A Learning Model of Financial Instability

Noah Williams*
Miami Herbert Business School
University of Miami
noahwilliams@miami.edu

November 22, 2024

Abstract

I develop an adaptive learning model where periods of stability lead to periods of instability, similar to Minsky's financial instability hypothesis. In tranquil times, investors increase their subjective risk-return assessments. They buy more stocks, driving up prices, and reinforcing the change in beliefs. Crashes arise endogenously as rapid booms increase perceptions of risk as well as returns. The model can help explain observed asset pricing phenomena. I also establish new results on instability and cycles in adaptive learning models. As long as agents put sufficient weight on new information, learning from asset returns leads to instability that drives endogenous booms and busts.

1 Introduction

1.1 Overview

In this paper, I develop an adaptive learning model that generates recurrent, endogenous fluctuations asset prices. Changes in expectations and risk assessments drive flows of funds in and out of risky assets, leading to unstable feedback between beliefs and returns. This dynamic instability generates booms and busts in asset prices reminiscent of Minsky's financial instability hypothesis.[1] As in Minsky's theory, periods of relative market tranquility

*I thank Fernando Alvarez, Gadi Barlevy, Marco Bassetto, Jeff Campbell, Cristina De Nardi, George Evans, Cars Hommes, Seppo Honkapohja, Alex Horenstein, Bruce McGough, Stefan Nagel, Anton Nakov, Jordan Rappaport, Laura Veldkamp, and especially Tom Sargent for helpful comments and discussion. Francesco Bianchi, the editor, and three anonymous referees provided valuable suggestions and feedback.

[1]Minsky (1975) initially developed the theory. More references are below.

lead agents to decrease their risk estimates and invest more in risky assets. But as investors become more levered, the market becomes more sensitive to changes in beliefs and asset positions, and small downturns quickly lead to large market crashes.

I study a model where agents can borrow or lend at fixed risk-free rate of return, and they can purchase a stock providing a claim to a risky dividend process. While asset returns are determined in equilibrium, investors believe that risky asset returns have a constant but unknown mean and variance. They estimate expected returns and return volatility based on past asset returns, and update these estimates over time. If agents had rational expectations and knew the shock distribution, then returns would indeed be i.i.d. with a constant mean and variance. All of the persistent fluctuations in the model are driven by learning dynamics.

In this environment, I study a learning rule where agents place a constant weight (or gain) on new information. By giving more weight to more recent observations, such rules are known to track drifting parameters (see Sargent and Williams, 2005), and provide robust estimates under model misspecification (see Evans, Honkapohja, and Williams, 2010). Both motivations are relevant here, as investors may believe that returns change over time, and the feedback from beliefs to returns makes this belief self-reinforcing.[2] The constant gain learning rule exhibits forgetting, which makes learning persistent. While the model is simple, it generates rich dynamics as simulations are characterized by repeated booms and crashes. I show that these price dynamics are driven by endogenous belief cycles. As long as agents put sufficient weight on new information, learning from asset returns results in a fundamental dynamic instability which drives the financial instability.

Since the equilibrium is unstable, shocks can lead to persistent fluctuations. For example, a sequence of positive shocks can make agents believe that mean returns have increased, and the volatility has fallen. In response, investors increase their risky asset positions, with the inflow of funds further driving up their prices. This process resembles an asset price bubble, where expectations of higher (risk-adjusted) returns lead, at least temporarily, to higher realized returns, and prices rapidly increase.

However once agents are sufficiently invested in stocks, the economy becomes fragile. A small change in portfolio allocations then has a large impact on prices, making the market prone to crash. In the model, booms end due to two distinct mechanisms: binding constraints and increased subjective volatility. The model incorporates borrowing constraints to ensure that an equilibrium exists. When demand for risky assets increases, these constraints may restrict flows into the risky asset, dampening price growth and leading to a market reversal. This is similar to models of credit frictions, where tightening constraints can lead to a crash.[3] A more novel source of market crashes, which is predominant in the model, is that a rapid

[2]Malmendier and Nagel (2016) find that constant gain learning also explains average survey expectations.

[3]See Eggertsson and Krugman (2012), Boz and Mendoza (2014), and Farboodi and Kondor (2023).

asset price boom may cause its own end by increasing uncertainty. Since agents are risk averse, their portfolio allocations reflect the perceived tradeoff between risk and returns. With a rapid increase in prices, agents' estimates of both the mean and variance of returns increase. The variance effect may dominate, leading to a cut in the portfolio share of risky assets, and a drop in prices.

No matter the cause of the initial asset price decline, any price reversal following a boom leads to a rapid market crash. Following a period of rapid growth, a negative return is influential in changing beliefs and, because of the fragility, market outcomes. As investors seek to sell stocks, the outflow of funds puts downward pressure on the stock price, which rapidly plummets until it his bottom. After having gone through a recent boom and bust, agents now view risky returns as highly volatile, and so are wary of investing more in stocks. As the market stabilizes, returns become less volatile, and investors start to forget the previous boom and bust episode. As they reduce their volatility estimate, they shift more assets to stocks, driving up prices and returns once again. The process starts slowly, but gains speed until a rapid boom leads to another crash. A fundamental dynamic instability due to learning drives outcomes similar to Minsky's financial instability.

While my main contribution is theoretical, I also explore the model's quantitative implications. Using a calibrated dividend growth process as an input, I show that the model can match the level and volatility of the price-dividend ratio, as well as matching the level and volatility of stock returns. Matching risk premia and asset price volatility with a relatively low degree of risk aversion is a relatively difficult task for a sparsely parameterized model like mine. The baseline calibration generates significant market crashes with plausible regularity, although less frequently than in the data. I also discuss empirical drawbacks of the model and how it can be reconciled with survey evidence on expected returns.

1.2 Implications and Related Literature

Minsky developed the financial instability hypothesis in a series of papers and books (Minsky, 1975, 1986, 1989, 1992). Kindleberger (1978) used the financial instability hypothesis as an organizing principle in his classic study of bubbles. Other early references include Taylor and O'Connell (1985) and Keen (1985). In Sethi (1992) financial instability results from switches between equilibria. Friedman and Laibson (1989)used a portfolio choice model with learning to interpret evidence of excess market volatility. In Eggertsson and Krugman (2012) changes in borrowing constraints can lead to a Minsky-style crisis of deleveraging. L'Huillier, Phelan, and Wieman (2024) obtain Minsky-style cycles from the interaction of credit constraints and sectoral reallocation of capital.

I draw on the literature of adaptive learning in macroeconomics, including Marcet and Sargent (1989) and Evans and Honkapohja (2001). Timmermann (1993) showed how learning

can generate volatility and predictability in returns. Like this paper, Branch and Evans (2011) consider agents who learn about risk and returns, and show how risk dynamics can generate recurrent booms and crashes. In their model, stochastic shocks under learning can lead to occasional escapes – in the sense of Cho, Williams, and Sargent (2002) and Williams (2019) – from the equilibrium. By contrast, the recurrent cycles in my model are driven by instability. Booms and crashes are typical outcomes in my environment, but rare events in their model. Nonetheless, there are clear similarities between the unstable dynamics that I find and the locally divergent dynamics which drive escapes.

My model builds on Adam, Marcet, and Nicolini (2016), who showed how learning can lead to empirically plausible volatility in prices and price-dividend ratios. Relative to their paper, I add a risk-free asset with a fixed interest rate, which introduces a portfolio problem. The resulting debt dynamics are crucial in generating the boom-bust cycles and instability that I find. Later learning papers, such as Adam, Marcet, and Beutel (2017) and Nagel and Xu (2022), added features to improve the empirical performance in matching data on consumption, asset prices, and survey measures of expectations.

My main methodological contribution to the learning literature is in showing that the standard convergence results only hold in my model for very small gain settings. As the gain increases, the equilibrium has a bifurcation where it loses stability, and a stable limit cycle arises. I characterize this instability by applying results on bifurcations in discrete dynamical systems from Kuznetsov (1998). To my knowledge, these results are novel in the economic learning literature.[4] Sargent and Williams (2005) studied related bifurcations in learning. Cho (2018) and Chien, Cho, and Ravikumar (2020) showed that small-gain limits in models with second-order dynamics may give a misleading characterization of stability. Evans and Honkapohja (2009) and Honkapohja and McClung (2023) documented instances of instability in learning with small gains.

A separate literature in finance uses psychologically-motivated models of expectation formation, as described by Gennaioli and Shleifer (2018). Barberis, Greenwood, Jin, and Shleifer (2018) show that extrapolation can lead to asset price bubbles in an environment like mine with a portfolio problem. Bordalo, Gennaioli, and Shleifer (2018) show that diagnostic expectations can lead to credit cycles driving asset price and output fluctuations. Jin and Sui (2022) show that learning and extrapolation can explain both asset market and survey expectation data. Greenwood, Hanson and Jin (2019) consider a model of credit cycles with feedback between investor beliefs and firms' creditworthiness. Maxted (2023) and Krishnamurthy and Li (2022) consider the interaction of sentiment and financial frictions in general equilibrium models. Bianchi, Ilut, and Saijo (2024) incorporate diagnostic expectations in business cycles, showing that the interaction between actions and beliefs nat-

[4]Beaudry, Galizia, and Portier (2020) apply related results to study endogenous business cycles.

urally generates repeated boom–bust cycles. Brianti and Cormun (2024) find that changes in survey forecast revisions unrelated to fundamentals predict boom-bust dynamics in macro variables, and develop a model of expectation-shock driven business cycles.

Like this literature, I focus on the role of beliefs in asset price booms and busts. In many of the earlier papers, the end of the bubble means a reversion to the fundamental asset valuations. In my model market crashes lead to a collapse in asset valuations, typically far below fundamental prices. Later, periods of stability lead to instability, with rapid asset price booms making crashes more likely. Such features are characteristic of Minsky's theory as well as experienced financial crises. For example, Greenwood, Hanson, Shleifer, and Sørensen (2022) show that periods of rapid credit and asset price growth greatly increase the likelihood that a financial crisis will follow. I show that financial instability is present in my model in a precise mathematical sense. As long as agents put sufficient weight on new information, a weight which may be quite small quantitatively, the model displays cyclical dynamics marked by recurrent booms and busts.

2 The Model and Rational Expectations Equilibrium

2.1 The Model

The model is a variation on a standard Lucas (1978) consumption-based asset pricing model, which was studied a learning context by Adam, Marcet, and Nicolini (2016). One key difference with the standard setup in those papers is that I study a small open economy version of the model, where agents can borrow and lend at a constant risk-free interest rate. This is a simple way of introducing (net) borrowing and lending, and hence asset flows, while retaining the simplicity of the representative agent framework.

There are two assets: stocks denoted S_t which trade at the price P_t, and bonds B_t that have a constant risk-free one-period gross return R. There is a unit supply of stock, so in equilibrium $S_t = 1$. Since agents can borrow and lend freely on the world market at a fixed return, domestic holdings of bonds B_t are endogenous. The driving stochastic process is an exogenous dividend stream D_t which accrues to stockholders. Like much of the recent literature, I assume that dividend growth is i.i.d.:

$$\log D_{t+1} = \log D_t + d + \sigma W_{t+1} \tag{1}$$

where $d > 0$ is the constant mean rate of dividend growth, $\sigma > 0$ is its standard deviation, and W_{t+1} is an i.i.d. shock. For ease of analysis, I assume (centered) binomial shocks: $W_t \in \{-1, 1\}$ each with probability 0.5. The discreteness is not crucial but simplifies the analysis.[5]

[5]Gaussian shocks lead to very similar outcomes, but cause some complications in the analysis.

