p_{ss'}^i \eta_i}.$$

Also, under these assumptions, firm's adjusted beliefs become a simple function of the stock returns and risk-free rate,

$$p_{sL}^f(X, s) = \frac{R_n(X, s, H) - R_b(X, s)}{R_n(X, s, H) - R_n(X, s, L)} \quad \text{and} \quad p_{sH}^f(X, s) = \frac{R_b(X, s) - R_n(X, s, L)}{R_n(X, s, H) - R_n(X, s, L)}.$$

These formulas are valid for any supply of funds schedule, $\hat{B}'(X, s)$, and simplify the computation of the equilibrium as there is no need to iterate over $\eta'_i(X, s, s')$ and $v_i(X, s, s')$.

Now if we assume there is no supply of funds to the investors, $\hat{B}'(X, s) = \hat{B} = 0$, which implies $\eta'_w(X, s, s') = \eta_w = 0$, the solution for the price per share is simplified,

$$\hat{q}(X, s) = \frac{\beta \hat{\pi}(X, s)}{1 - \beta},$$

which implies the following expression for asset returns,

$$R_n(X, s, s') = \frac{\pi(\psi(X, s, s'), s')}{\beta \hat{\pi}(X, s)} = \frac{1}{\beta} \frac{\bar{g}_{s'}(\alpha \varepsilon'(X, s))^{\frac{\alpha}{1+\nu-\alpha}} - (\alpha \varepsilon'(X, s))^{\frac{1+\nu}{1+\nu-\alpha}}}{\bar{g}_s(\alpha \varepsilon)^{\frac{\alpha}{1+\nu-\alpha}} - (\alpha \varepsilon)^{\frac{1+\nu}{1+\nu-\alpha}}}.$$

In addition, the market clearing condition for bonds reads

$$\sum_{i=1}^I \eta_i \tilde{b}_i(X, s) = 0.$$

These expressions are valid for any value of γ, κ or S . They facilitate the implementation of the equilibrium as, of course, we do not need to track η_w as a state variable, reducing the dimensionality of the state space.

Finally, by assuming altogether $\gamma = 1, S = 2, \kappa = \infty, \hat{B}'(X, s) = \hat{B} = 0$ as well as $I = 2$, one can use the expressions in Lemma 2 to work out the above market clearing condition for bonds, and reach an expression for $R_b(X, s)$ also as a function of stock returns,

$$R_b(X, s) = \frac{R_n(X, s, H)R_n(X, s, L)}{(\eta_1 p_{sL}^1 + \eta_2 p_{sL}^2)R_n(X, s, H) + (\eta_1 p_{sH}^1 + \eta_2 p_{sH}^2)R_n(X, s, L)}.$$

These extra assumptions make the numerical computation of the equilibrium significantly more efficient, as

one only needs to iterate over $\varepsilon'(X, s)$ within a state space with lower dimensionality. In addition, there is no need to use the bisection method to find the $R_b(X, s)$ that clears the market.

If we impose an upper bound on the portfolio decisions, $\kappa < \infty$, to obtain speculative behavior, the numerical implementation becomes less efficient. In particular, one also needs to iterate over $\eta_i(X', s, s)$, check whether the optimal interior solution for $\tilde{b}_i(X, s)$ is outside the bound (and replace accordingly), and run the bisection method to find $R_b(X, s)$ that clears the bonds market. Nonetheless the state space is now reduced due to the assumption that $\hat{B}'(X, s) = \hat{B} = 0$.

C.2 Specialization: $\hat{B}'(X, s) = \phi_s(\hat{w}h + R\hat{B})$

We consider in this subsection the case $\hat{B}'(X, s) = \phi_s(\hat{w}h + R\hat{B})$. Recall that the algorithm outlined above compute the endogenous objects by taking $B'(X, s)$ as given. Also recall that the state space is represented by $X = \{\varepsilon, \{\eta_i\}_{i=1}^{I-1}, \eta_w\}$. Hence,

$$\hat{B}'(X, s) = \phi_s(\hat{w}(\varepsilon)h(\varepsilon) + \eta_w(\hat{q}(X, s) + \hat{\pi}(X, s))), \quad (15)$$

which coupled with equation (14) determines both $B'(X, s)$ and $q(X, s)$ for each (X, s) . In this linear case, it is straight forward to obtain the solution. Indeed, one can further develop (14) by plugging (15) into, and after rearranging terms, obtain the following expression:

$$\hat{q}(X, s) = \frac{(\beta(1 - \eta_w) + \bar{g}_s \phi_s \eta_w) \hat{\pi}(X, s) + \bar{g}_s \phi_s \hat{w}(\varepsilon) h(\varepsilon)}{1 - \beta(1 - \eta_w) - \bar{g}_s \phi_s \eta_w}.$$

By plugging this expression back into (15), one obtains $B'(X, s)$ as a direct function of the state space, and can solve the model by applying the algorithm outlined above.

If the supply of funds were not linear, an intermediate step in the algorithm would be necessary to solve for both $\hat{q}(X, s)$ and $\hat{B}'(X, s)$ simultaneously.