Agents have identical time separable power utility over consumption C_t with risk aversion parameter γ and discount factor β. Agents choose how much to consume and how to allocate their portfolio between stocks and bonds:

$$\max E_0 \sum_{t=0}^{\infty} \beta^t \frac{C_t^{1-\gamma}}{1-\gamma}$$

subject to the budget constraint:

$$P_t S_t + B_t + C_t = (P_t + D_t) S_{t-1} + R B_{t-1}. \tag{2}$$

Defining the gross return on stocks as $Z_t = (P_t + D_t)/P_{t-1}$, the optimality conditions lead to the standard Euler equations:

$$\begin{aligned} C_t^{-\gamma} &= \beta R E_t \left[C_{t+1}^{-\gamma} \right], & (3) \\ C_t^{-\gamma} &= \beta E_t \left[C_{t+1}^{-\gamma} Z_{t+1} \right]. & (4) \end{aligned}$$

Below I study arbitrary beliefs, so I impose a portfolio or leverage constraint to ensure that an equilibrium exists. The constraint limits debt, restricting the share of wealth (greater than unity) that an agent can invest in the risky asset. In general the Euler equations are inequalities, with (3)-(4) holding when the leverage constraint does not bind.

2.2 Rational Expectations Equilibrium

I now solve for a rational expectations equilibrium with zero debt, which serves as an important reference point. In equilibrium the stock holdings equal the share supply $S_t \equiv 1$, so the aggregate resource constraint is:

$$C_t = D_t + R B_{t-1} - B_t.$$

Equilibrium consumption is equal to dividends plus the net change in bonds, after accounting for interest costs. For an arbitrary interest rate R, the economy have a positive or negative net bond position. But I set R at the closed-economy equilibrium level, so equilibrium outcomes match those when bonds are in zero net supply.

The simple structure of the economy, with power utility and i.i.d. dividend growth, allows me to completely solve for the zero-debt equilibrium. The calculations, which are standard, are detailed in Appendix A.1. In equilibrium, the price-dividend ratio is constant, at a value δ^* given in the appendix, and the gross return on stocks is i.i.d.:

$$\begin{aligned} P_t &= \delta^* D_t, \\ \log Z_t &= \log(1+\delta^*) - \log \delta^* + d + \sigma W_t. & (5) \end{aligned}$$

For later reference, it is useful to describe the agent's decision rules. Define an agent's wealth X_t as the right side of the budget constraint (2):

$$X_t = RB_{t-1} + (P_t + D_t)S_{t-1},$$

and let v_t be the share of post-consumption wealth invested in stocks:

$$v_t = \frac{P_t S_t}{X_t - C_t}.$$

In equilibrium $v_t \equiv 1$ and $X_t = P_t + D_t = (1 + \delta^*)D_t$, so $C_t = \frac{1}{1+\delta^*} X_t = D_t$. That is, agents consume a constant fraction of wealth and invest a constant fraction (unity) of post-consumption wealth in stocks. These decision rules are special cases of Hakansson (1970), which is a discrete time version Merton (1969). This equivalence is natural, as Hakansson studied consumption and portfolio choice problems with i.i.d. returns, as they are in the rational expectations equilibrium. The zero-debt equilibrium returns are determined so that the agent invests entirely in stocks.

3 Equilibrium for Arbitrary Beliefs

3.1 Self-Confirming Equilibrium

As a step toward learning, I first analyze the consequences of agents having arbitrary (but fixed) beliefs. To make their consumption and portfolio decisions, agents do not need to know the process for dividends or understand the determinants of equilibrium prices. All they need is a return forecast distribution. I specify that agents assume that log returns are i.i.d. binomially distributed, with mean m and standard deviation s.[6] We have already seen that in the rational expectations equilibrium, returns are i.i.d. with standard deviation σ and mean μ given by:

$$\mu \equiv E \log(Z_t) = \log(1 + \delta^*) - \log \delta^* + d.$$

If $m = \mu$ and $s = \sigma$, then returns are indeed i.i.d., so agents' beliefs are correct, and outcomes are the same as under rational expectations.

Following Sargent (1999), I adopt the notion of a self-confirming equilibrium (SCE) to describe such a situation.[7] In a SCE, agents optimize given their beliefs and their beliefs are correct *along the equilibrium path*, but may be incorrect for events not observed in

[6] As discussed above, very similar behavior results if shocks are normal and agents have a subjective normal distribution.

[7] The self-confirming equilibrium concept arose in game theory, as in Fudenberg and Levine (1998). Sargent (1999) formulated the concept in a parametric form for adaptive learning models.

equilibrium. Therefore $(m, s) = (\mu, \sigma)$ is a SCE in this model, and in fact is the unique one. When $m \neq \mu$ and/or $s \neq \sigma$ then equilibrium returns exhibit temporal dependence due to changes in net asset positions. In such cases, agents' beliefs are incorrect, failing to account for the serial correlation in returns. Nonetheless, one can still solve for an asset market equilibrium with incorrect beliefs, as I do now.

3.2 Equilibrium for Arbitrary Beliefs

With general subjective beliefs $\theta = (m, s)$, agents solve the consumption portfolio choice problem. As discussed more below, I impose a portfolio constraint $0 \leq v_t \leq \overline{v}$. The upper bound reflects the leverage constraint, and since the equilibrium requires domestic agents to hold stocks, I rule out short sales as well. The decision rules follow from Hakansson (1970), and have the same form as before: $C_t = c(\theta)X_t$, and $v_t = v(\theta)$.[8] Agents consume a constant fraction of their wealth X_t and invest a constant fraction in stocks, where the fractions depend on their perceptions of the stock returns. Expressions for the decision rules are given in Appendix A.2. In the absence of binding constraints, the optimal fraction invested in stocks increases in expected returns m and decreases in perceived risk s, as expected.

I now solve for an equilibrium. The value of stock holdings at date t is:

$$P_t S_t = (1-c)vX_t = (1-c)v(RB_{t-1} + S_{t-1}(P_t + D_t)),$$

where I suppress the dependence on the beliefs θ. Therefore the stock price must satisfy:

$$P_t = \frac{(1-c)v(RB_{t-1} + S_{t-1}D_t)}{S_t - (1-c)vS_{t-1}}.$$

Since prices and dividends are growing over time, it is convenient to focus on the price-dividend ratio. Imposing the equilibrium conditions that the agent always holds the entire supply of stocks, $S_t = S_{t-1} = 1$, the equilibrium price-dividend ratio is then:

$$\frac{P_t}{D_t} = \frac{(1-c(\theta))v(\theta)}{1-(1-c(\theta))v(\theta)} \frac{RB_{t-1} + D_t}{D_t} \equiv \delta(\theta)\omega_t. \tag{6}$$

The first fraction on the right defines the pricing function $\delta(\theta)$ and the second defines the relative net asset position $\omega_t = 1 + RB_{t-1}/D_t$. The pricing relation (6) holds as long as it is well-defined and nonnegative, which is ensured by the portfolio constraints specified later.

In the zero-debt rational expectations equilibrium, $\delta(\theta) = \delta^*$, stock prices are driven entirely by dividends, the relative asset position is $\omega_t \equiv 1$, and the price-dividend ratio is constant. But in general, agents can take either long or short positions in bonds, up to the portfolio limits. Thus the total demand for stocks depends on the net asset position ω_t, and

[8] Hakansson (1970) had no portfolio constraints, but it straightforward to show that his results extend.

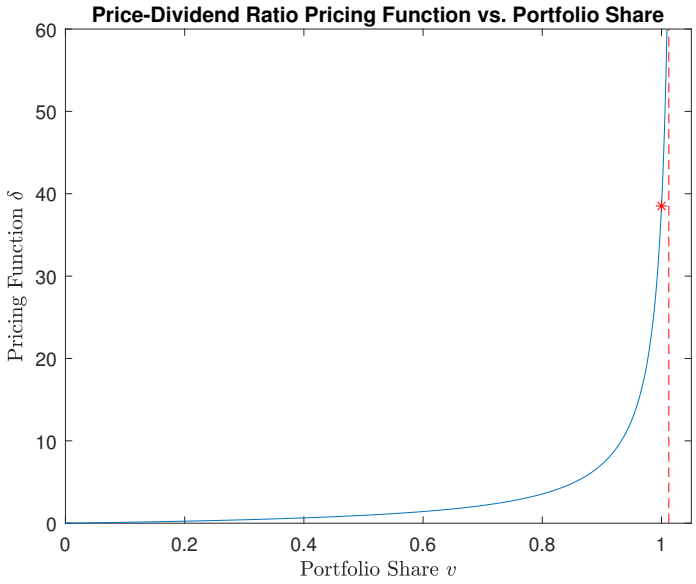


Figure 1: The pricing function δ versus the portfolio share v for fixed beliefs $\theta = (\mu, \sigma)$. The SCE level of $v = 1$ is shown with an asterisk $*$, while the dashed line is the upper bound $\overline{v}$.

this is reflected in prices and returns. Later I introduce learning, where belief variation also contributes to variation in prices. Expressions for stock returns Z_t and the dynamics of the net asset position ω_t with arbitrary beliefs are given in Appendix A.2.

3.3 Portfolio Constraints

In the model stock prices are very sensitive to the risky-asset portfolio share v. This is particularly true for v near the self-confirming and rational expectations equilibrium value of unity. As agents invest almost their entire wealth in stocks, or even borrow to increase their leverage in risky assets, the stock price increases dramatically. This effect is illustrated in Figure 1, which plots the pricing function δ as a function of the portfolio share v. In particular, I fix beliefs at the SCE values $\theta = (\mu, \sigma)$ and vary the portfolio share v.

The figure shows that as $v \to \frac{1}{1-c}$ the stock price explodes, with $\delta \to \infty$. While the pricing function is finite for portfolio weights slightly above unity, there is a limit to how much leverage the economy can sustain. As I show in Appendix A.2, with sufficient leverage the demand for stocks outstrips any fixed supply at a finite price. How close the divergence threshold is to the rational expectations equilibrium depends on the consumption/wealth ratio. Even with a relatively high value of $c = 0.1$, the prices diverge at $v = 1.11$.

Moreover, stock prices are very sensitive to changes in the portfolio share near the SCE value of $v = 1$. Since the pricing function δ has a steep slope in this region, very small changes in portfolio positions lead to large changes in stock prices. In particular, I show in

Appendix A.2 that the derivative at the pricing function $\delta(v)$ at the SCE value of $v = 1$ is $\delta'(1) = \delta^*(1 + \delta^*)$. As discussed below, the mean price/dividend ratio in the data is approximately 29, so matching this value would imply a slope of 870 at $v = 1$. To first order, 1 percentage point fall in the portfolio share to $v = 0.99$ would lead to a 30% drop in the price dividend ratio to 20.3.

Beyond ensuring a finite pricing function, equilibrium requires that the representative agent can repay debt, so that $RB_{t-1} + D_t \geq 0$ in all states of the world. This in turn guarantees that the net-asset position ω_t is nonnegative, and so the price P_t is as well. Along with simplicity, this is a primary reason I assumed a dividend process with finite support.[9] In Appendix A.2, I show that such a debt limit can be written as a constraint on the portfolio share $v \leq \overline{v}$. I allow for the loosest portfolio restriction consistent with equilibrium.

If the portfolio restriction failed, agents would not be able to repay the principal and interest on their debt. This is a situation Minsky (1986) refers to as "speculative finance," where investors need to roll over their debt liabilities, and thus the stock of debt would increase. Minsky emphasizes how this can lead to fragility of the financial system. Minsky (1986) also discusses a further stage of "Ponzi finance," where agents cannot repay event the interest on their debt and so would need to sell off assets. In my setting with one-period debt, principal must be repaid each period, so these speculative and Ponzi stages merge. If $RB_{t-1} + D_t < 0$, agents would want to sell their stock holdings in order to repay their debt. But in a representative agent equilibrium they must hold the entire unit supply. This need for liquidity would drive the stock price to zero.

4 Adaptive Learning

Having determined equilibrium prices and returns for fixed beliefs, I now describe how agents update their beliefs over time. I assume that agents learn adaptively, meaning that when they form their decision rules agents treat their beliefs as constant, but then they update their beliefs over time. This approach, which rules out active experimentation, is what Kreps (1998) calls "anticipated utility," and it dominates the literature on adaptive learning in macroeconomics. By contrast, a fully Bayesian approach to learning would forecast not only future outcomes but future beliefs, and so requires a higher degree of rationality than under full information rational expectations. Kreps (1998) recommends the anticipated utility approach as a normative description of how people solve uncertain dynamic decision problems, where they construct a model, act, and adapt. Cogley and Sargent (2008) discuss how the fully Bayesian approach is often intractable, while the simpler anticipated utility

[9] With unbounded dividend growth, if agents began a period in debt, there would be a nonzero probability that the dividend realization would be too low for repayment.

approach may provide a good approximation. Anticipated utility is also consistent with Minsky (1975), who argues that, "Financing is often based on an assumption 'that the existing state of affairs will continue indefinitely', but of course this assumption proves false."

At each date I solve for the equilibrium prices with $\theta = \theta_t = (m_t, s_t)$, the agents' current estimates of the mean and standard deviations of returns. Then after observing the realized returns, agents update their beliefs. Since agents estimate means and variances, I assume that they update their beliefs using a recursive procedure for estimating the sample counterparts:

$$m_{t+1} = m_t + \varepsilon(\log(Z_t) - m_t), \tag{7}$$

$$s^2_{t+1} = s^2_t + \varepsilon([\log(Z_t) - m_t]^2 - s^2_t). \tag{8}$$

Here $\varepsilon > 0$ is a constant known as the gain, reflecting the weight on new information.

Setting $\varepsilon = \frac{1}{t+1}$ would reflect a subjective belief that the mean and volatility of returns are constant, so more observations provide more information about the unchanging underlying parameters. But if agents believe that the return process changes over time, then more recent data would be more informative than observations further in the past. Constant gain learning rules discount past observations, allowing beliefs to better track moving parameters. Sargent and Williams (2005) showed that constant gain learning rules are Bayesian optimal filters when the underlying parameters follow a random walk.[10] Alternatively, if agents are uncertain about the return process, they may opt to use learning rules which are robust to model misspecification. Evans, Honkapohja, and Williams (2010) show that constant gain learning rules are robust optimal filters. Both of these motivations are reflected in discussion by Minsky (1989) states, "[M]odels of system performance that help form the expectations of businessmen and bankers are affected by the recent performance of the economy... The critical agents are unsure how the economy will perform, because they are unsure of the effect of recent institutional and environmental changes."

The dynamics of beliefs make the economy self-referential. Based on their beliefs, agents make decisions that determine equilibrium prices and returns, which in turn affect future beliefs. Using the expressions for the pricing functions and dynamics of the asset position ω_t in Appendix A.2, the feedback between beliefs and outcomes is apparent:

$$\log Z_t = \log \frac{1 + \delta(\theta_t)(1 + q(\theta_{t-1})\omega_{t-1}\chi_t)}{\delta(\theta_{t-1})\omega_{t-1}\chi_t} \equiv \log Z(\theta_t, \theta_{t-1}, \omega_{t-1}, \chi_t) \tag{9}$$

$$\omega_t = 1 + q(\theta_{t-1})\omega_{t-1}\chi_t. \tag{10}$$

Here the weighting function $q(\theta)$ is a combination of the pricing function and decision rules specified in the appendix. The second equality in (9) defines the stock return function $Z(\theta_t, \theta_{t-1}, \omega_{t-1}, \chi_t)$, which makes clear how returns at date t depend on current and lagged

[10] Malmendier and Nagel (2016) show that constant gain learning explains average survey expectations.

beliefs, net asset positions ω_{t-1}, and the fundamental dividend shocks χ_t. The economy is then governed by joint dynamics of beliefs (7)-(8) and outcomes (9)-(10).

In the Appendix B.1, I apply some well-known results from the learning literature to show that in the limit as the gain ε goes to zero, the beliefs can be approximated by a differential equation whose stable point is the self-confirming equilibrium. This suggests that for small gain, agents' beliefs and the price-dividend ratio should be concentrated around the self-confirming levels. However, as I discuss next, simulations of the learning model suggest behavior very different from stability and convergence.

5 Financial Instability

Figure 2 shows a representative simulation of the calibrated baseline model. The top panel plots the price-dividend ratio from a long simulation, while the lower panels show a closeup on a typical boom-bust cycle (shown in the shaded area). The model displays repeated episodes of booms and busts in asset prices, driven by buildups, spikes, and collapses in asset flows. These flows in turn are driven by changes in agents' beliefs. Beliefs change rapidly during the boom-and-bust period, but this episode slowly fades from memory during more tranquil times. Periods of stability lead to forgetting, which brings on periods of instability.

In more detail, the top panel of Figure 2 plots the realized equilibrium price-dividend ratio in a long sample of over 1,000 periods (years) under the calibration described in Section 7 below. The model generates recurrent booms and crashes, which occur regularly but at random intervals. During booms, the price-dividend ratio can increase by 30-50% with the span of a years, then crash by a similar or larger magnitude in a short span. Below I show that the model can match the mean and variance of the observed price-dividend ratio, and can come close to matching the frequency of large market crashes.

The lower panels zoom in on the boom-bust cycles in the shaded region to better understand the dynamics of the model. The panels in the middle row show beliefs m_t and s_t, along with the portfolio share invested in stocks v_t. The bottom left panel shows the log stock return Z_t, which agents forecast in the model. The final two panels decompose the decompose the movements in the price-dividend ratio. In equation (6) the price-dividend ratio is the product of the pricing function $\delta(\theta_t)$, which is shown in the bottom center panel, and the net asset position ω_t, shown in the bottom right. In the SCE, ω_t is identically equal to one. But as the agent accumulates bonds, ω_t is greater than one which acts as a multiplier on how beliefs affect prices through the pricing function $\delta(\theta)$.

The shaded region begins in period 240 in the aftermath of a crash. From the convergence results in Appendix B.1, we expect agents' beliefs (m_t, s_t) to converge to the self-confirming levels and the portfolio share v_t to converge to one. We observe this general trend for roughly

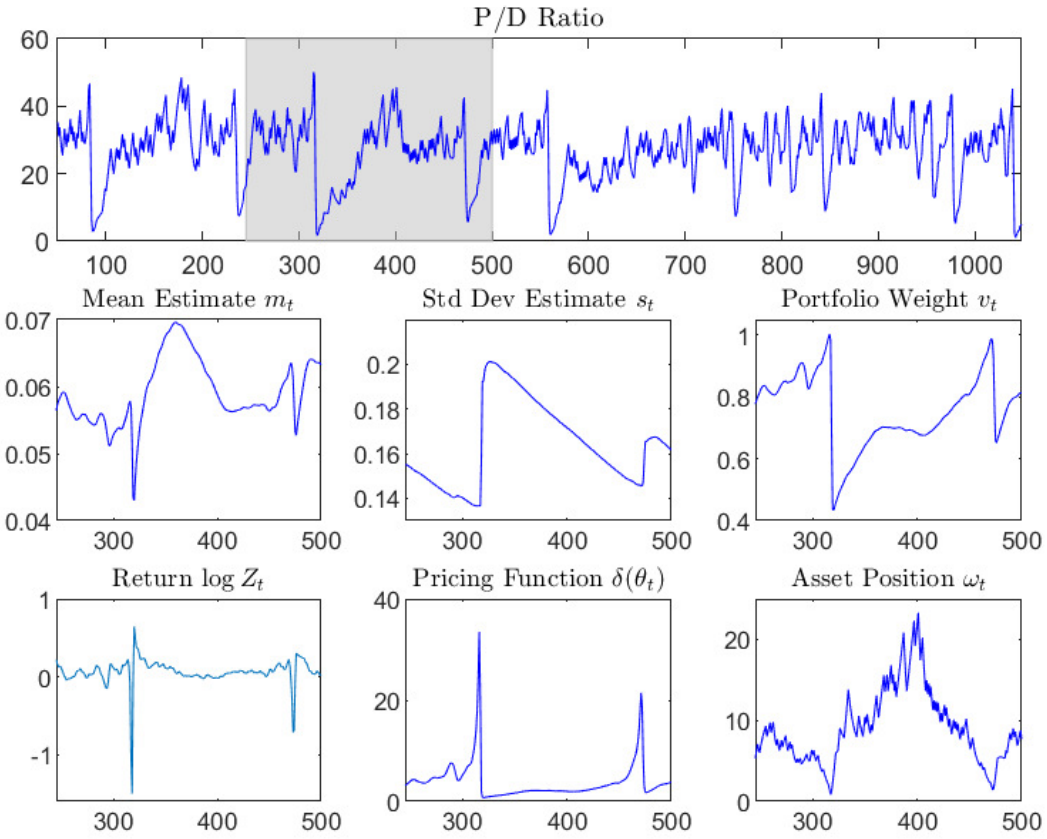


Figure 2: Representative simulation for the baseline calibration. The top left panel plots a long simulation of the price-divided ratio, with the lower panels focusing on the shaded region. The middle plots the beliefs m_t and s_t, along with the portfolio share v_t. The bottom row shows the log return Z_t, the pricing function $\delta(\theta_t)$, and the net asset position ω_t.

the next 50 periods. Volatility estimates start high, as agents had recently been through a period of large fluctuations, and they decline monotonically over this span. Overall, beliefs about volatility move more and have a larger impact than expected returns. The periods following a crash see a quick recovery in prices, so mean estimates are also high and drift downward slowly over this period. Driven mostly by the declining volatility estimates, agents increase their stock holdings v_t over time. Throughout this period, price-dividend ratios increase, driven by increases in the pricing function $\delta(\theta_t)$, which has a large multiplier effect through ω_t, as investors have accumulated bonds that they slowly sell. Investors slowly forget the experience of a market crash, as a period of stable and gradually increasing prices makes them more willing to invest in stocks. As they do, they bid up asset prices. Since agents are

backward looking, these higher realized returns lead to increases in expected returns.

Around period 300 the trends in the economy accelerate. As more funds flow into stocks, asset prices increase rapidly. After trending downward, the expected return m_t starts to increase, and the portfolio share v_t increases to near unity. The pricing function δ spikes as the portfolio share v_t increases into the steeply sloped region from Figure 1. The pricing function $\delta(\theta_t)$ first doubles in the span of a couple of periods, then doubles again in a couple more periods. However the net asset variable ω_t declines as bond holdings fall, dampening the impact of the increase in δ. The price-dividend ratio increases from 30 to 50 in the span of a few periods, before crashing to values around 2.

While the asset price booms are generated by flows into the stock market, the market crashes when the inflow stops, which happens for two main reasons. If the portfolio share v_t hits the constraint $\bar{v}$ during the market boom, the inflow of funds into stocks will slow or stop, generating a negative return surprise. This leads agents to revise down their estimate of mean returns and thus cut v_t. However binding constraints are rare in the calibrated model, driving only the final crash in the top panel of the Figure 2.

More surprisingly, rapid asset price booms can directly lead to market crashes by increasing uncertainty. Agents trade off the mean and variance of returns in their investment decisions, so if the perceived variance increases more rapidly than expected return, they cut the risky-asset portfolio share. During the boom, the economy enters the steeply-sloped region in Figure 1, so small changes in portfolio shares driven by either cause lead to large price changes. Asset flow reversals have a large impact and lead to abrupt crashes.

Elevated subjective uncertainty causes all of the other crashes Figure 2, including those in the closeup shaded area. The portfolio share peaks at a value just over one in period 316, which is below the constraint. This is also the period which marks the peak in the pricing function $\delta(\theta)$, with the price-dividend ratio peaking in the previous period. The rapid run-up in prices leads to an increase in the estimated variance, which leads to a slight cut in portfolio share in the next period. While the portfolio change is small (from 1.003 to 0.985), it leads to drop of over 60% in the price-dividend ratio, reflecting the steepness of the pricing function. After the recent boom, this drop in prices is influential. The estimated volatility increases sharply, leading to a further cut in the portfolio share. The desire to withdraw funds from the stock market further depresses prices. In the span of a couple of periods the stock market crashes, bottoming out at a portfolio share slightly below $v_t = 0.45$ and a price-divided ratio around 2, a decline of over 95% from its peak.

Figure 2 shows that there is a distribution of crash severity. This particular crash in period 318 is a severe one, but crashes of similar magnitude occur in periods 86, 560, and 1041. Continuing in the shaded area, after bottoming out in period 318, the price-dividend ratio recovers substantially in the following couple of periods, which leads to a spike in returns

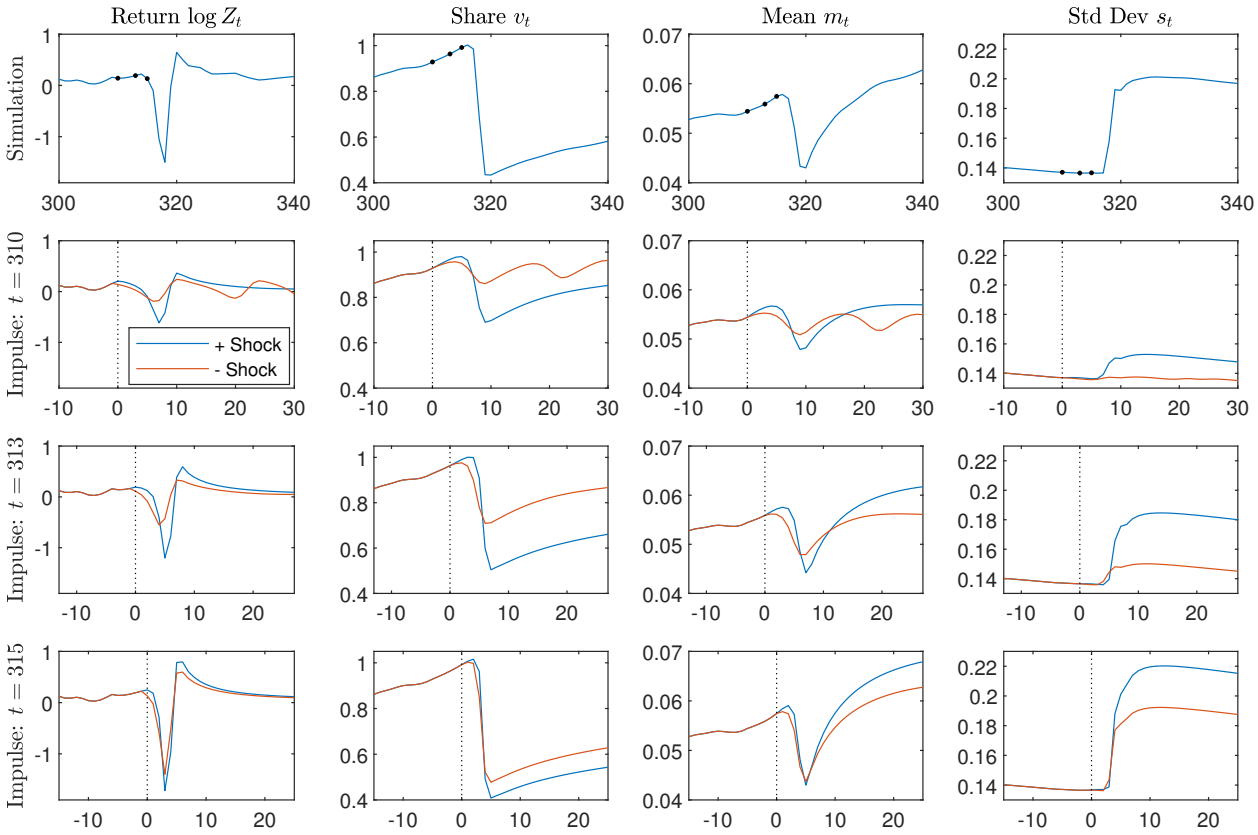


Figure 3: Representative simulation (top row of panels) and conditional impulse responses at different points ($t = 310$, 315, and 318 in the following rows) for the baseline calibration.

Z_t (bottom left panel), which in turn drives an increase in expected returns m_t (center-left) and volatility s_t (center-right). Slowly, the experience of the crash fades, assets in the stock market build up again, and another boom begins. This boom-bust cycle is less severe, with the portfolio share peaking at 0.99 and bottoming out at 0.65, but still quite strong.

Figure 3 illustrates the impact of stochastic shocks and endogenous model dynamics. The top row of panels show the log of the stock return Z_t, beliefs m_t and s_t, and the portfolio share v_t from part of the boom-bust cycle highlighted in Figure 2. The lower panels show conditional impulse responses starting at different dates into the asset market boom – periods 310, 313, and 315 – marked with dots in the top panels. For example, the second row takes the state of the model in period 310, feeds in a positive (blue line) and negative (red line) dividend shock for that period only, and traces out the dynamics in the absence of further shocks. The third and fourth rows do the same for periods 313 and 315.

In period 310, midway through the boom phase, we see very different responses to positive and negative shocks. A negative shock leads to dampened fluctuations, with no substantial boom or crash, and little change in beliefs. But a positive shock reinforces the dynamics in the

ongoing market boom, with the portfolio share continuing to increase for several periods, before dropping rapidly, leading to larger negative returns and an increase in perceived volatility s_t. Thus at this point in the boom, shocks are influential in influencing the market dynamics. A few periods later at $t = 313$, the simulation had seen a sequence of positive shocks, so the market boom had continued. Now the positive and negative impulses have qualitatively similar dynamics, with differing magnitudes. The positive shock leads to a larger ongoing boom and a more severe crash, which in turn drives a much larger increase in volatility s_t. A negative shock now leads to a less severe boom and bust. By the time we get to $t = 315$, when the market boom has continued, the differences in the impulses are minimal. The force of the endogenous model dynamics is stronger, so there is less variation due to the shocks. The boom continues for only a couple of periods before a severe crash.

The figures illustrate the interaction of stochastic shocks and the endogenous nonlinear dynamics in asset price fluctuations. Shocks play a key role in the timing and magnitude of market fluctuations, but the endogenous belief changes and associated model dynamics ultimately drive the boom-bust cycles. The sharp price movements are driven by relatively small initial changes in beliefs, as booms lead the economy into the fragile region of very responsive asset prices. Periods of stability lead to periods of instability, and rapid booms lead to market crashes, as in Minsky's theory. In the next section I theoretically characterize these model dynamics, later exploring the quantitative implications in more detail.

6 Dynamic Instability Drives Financial Instability

While the self-confirming is stable in the small gain limit, the cyclical behavior just illustrated seems to be at odds with this convergence result. I reconcile these apparently contradictory results in this section by showing that the SCE is only stable for very small gain. As the gain increases, the equilibrium has a bifurcation where it loses stability, and a stable limit cycle arises. This bifurcation occurs at a threshold gain value which is quantitatively very small, much smaller than empirically plausible values. I also show that the economy may experience long periods of volatility before finally settling down to the stable limit cycle.

6.1 Discrete Nonlinear Dynamical System

In this section, I analyze the model as a discrete-time nonlinear dynamical system, studying the dynamics of the model with a fixed gain. Most of the learning literature, including the convergence result in Theorem 2 in the appendix, is based on small gain limits. I show that such small gain analysis is misleading in this model, as the radius of stability is quantitatively

very small. The analysis in this section are novel in the learning literature.[11]

The dynamic system governing the economy is given by joint dynamics of beliefs in equations (7)-(8) and outcomes in (9)-(10). . The convergence results take small gain limits, and therefore make a sharp distinction between the beliefs and the state variables. As the gain ε shrinks, the dynamics of the state are "fast" in comparison to the beliefs. Therefore to analyze the beliefs, we can average over the state evolution.

But for my model (and others like it), such an approximation is only valid for very small gain ε. For larger (but still empirically small) gains, the state and belief dynamics are more closely intertwined. Therefore I now stack the current beliefs, lagged beliefs, and state into a larger vector $X_t = [\theta_t, \theta_{t-1}, \omega_{t-1}]'$.[12] Then the second-order dynamics of the combined beliefs and outcomes from (7)-(10) can be written as a first-order vector system:

$$X_{t+1} = G(X_t, \chi_t). \tag{11}$$

Direct analytic characterization of the discrete-time nonlinear stochastic difference equation system (11) is difficult. To provide analytic tractability on the expected behavior of the system, and particularly its behavior local to the self-confirming equilibrium, I study the dynamics of the averaged system:

$$X_{t+1} = \bar{G}(X_t) = \int G(X_t, \chi) dF(\chi), \tag{12}$$

where I average over the i.i.d. stochastic shocks χ. In Appendix B.2 I provide more discussion and show that, for small shocks, the deterministic averaged system characterizes the mean of the stochastic system in a neighborhood of the SCE. This small noise approximation employs less averaging than in the small gain limit, which also averages over the dynamics of the state variable ω_t. The feedback between beliefs θ_t and the state variable ω_t, which captures the net asset positions and governs assets flows, drives the boom-bust cycle. Averaging over the net asset position misses this key interaction.

I now apply methods for discrete-time dynamical systems to analyze the averaged system (12). With the SCE beliefs $\overline{\theta} = (\mu, \sigma)$, denote the SCE state $\bar{X} = [\overline{\theta}, \overline{\theta}, 1]'$. The SCE state is a fixed point of the averaged system:

$$\bar{X} = \bar{G}(\bar{X}).$$

Applying standard linearization techniques, the local stability of the SCE state $\bar{X}$ is determined by the Jacobian $\bar{G}'(\bar{X}) = \int \frac{\partial G}{\partial X}(\bar{X}, \chi) dF(\chi)$. In other words, locally around the SCE, the behavior of the dynamical system (12) is determined by its linearization.

[11] Cho (2018) and Chien, Cho, and Ravikumar (2020) show that small-gain limits in models with second-order dynamics may give misleading stability results. Evans and Honkapohja (2009) and Honkapohja and McClung (2023) documented similar examples of instability with small gain.

[12] Using equation (9), we write the return Z_t in terms of current and lagged beliefs, net assets, and shocks.

I now show that when the gain ε is small enough, all eigenvalues of the Jacobian are inside the unit circle. For small gain the SCE is stable, as in the convergence results. But I also show that there is threshold level of the gain $\overline{\varepsilon}$ where the SCE loses its stability. At $\overline{\varepsilon}$ the system has a pair of complex eigenvalues on the unit circle, and for $\varepsilon > \overline{\varepsilon}$ the SCE is surrounded by a stable, closed invariant curve. This is a Neimark-Sacker bifurcation, the discrete-time equivalent of a Hopf bifurcation in continuous time. For sufficient weight on new information, the SCE is unstable, but the system has a stable limit cycle.

The following theorem collects results from Kuznetsov (1998), establishing the existence and local properties of the two-dimensional center manifold from the five-dimensional dynamics for X_t, and providing conditions for the stability of the limit invariant curve. To emphasize the dependence of the dynamics on the gain, I denote the key Jacobian as $J(\varepsilon) = \bar{G}'(\bar{X}; \varepsilon)$. The proof of the following is in Appendix B.2.

Theorem 1 *The averaged dynamics of the learning model (12) has a fixed point $\bar{X}$ corresponding to the SCE state. There exists a threshold value of the gain $\overline{\varepsilon}$ such that:*

1. *For $\varepsilon < \overline{\varepsilon}$, all eigenvalues of the Jacobian matrix $J(\varepsilon)$ are inside the unit circle, and hence the SCE state $\bar{X}$ is locally asymptotically stable.*

2. *At $\varepsilon = \overline{\varepsilon}$, the Jacobian $J(\varepsilon)$ has a pair of complex eigenvalues on the unit circle. Denote these eigenvalues as $\lambda_{1,2} = e^{\pm i\psi_0}$, for some real value ψ_0.*

3. *For sufficiently small $\varepsilon > \overline{\varepsilon}$, denote the pair of complex eigenvalues as $\lambda_{1,2}(\varepsilon) = r(\varepsilon)e^{\pm i\psi(\varepsilon)}$, where $r(\overline{\varepsilon}) = 1$ and $\psi(\overline{\varepsilon}) = \psi_0$. Assume:*

   (a) $r'(\overline{\varepsilon}) \neq 0$,

   (b) $e^{ik\psi_0} \neq 0$ for $k = 1, 2, 3, 4$, and

   (c) the term $a(\overline{\varepsilon})$ defined in equation (19) below satisfies $a(\overline{\varepsilon}) < 0$.

   Then there is a neighborhood of $\bar{X}$ in which a unique, closed, and attracting invariant curve bifurcates from $\bar{X}$ as ε passes through $\overline{\varepsilon}$.

Part 3 of the theorem gives the existence and stability of the limit cycle. The nondegeneracy conditions in part 3(a) and 3(b) allow for local inversions and ensure existence. The stability condition in part 3(c) is based on a higher-order local approximation around the SCE, since the system is neutrally stable with eigenvalues on the unit circle at $\overline{\varepsilon}$. In practice, I need to verify these conditions numerically, as they depend on matrices of derivatives evaluated at the SCE, as well as their eigenvalues and eigenvectors.

Figure 4 shows the maximal eigenvalue modulus from the Jacobian $J(\varepsilon)$ evaluated at the SCE versus the learning gain ε. The threshold gain for instability is approximately

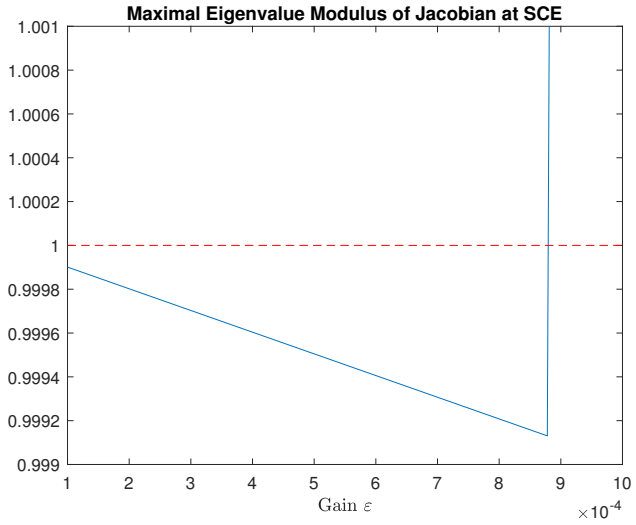


Figure 4: Maximal eigenvalue modulus of the Jacobian $J(\varepsilon)$ at the SCE versus the gain ε.

$\overline{\varepsilon} = 8.8 \times 10^{-4}$. This value is very small, as the calibrated value of $\varepsilon = 0.0051$ discussed below and shown in the figures is almost six times larger. For gains below $\overline{\varepsilon}$, all eigenvalues are inside the unit circle, as in part 1 of the theorem. As in part 2 of the theorem, exactly at $\overline{\varepsilon}$, the Jacobian has a pair of complex eigenvalues on the unit circle.[13] By varying ε, it is also straightforward to verify numerically that the nondegeneracy and stability conditions (a)-(c) of part 3 the theorem hold. Thus as the gain passes through $\overline{\varepsilon}$, there is a closed invariant curve which bifurcates from the SCE, leading to limiting cyclical dynamics.

Figure 5 illustrates how the bifurcation arises by analyzing the averaged dynamics for two different gain settings on either side of the threshold $\overline{\varepsilon}$. The top row of panels shows the beliefs and portfolio shares for simulations of the averaged dynamics for $\varepsilon = 5 \times 10^{-4}$ and 9×10^{-4}. The top left panel shows the trajectories of beliefs over long simulations of one million periods, where both simulations are initialized from a common point (marked with an asterisk) in a neighborhood of the SCE (marked with a circle). In both cases the beliefs follow non-monotone dynamics as they converge back toward the SCE, with large movements in the estimated volatility s_t. The smaller gain is below the bifurcation threshold, so the beliefs converge back to the SCE. The larger gain is over the stability threshold, so the beliefs enter a cycle, shown as the solid mass of points at the bottom of the figure.

The top right panel shows the portfolio shares from these simulations, focusing on the first 500 periods. In both cases, the non-monotonicity in beliefs is associated with an initial cut in the portfolio share v_t from near one to a lower level, followed by a rise. The smaller, stable gain shows a smaller drop in the portfolio share, and rapid convergence back to the

[13] Precisely, the eigenvalues are $\lambda_{1,2} = (0.5119 + 0.8590i, 0.5119 - 0.8590i) = e^{\pm 1.0334i}$, so $\psi_0 = 1.0334$.

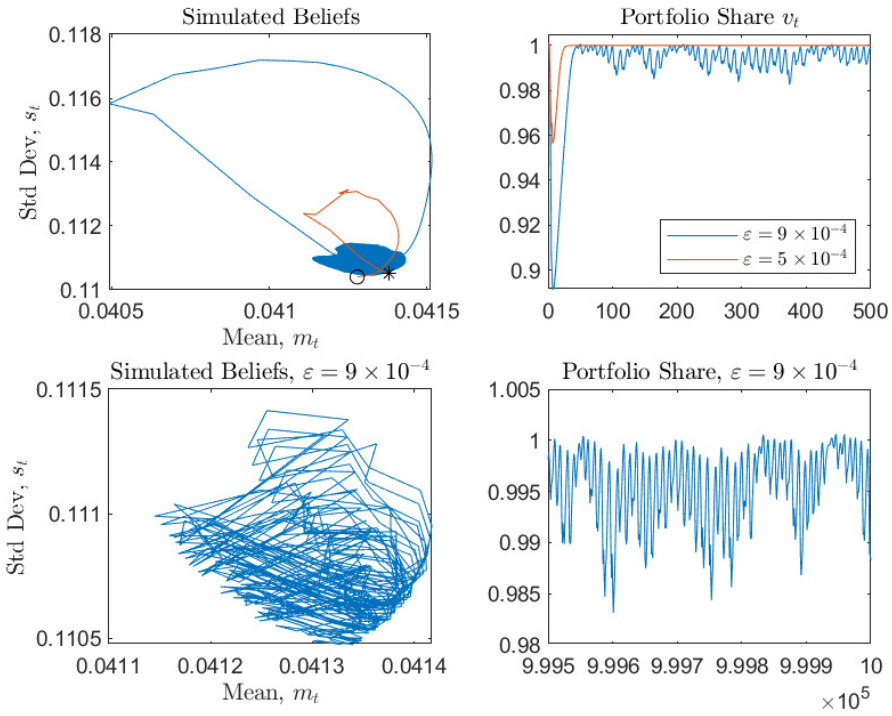


Figure 5: Beliefs and portfolio shares from the averaged dynamics for gains on either side of the stability threshold. Top left: simulated beliefs from 1 million periods initialized at the same point (marked '*', while the SCE is marked 'o'). Top right: the resulting portfolio shares for the first 500 periods. Bottom left: beliefs for the larger gain over the last 500 periods. Bottom right: portfolio share for the larger gain over the last 500 periods.

SCE value of one. There is a larger initial drop with the larger, unstable gain, and an increase back toward one. But rather than converging to this SCE value, the portfolio share oscillates over a relatively small range as the economy enters a limit cycle. To illustrate the cycle more clearly, the bottom row of the figure shows the outcomes for the end of the sample for the larger, unstable gain. The left panel shows the belief trajectory for the last 500 of the one million periods, showing that the beliefs follow an irregular cyclical path. The magnitude of the variation in this limit cycle is relatively small, as the bottom right panel shows that the portfolio share oscillates over a small range just below unity.

For even very small gains, learning about risk and return leads to cycles in beliefs, and hence outcomes. For gains close to the threshold $\bar{\varepsilon}$, these cyclical fluctuations are quite small in magnitude. But as the gain increases, the limit cycle dynamics become more pronounced and drive larger fluctuations in asset prices. Figure 6 shows the same plots as in the previous figure, but now for the calibrated gain $\varepsilon = 0.0051$ over ten million periods.

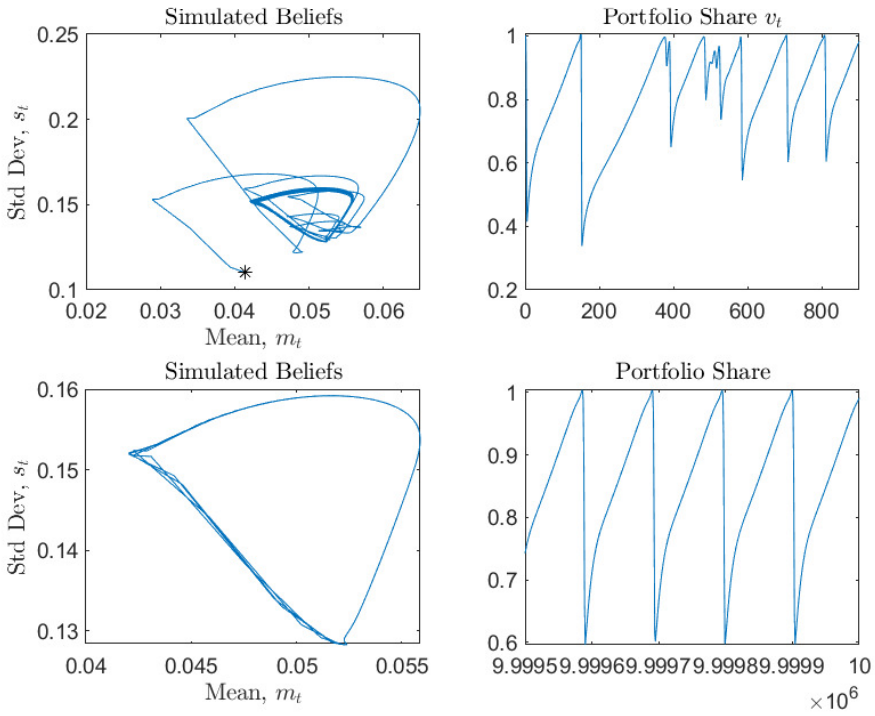


Figure 6: Beliefs and portfolio shares from the averaged dynamics for the calibrated gain setting of $\varepsilon = 0.0051$. Top left: simulated beliefs from 10 million periods initialized at a point (marked '*') near the SCE. Top right: portfolio share for the first 900 periods. Bottom left: beliefs over the last 500 periods. Bottom right: portfolio share over the last 500 periods.

The top left panel shows that the beliefs experience larger fluctuations for the first several hundred periods, looping around far from the SCE, before settling into a stable limit cycle. These belief fluctuations lead to large changes in the portfolio share in the top right panel. The share v_t experiences a couple dramatic falls early in the sample, then has a period of erratic fluctuations, before settling into a regular cycle. The closed invariant curve of the limit cycle is clear in the bottom left panel, which shows the trajectory of beliefs over the last 500 periods. The belief cycle drives a regular cycle in outcomes, with the portfolio share (bottom right) recurrently collapsing to 0.6, recovering to near 1.0, and then repeating.

Even at gain settings that are quantitatively very small, the model exhibits instability the drives persistent booms and busts in asset prices. With sufficient weight on new information, learning from asset returns generates recurrent cycles. The deterministic averaged system displays notable volatility, which is more pronounced before it settles into the limit cycle. Thus it is no surprise that the stochastic system, which is perturbed from these smooth cyclical dynamics, displays persistent volatility marked by recurrent booms and busts.

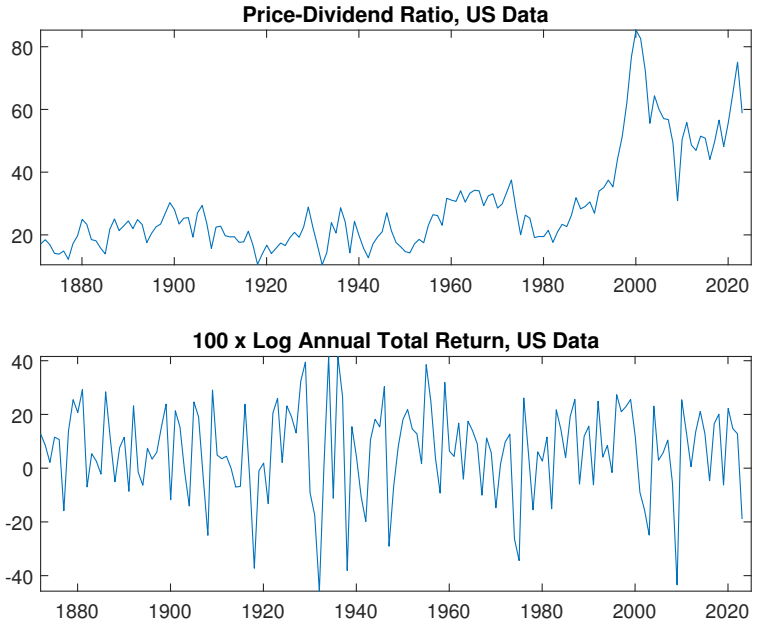


Figure 7: Top panel: Price-dividend ratio (annual) in US data. Bottom panel: Log annual total return on stocks. Both from the Shiller dataset.

7 Implications of Instability

I now explore in more detail the quantitative implications the model, showing that it can help explain observed asset pricing data. The model can generate an empirically plausible mean and variance of the price-dividend ratio, and a plausible frequency of market crashes.

7.1 Data

To calibrate the model, I use Robert Shiller's data set on the S&P 500 prices and dividends.[14] The calibration is based on data from 1871-2022 at an annual frequency. The empirical price-dividend ratio, shown in Figure 7, is very volatile, with a mean of 28.97 and standard deviation of 15.53. The calibration below uses the log price-dividend ratio, which has a mean of 3.25 and standard deviation of 0.46. Although the secular increase in the P/D ratio around the year 2000 swamps much of the variation in the earlier sample, declines of 20% or more in short spans are common. The bottom panel plots the log total return, the data counterpart of $100 \times \log Z_t$, also at an annual frequency. These returns are highly variable, with a mean of 6.51 and a standard deviation of 16.90, with multiple booms and crashes with moves of 20 percentage points or more in either direction. Dividend growth is the driving exogenous process in the model. In Shiller's data, the mean and standard deviation

[14]The data was accessed from: http://www.econ.yale.edu/∼shiller/data.htm.

Statistic	Data	Baseline	Closed (Base)	Closed (Opt)	REE	Two Gains
$E(\log P/D)$	3.25	3.16	3.70	3.21	3.02	3.43
$\sigma(\log P/D)$	0.46	0.46	0.15	0.35	0	0.50
$100 \times E(\log Z)$	6.51	6.44	4.10	6.50	6.31	5.32
$100 \times \sigma(\log Z)$	16.90	16.93	11.17	17.90	11.04	18.43
Error	–	0.031	0.706	0.236	1.365	0.234
β	–	0.979	0.979	0.903	0.960	0.975
γ	–	3.259	3.259	4.721	3.000	1.174
ε	–	0.0051	0.0051	0.0276	–	0.0010 0.0103
Crash Frequency	25.25	38.30	–	422.35	–	27.30
$100 \times E(\log \Delta C)$	0.86	1.58	1.58	1.55	1.56	1.56
$100 \times \sigma(\log \Delta C)$	1.27	10.88	11.18	17.76	11.04	13.90

Table 1: Summary statistics from data and simulations of the baseline model, the closed economy version with the baseline parameters, the closed economy version with optimized parameters, the model under rational expectations, and a version with two learning gains.

of annual dividend growth rates are $100 \times (d, \sigma) = (1.56, 11.104)$. Since dividend growth is relatively volatile, the model is able to generate larger risk premia, even under rational expectations, compared to a calibration based on aggregate consumption data. The obvious downside, discussed below, is that the consumption in the model is overly volatile.

7.2 Calibration and Fit

Table 1 summarizes that data and the results of calibrating different versions of the baseline model. The first column provides the four calibration targets: the mean and standard deviation of the log price-dividend ratio and the stock returns from the Shiller data set. To calculate the corresponding statistics from the simulated models, I run 1,000 simulations of 2,000 periods each. I discard the first 1,000 periods of each run in computing the sample statistics, averaging over the 1,000 runs.

I consider optimized parameterizations of different versions of the model: the baseline model (open economy with a fixed interest rate), a closed economy version (where the interest rate adjusts each period to clear the bond market with zero net supply), and the rational expectations version of the model.[15] To illustrate the role of the portfolio adjustment and

[15]The open economy version fixes R and uses the optimality conditions to determine v. The closed economy version has $v = 1$ in equilibrium, so the optimality conditions determine R.

the open versus closed economy, I first report the results for the baseline parameterization in a closed economy, then optimize over the parameters to best fit the data using the closed economy model. The final column considers a variation with separate gains for the mean and variance discussed in Section 7.6 below.

For any parameter vector, I use as a measure of fit the norm of the relative error over the calibration targets. The table reports the average simulated statistics for optimized versions of the model. For the learning models, I optimize over the parameters $(\beta, \gamma, \varepsilon)$ to best match the four data moments. In the baseline model, I set the risk-free rate to its zero-debt self-confirming equilibrium level, which here gives $100 \times \log R = 0.87$, close to the conventional calibration of a 1% risk-free rate. In the rational expectations version of the model, the price-dividend ratio is constant at δ^* and the standard deviation of stock returns is given by σ, the exogenous dividend growth volatility. Furthermore, $E \log Z = \mu$, which is pinned down by δ^* and d, the exogenous mean dividend growth rate. I find the δ^* that best fits the two moments of the mean price-dividend ratio and mean stock return, which in turn pins down a one-dimensional (β, γ) surface. The table sets $\gamma = 3$, a value close to the optimized learning models, and reports the corresponding β. The fifth row of the table provides the error measure, the average relative error across the four data moments, for each model specification. The next row gives the average duration between market crashes, defined as a 30% decline in the price-dividend ratio. The bottom two rows give the mean and standard deviation of consumption growth.

The baseline learning model is able to fit the targeted moments quite well with preference parameters within the commonly used range. The means of the P/D ratio and stock return are very close to but slightly lower than their values in the data, while the model nails their volatilities. Using the baseline parameters in a closed economy version of the model, which is similar to Adam, Marcet, and Nicolini (2016) and does not feature the debt dynamics, results in notably worse fit: the price-dividend ratio is too high, stock returns are lower, and risky assets are substantially less volatile. Thus the debt dynamics and portfolio adjustment play a key role in generating booms, busts, and asset price volatility.

Re-optimizing the parameters in the closed economy model improves its fit, especially for the means of the price-dividend ratio and the stock returns. But the model misses on the volatilities, overshooting on stock returns and undershooting on the P/D ratio. These optimized preference parameters differ substantially from the baseline calibration, with a much lower discount factor β, a higher risk aversion coefficient, and a much larger learning gain. These findings are roughly similar to Adam, Marcet, and Nicolini (2016), although the learning specification here differs slightly and they target more data moments. With a risk aversion coefficient of $\gamma = 5$, their model is able to match the mean and volatility of the price-dividend ratio, but cannot match the mean stock returns.

The debt dynamics of the baseline model drive the booms and crashes, bringing the simulations closer to the data. Table 1 shows that large market crashes, defined as a 30% drop in the price-dividend ratio, occurred on average every 25 years in the data. In the baseline model, such crashes occur less frequently, roughly every 38 years. However that the baseline model can generate them at all is notable, as crashes of this magnitude were hardly ever observed in the closed economy model, and of course cannot happen in the rational model where the price-dividend ratio is constant. Similar results hold for other market crash definitions, whether using different thresholds or stock returns instead of the P/D ratio.

Both learning models fit the data much better than the rational expectations model. As noted above, dividends are relatively volatile, so under this calibration the rational model fits the mean price-dividend ratio and the mean stock return reasonably well, at the cost of counterfactual consumption volatility. But the rational model generates lower stock return volatility, driven entirely by the exogenous dividends, and a constant price-dividend ratio.

Overall, learning can help explain observed asset price volatility. In addition, debt dynamics and the associated shifts in asset allocations can generate asset market volatility, plausible risk prices, and recurrent market crashes. The financial instability model can substantively contribute to our understanding of asset market data.

7.3 Differences Between Model and Data

While the model fits along a number of dimensions, it also has some notable misses. In some ways this is unsurprising, as my focus is mostly theoretical, so the model is kept simple with only a few free parameters. These drawbacks suggest avenues for future work.

The first drawback is that consumption growth is overly volatile in the model, which is a consequence of using empirical dividend growth as the driving shock process. Table 1 shows that all of the models generate very similar consumption growth, which is nearly equal to the mean of the calibrated exogenous dividend growth process.[16] In the rational model, consumption is equal to dividends, and it is nearly so in all of the other models. All models also generate relatively high consumption growth volatility. In the rational model, again this is inherited from the calibration. In my baseline learning model, consumption growth is also driven by changes in beliefs and net asset positions, and I find slightly lower volatility than in the rational model. The closed economy learning model generates slightly higher consumption growth volatility, which is magnified with the enhanced volatility of the optimized parameterization. This suggests that the debt dynamics of the open economy model slightly dampen consumption volatility. Nonetheless, consumption volatility is implausibly high in all cases, which reflects equating consumption and dividends as in Lucas (1978).

[16]The data column uses annual real per capita personal consumption expenditure growth from 1930-2022.

Considering stocks as levered equity claims, as in Abel (1999), would break the link between consumption and dividends and help to close this gap.

In addition, the dynamics of returns in my model generate patterns of return predictability that are much stronger than in the data. The model generates return predictability from persistent changes in beliefs θ_t and net asset positions ω_t, with the main indicator of the market condition being the portfolio share v_t. During market recoveries and booms, the portfolio share tends upward, and returns tend to be positively correlated. But a market boom today also predicts a crash in the future, as the economy enters the fragile region. In Appendix C.1, I show that both the autocorrelation of returns and long-term reversal are much stronger in the model than in the data, where such evidence is weak at best. Adding additional sources of shocks or incorporating heterogeneity would dampen this predictability.

7.4 Open Economy versus Closed Economy

The baseline model is a small open economy, while most macro asset pricing models consider the US as a closed economy. In this section I explore the importance of the open economy assumption. What matters for the asset price dynamics is the debt dynamics and nontrivial changes in asset positions, which the closed economy model lacks. I show that intermediate specifications, which allow for some interest variation but maintain a nontrivial portfolio choice, maintain the same qualitative features of the baseline small open economy model.

The difference between the baseline and closed economy models is illustrated in Figure 8, which plots a representative simulation of each specification with the same baseline parameter values and the same shock realizations. The left column shows the log returns, which in the baseline model are volatile and display notable outliers, with periods of rapid market crashes as well as large recoveries and market spikes. The center column shows the price-dividend ratio, which displays the patterns of booms and crashes we saw above in the baseline model. The top right panel shows the dynamics of the portfolio share v_t, which increases toward one in periods of high returns and P/D ratios, then falls substantially during market crashes with large drops in the P/D ratio and negative returns.

By contrast, both returns and the price-dividend ratio in the closed economy model are much less volatile. The tight range of return fluctuations is due to the variations in returns being driven mostly by the binary dividend shocks, with changes in beliefs playing a much smaller role. The price-dividend ratio drifts over time with changes in beliefs (through $\delta(\theta_t)$), but does not display either the overall volatility or the periods of rapid change that we observed in the data. In the baseline model, the net asset position and resulting debt dynamics also affect returns (through the variable ω_t) and magnify the impact of belief variation. But in the closed economy model, the interest rate adjusts so that debt is always zero. The bottom right panel shows the time series of the risk-free interest rate in the

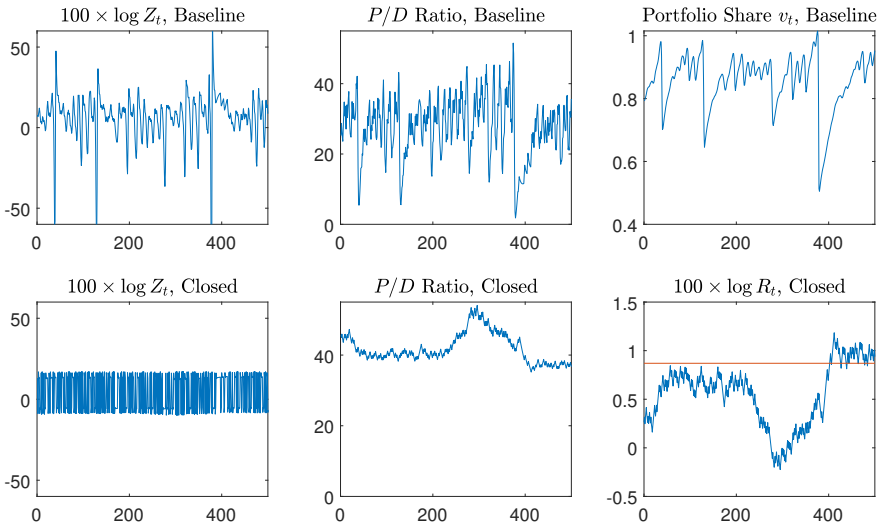


Figure 8: Log returns (left column), price-dividend ratio (center column), along with portfolio shares and interest rates (right column) from simulations of the baseline model (top row) and the closed economy model (bottom row).

closed economy model, along with the constant value in the open economy model. The real interest rate has a strong negative correlation with the price-dividend ratio in the closed economy model. After period 200, the P/D ratio drifts slowly upward, while the interest rate falls steadily and drops below zero. Then after around period 300, the P/D ratio slowly drifts down as the real interest rate increases back to the open economy level. Overall, the price-dividend ratio movements are slow and the fluctuations are small.

Table 1 also reports the implications of optimizing parameters for the closed economy model. To better fit the data, this specification increases the risk aversion coefficient to increase risk premia, and increases the gain to speed up the dynamics. While these changes increase the volatility of prices and returns, they do not lead to plausible boom and crash dynamics. Some market crashes do occur in the simulations, but they are extremely rare.

In Appendix C.2 I report some results from an intermediate case between the small open and closed models. Here the interest rate adjusts each period based on household asset demand, but not enough to ensure that there is zero debt, as in the closed economy. I find that the qualitative features of the closed economy model persist in many intermediate cases.

The debt dynamics and asset allocation features of my baseline open economy model, where funds flow in and out of risky assets in response to changes in beliefs, are crucial for the asset pricing implications. The small open economy assumption was a simple means of introducing borrowing and lending. My results suggest that a nonzero but imperfectly elastic supply of bonds, which is surely a relevant case empirically for large countries on the

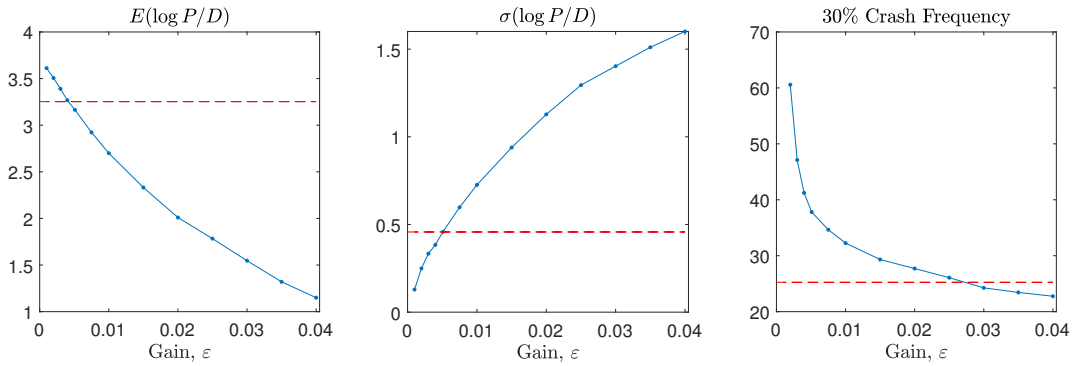


Figure 9: Mean log price-dividend ratio (left panel) standard deviation of the log price-dividend ratio (center panel), and average duration between crashes (right panel) from simulations in the model for different values of the learning gain. Values from the data are shown with dashed lines.

world market like the United States, generates similar implications.

7.5 Quantitative Impact of the Speed of Learning

The theoretical analysis of the model showed the importance of the gain ε, which governs the speed of learning. For the calibration in the previous section, I chose the gain ε to best fit moments from asset pricing data.In this section I illustrate in more detail how the gains affect outcomes in the model, showing the tradeoffs balanced in the calibration process.

Figure 9 shows how the mean log price-dividend ratio (left panel), its volatility (center panel), and the frequency of market crashes (right panel) vary with the gain. In each panel I plot the relevant statistic computed as above, where for each gain I run 1,000 simulations of 2,000 periods each, using the last 1,000 periods in the calculations. I also show the values from the data with a dashed line. The left and center panels illustrate a key tradeoff which the calibration balances: as the gain decreases the mean price-dividend ratio increases, but the volatility falls. Consistent with the convergence result, as the gain decreases to zero the price-dividend ratio converges to its self-confirming equilibrium value, which in this parameterization is $\log(P/D) = \log \delta(\bar{\theta}) \approx 3.6$, and the simulated distribution becomes increasingly concentrated on this value. The calibration balances the mean and volatility of the price-dividend ratio as targets, and is able to nearly match them at a gain of $\varepsilon = 0.0051$.

The right panel of Figure 9 shows the average duration between market crashes, again defined as a 30% drop in the price-dividend ratio. In the calibrated model, market crashes occur less frequently than in the data. Matching the observed crash frequency would require a larger gain around 0.025, increasing the speed of belief adjustment. Increases in the gain

increase the volatility of asset prices. This increases the local fluctuations around the self-confirming equilibrium and also makes large market booms and crashes more likely. However booms and crashes are asymmetric: booms are shorter-lived, while crashes take a long time to recover from. Therefore increases in the gain make crashes more frequent but decrease the average price-dividend ratio as more time is spent recovering from a crash. While the model is successful at matching the mean and volatility of asset prices and generates plausible market crashes, it has difficulty matching the speed of asset market dynamics.

7.6 Beliefs and Survey Expectations

In my model, the subjective excess return is essentially uncorrelated with the price-dividend ratio. However, Greenwood and Shleifer (2014), Adam, Marcet, and Beutel (2017), and Barberis, Jin, and Shleifer (2018) among others, have shown that survey measures of expected returns are positively correlated with price-dividend ratios. In Appendix D, I discuss how to reconcile the model with the survey evidence. The key in the model is the subjective risk-return tradeoff, which determines the portfolio position v_t, and drives asset market dynamics. I argue that some survey data may reflect such subjective risk-adjusted returns. I also show that a small variation on the model, allowing for different gains for the mean and variance as in the last column of Table 1, brings its implications closer to the survey evidence.

8 Conclusion

I have developed a simple model where learning by investors leads to recurrent booms and busts in asset prices, capturing aspects of Minsky's financial instability hypothesis. Periods of tranquility lead investors to increase their estimates of the risk-return tradeoff, investing more in risky assets, which drives up their price. Once agents are sufficiently concentrated in stocks, the economy becomes fragile, and minor portfolio shifts can trigger substantial price changes. Market booms can sow the seeds of their destruction, as rapid increases in stock prices may increase agents' uncertainty, causing them to cut their portfolio allocation to stocks. This slowing of asset flows leads to an abrupt market crash, further spiking perceived uncertainty. During the recovery, agents observe positive asset market performance and gradually forget the crash, setting the stage for another boom. The cycle of booms and crashes is driven by a fundamental instability that arises when investors learn from prices.

While my focus is theoretical, analyzing how a learning model can generate financial instability, the model has some quantitative success. The learning model can help explain observed market phenomena: high but volatile price-dividend ratios, sizeable premia on risky assets (with relatively low risk aversion), and recurrent episodes of large market crashes. The

model misses along some empirical dimensions, as it was intentionally simplified to focus on the role of belief dynamics. The underlying rational expectations equilibrium features a constant price-dividend ratio (so that all fluctuations are due to learning), output growth is exogenous (to focus on the role of asset prices), the risk-free interest rate is constant (so that changes in stock returns are entirely driven by investment flows), and there is a representative agent (to simplify the characterization and analysis). Relaxing each of these simplifications would be important for many reasons: to better analyze the interactions between financing and the real economy through investment, to better match observed behavior, and to analyze potential policy responses to limit asset market bubbles and crashes.

While a discussion of potential policy responses is outside the scope of this paper, a natural question is whether policymakers could dampen the boom-bust cycles. As may be expected, a higher interest rate lowers excess returns on stocks and leads to less frequent boom-and-bust cycles, but the cycles can become more severe. Booms and crashes may be more infrequent when borrowing costs are higher, but when a boom occurs it may lead to larger spikes in prices and returns. This suggests policymakers face tradeoffs more complex than simply "leaning against the wind" of market bubbles.

Nonetheless, even my simple model displays relatively complex behavior and captures key aspects of market dynamics, as in Minsky's theory. Periods of stability lead to periods of instability. Recurrent booms and busts are driven by recurrent, endogenous shifts in beliefs, whose effects are amplified by debt dynamics. Larger and more rapid booms make more severe crashes inevitable. Financial instability results from dynamic instability in the feedback between beliefs and prices. As long as investors put sufficient weight on new information, booms and busts follow.

References

[1] Abel, Andrew B (1999) "Risk Premia and Term Premia in General Equilibrium," *Journal of Monetary Economics*, 43: 3-33.

[2] Adam, Klaus, Albert Marcet, and Johannes Beutel (2017) "Stock Price Booms and Expected Capital Gains," *American Economic Review*, 107: 2352–2408.

[3] Adam, Klaus, Albert Marcet and Juan Pablo Nicolini (2016) "Stock Market Volatility and Learning," *Journal of Finance*, 71: 33-82.

[4] Barberis, Nicholas, Robin Greenwood, Lawrence Jin, and Andrei Shleifer (2018) "Extrapolation and Bubbles," *Journal of Financial Economics*, 129: 203-227.

[5] Beaudry, Paul, Dana Galizia, and Franck Portier (2020) "Putting the Cycle Back into Business Cycle Analysis," *American Economic Review*, 110: 1–47.

[6] Bianchi, Francesco, Cosmin Ilut, and Hikaru Saijo (2024) "Diagnostic Business Cycles," *Review of Economic Studies,* 91: 129–162.

[7] Bordalo, Pedro, Nicola Gennaioli, and Andrei Shleifer (2018) "Diagnostic Expectations and Credit Cycles," *Journal of Finance*, 73: 199–227.

[8] Boz, Emine and Enrique G. Mendoza (2014) "Financial Innovation, the Discovery of Risk, and the U.S. Credit Crisis," *Journal of Monetary Economics*, 62: 1-22.

[9] Branch, William A. and George Evans (2011) "Learning about Risk and Return: A Simple Model of Bubbles and Crashes," *American Economic Journal: Macroeconomics*, 3: 159-191.

[10] Brianti, Marco and Vito Cormun (2024) "Expectation-Driven Boom-Bust Cycles," *Journal of Monetary Economics*, 146: 103575.

[11] Chien, YiLi, Cho, In-Koo, Ravikumar, B. (2020) "Convergence to Rational Expectations in Learning Models: A Note of Caution," Federal Reserve Bank of St. Louis Working Paper 2020-027.

[12] Cho, In-Koo (2018) "Learning Growth Rate: Spurious Stability and Hidden Discontinuity," Working paper, Emory University.

[13] Cho, In-Koo, Noah Williams, and Thomas J. Sargent (2002) "Escaping Nash Inflation," *Review of Economic Studies*, 69: 1-40.

[14] Cogley, Timothy and Thomas J. Sargent (2008) "Anticipated Utility and Rational Expectations as Approximations of Bayesian Decision Making," *International Economic Review*, 49: 185-221.

[15] Eggertsson, Gauti and Paul Krugman (2012) "Debt, Deleveraging, and the Liquidity Trap: A Fisher-Minsky-Koo Approach," *Quarterly Journal of Economics,* 127: 1469–1513.

[16] Evans, George and Seppo Honkapohja (2001) *Learning and Expectations in Macroeconomics,* Princeton University Press.

[17] Evans, George and Seppo Honkapohja (2009) "Robust Learning Stability with Operational Monetary Policy Rules," in *Monetary Policy under Uncertainty and Learning*, (Klaus Schmidt-Hebbel and Carl E. Walsh, eds.) Central Bank of Chile, Santiago, Chile.

[18] Evans, George, Seppo Honkapohja, and Noah Williams (2010) "Generalized Stochastic Gradient Learning," *International Economic Review,* 51: 237-262.

[19] Farboodi, Maryam and Peter Kondor (2023) "Cleansing by Tight Credit: Rational Cycles and Endogenous Lending Standards," *Journal of Financial Economics,* 150: 46-67.

[20] Friedman, Benjamin M. and David I. Laibson (1989) "Economic Implications of Extraordinary Movements in Stock Prices," *Brookings Papers on Economic Activity*, 1989 (2): 137-72.

[21] Fudenberg, Drew and David Levine (1998) *The Theory of Learning in Games,* MIT Press, Cambridge.

[22] Gennaioli, Nicola, and Andrei Shleifer (2018) *A Crisis of Beliefs: Investor Psychology and Financial Fragility,* Princeton University Press, Princeton, NJ.

[23] Greenwood, Robin, and Andrei Shleifer (2014) "Expectations of Returns and Expected Returns," *Review of Financial Studies,* 27: 714-746.

[24] Greenwood, Robin, Samuel G. Hanson and Lawrence J. Jin (2019) "Reflexivity in Credit Markets," NBER Working Paper 25747.

[25] Greenwood, Robin, Samuel G. Hanson, Andrei Shleifer and Jakob Ahm Sørensen (2022) "Predictable Financial Crises," *Journal of Finance*, 88: 863-921.

[26] Hakansson, Nils H. (1970) "Optimal Investment and Consumption Strategies Under Risk for a Class of Utility Functions," *Econometrica*, 38: 587-607

[27] Honkapohja, Seppo and Nigel McClung (2023) "On Robustness of Average Inflation Targeting," Working paper, Aalto University School of Business.

[28] Jin, Lawrence (2015) "A Speculative Asset Pricing Model of Financial Instability," working paper, California Institute of Technology.

[29] Jin, Lawrence and Pengfei Sui (2022) "Asset Pricing with Return Extrapolation," *Journal of Financial Economics* 145: 273–295.

[30] Keen, Steve (1995) "Finance and Economic Breakdown: Modeling Minsky's 'Financial Instability Hypothesis," *Journal of Post Keynesian Economics*, 17: 607-35.

[31] Kindleberger, Charles (1978) *Manias, Panics and Crashes*, Basic Books, New York.

[32] Kreps, David (1998) "Anticipated Utility and Dynamic Choice," in *Frontiers of Research in Economic Theory* (D. P. Jacobs, E. Kalai, and M. I. Kamien, eds.), Cambridge University Press.

[33] Krishnamurthy, Arvind and Wenhao Li (2022) "Dissecting Mechanisms of Financial Crises: Intermediation and Sentiment," *Review of Financial Studies,* forthcoming.

[34] Kushner, Harold J. and George G. Yin (1997) *Stochastic Approximation Algorithms and Applications,* Springer-Verlag, New York.

[35] Kuznetsov, Yuri A. (1998) *Elements of Applied Bifurcation Theory*, 2nd Edition, Springer, New York.

[36] L'Huillier, Jean-Paul, Gregory Phelan and Hunter Wieman (2024) "Technology Shocks and Predictable Minsky Cycles," *Economic Journal*, 134: 811–836.

[37] Lucas, Robert E., Jr. (1978) "Asset Prices in an Exchange Economy," *Econometrica*, 45: 1429-45.

[38] Malmendier, Ulrike and Stefan Nagel (2016) "Learning from Inflation Experiences," *Quarterly Journal of Economics,* 131: 53–87.

[39] Marcet, Albert and Thomas J. Sargent (1989) "Convergence of Least Squares Learning Mechanisms in Self-Referential Linear Stochastic Models," *Journal of Economic Theory,* 48: 337-368.

[40] Maxted, Peter (2023) "A Macro-Finance Model with Sentiment," *Review of Economic Studies*, forthcoming.

[41] Merton, Robert (1969) "Lifetime Portfolio Selection under Uncertainty: The Continuous Time Case," *Review of Economics and Statistics*, 51: 247-257.

[42] Minsky, Hyman (1975) *John Maynard Keynes*, Columbia University Press

[43] Minsky, Hyman (1986) *Stabilizing an Unstable Economy*, Yale University Press.

[44] Minsky, Hyman (1989) "Economic Implications of Extraordinary Movements in Stock Prices: Comments," *Brookings Papers on Economic Activity*, 1989 (2): 173-82.

[45] Minsky, Hyman (1992) "The Financial Instability Hypothesis," Working paper, Levy Economics Institute.

[46] Nagel, Stefan and Zhengyang Xu (2022) "Asset Pricing with Fading Memory," *Review of Financial Studies,* 35: 2190–2245.

[47] Sargent, Thomas J. (1999) *The Conquest of American Inflation,* Princeton University Press.

[48] Sethi, Rajiv (1992) "Dynamics of Learning and the Financial Instability Hypothesis," *Journal of Economics (Zeitschrift fur Nationalokonomie)*, 56: 39-70.

[49] Taylor, Lance and Stephen A. O'Connell (1985) "A Minsky Crisis," *Quarterly Journal of Economics*, 100: 871-85.

[50] Sargent, Thomas J. and Noah Williams (2005) "Impacts of Priors on Convergence and Escapes from Nash Inflation," *Review of Economic Dynamics*, 8: 360-391.

[51] Timmermann, Alan G. (1993) "How Learning in Financial Markets Generates Excess Volatility and Predictability in Stock Prices," *Quarterly Journal of Economics*, 108: 1135-4563.

[52] Williams, Noah (2019) "Escape Dynamics in Learning Models," *Review of Economic Studies*, 86: 882-912.

A Intermediate Calculations

A.1 Rational Expectations Equilibrium

Since in equilibrium consumption is equal to dividends, $C_t = D_t$, the Euler equation for bonds (3) implies:

$$1 = \beta R E_t \left[\left(\frac{D_{t+1}}{D_t} \right)^{-\gamma} \right] = \beta R \exp(-\gamma d) E_t \exp(-\gamma \sigma W_{t+1}),$$

so defining:

$$\phi = E_t \exp(-\gamma \sigma W_{t+1}) = \frac{1}{2}[\exp(-\gamma\sigma) + \exp(\gamma\sigma)]$$

we then have:

$$\log R = -\log \beta + \gamma d - \log \phi.$$

Given preferences and dividends, I use this equation to set the interest rate R as the value consistent with the zero-debt equilibrium.

Similarly, the Euler equation for stocks (4) can be written:

$$P_t = \beta E_t \left[\left(\frac{D_{t+1}}{D_t} \right)^{-\gamma} (P_{t+1} + D_{t+1}) \right].$$

Then guessing that the price-dividend ratio is constant, $P_t = \delta^* D_t$, delivers:

$$\delta^* = \beta(1+\delta^*)E_t\left(\frac{D_{t+1}}{D_t}\right)^{1-\gamma} = \beta(1+\delta^*)E_t \exp\left((1-\gamma)(d+\sigma W_{t+1})\right).$$

So defining:

$$\bar{\phi} = E_t \exp((1-\gamma)\sigma W_{t+1}) = \frac{1}{2}[\exp(-(1-\gamma)\sigma) + \exp((1-\gamma)\sigma)],$$

I solve for the equilibrium price-dividend ratio:

$$\delta^* = \frac{\beta \exp((1-\gamma)d)\bar{\phi}}{1-\beta \exp((1-\gamma)d)\bar{\phi}}.$$

Therefore the gross return on stocks is given by:

$$Z_t = \frac{P_t + D_t}{P_{t-1}} = \frac{(1+\delta^*)D_t}{\delta^* D_{t-1}}.$$

Taking logs and using (1) gives the expression (5) in the text.

A.2 Equilibrium for Arbitrary Beliefs

With arbitrary beliefs $\theta = (m, s)$, Hakansson (1970) shows that the optimal portfolio fraction $v(\theta)$ solves the static maximization problem:

$$U(v(\theta)) = \max_{0 \le v \le \overline{v}} E^\theta \left[\frac{(vZ(\theta) + (1-v)R)^{1-\gamma}}{1-\gamma}\right], \tag{13}$$

where E^θ represents expectation over the random gross returns $Z(\theta)$, defined by the subjective belief distribution with parameters θ. Hakansson (1970) does not include portfolio constraints, but due to the homogeneity the results extend to restrictions of this form. Then the optimal fraction of wealth to consume $c(\theta)$ is:

$$c(\theta) = 1 - [\beta U(v(\theta))(1-\gamma)]^{\frac{1}{\gamma}}. \tag{14}$$

With a subjective binomial distribution, we can write the possible return outcomes as $Z_H = \exp(m+s)$ and $Z_L = \exp(m-s)$. Then we can solve for the optimal portfolio weight:

$$v(\theta) = \max\left\{\min\left\{\frac{R\left((R-Z_L)^{-\frac{1}{\gamma}} - (Z_H - R)^{-\frac{1}{\gamma}}\right)}{(R-Z_L)^{1-\frac{1}{\gamma}} + (Z_H-R)^{1-\frac{1}{\gamma}}}, \overline{v}\right\}, 0\right\}$$

so that:

$$c(\theta) = 1 - \left(\frac{\beta}{2}\right)^{\frac{1}{\gamma}} \left((v(\theta)(Z_L - R) + R)^{1-\gamma} + (v(\theta)(Z_H - R) + R)^{1-\gamma}\right)^{\frac{1}{\gamma}}.$$

With expression (6) for the price-dividend ratio, we can solve for stock returns:

$$Z_t = \frac{P_t + D_t}{P_{t-1}} = \frac{\frac{P_t}{D_t} + 1}{\frac{P_{t-1}}{D_{t-1}} \frac{D_{t-1}}{D_t}} = \frac{1 + \delta(\theta)\omega_t}{\delta(\theta)\omega_{t-1}\chi_t}, \tag{15}$$

where

$$\chi_t = \frac{D_{t-1}}{D_t} = \exp(-d - \sigma W_t).$$

When $\omega_t = \omega_{t-1} = 1$, this reduces to the rational expectations expression, with returns driven by dividend growth. But in general, agents' bond positions affect their wealth X_t through interest payments made or received, and thus their demand for current stock purchases. Therefore prices and returns are in general correlated over time, with prices depending on the asset position ω_t and returns depending on the dynamics of ω_t. Working through the dynamics of the budget constraint, I can write the dynamics of the asset position as:

$$\omega_t = 1 + q(\theta)\omega_{t-1}\chi_t \tag{16}$$

where I define the weight function:

$$q(\theta) = R(1 - c(\theta))(1 - v(\theta))(1 + \delta(\theta)).$$

At the SCE, all wealth is invested in stocks so $v(\theta) = 1$ and $q(\theta) = 0$. Thus the net asset position is constant at $\omega_t \equiv 1$. If the agent holds bonds, then $0 < v(\theta) < 1$ so $q(\theta) > 0$ and the net asset position ω_t increases over time, with a stochastic growth rate. When the agent borrows, so $v(\theta) > 1$ then $q(\theta) < 0$, and the net asset position decreases.

I now provide more detail which explains the shape and slope of the pricing function $\delta(v)$ in Figure 1. As I discuss in the text, the pricing function diverges as $v \to \frac{1}{1-c}$. This is easiest to demonstrate when the agent comes into the period with no debt so $B_{t-1} = 0$, and as in the text I now focus directly on the dependence on the portfolio share $\delta(v)$ rather than on beliefs. Then the demand for stocks can be written

$$P_t S_t = (1 - c)v S_{t-1}(P_t + D_t).$$

Then for $v = \frac{1}{1-c}$ we would need $P_t S_t = S_{t-1}(P_t + D_t)$ or

$$\delta S_t = (1 + \delta) S_{t-1},$$

which clearly cannot hold with finite δ for any fixed share supply $S_t = S_{t-1}$. The demand for stocks outstrips the supply at any finite price, which means that the share price must diverge as $v \to \frac{1}{1-c}$. Equivalently, if $S_{t-1} = 1$, the demand for current shares is:

$$S_t^d = \frac{(1 - c)v(1 + \delta)}{\delta}.$$

As $v \to \frac{1}{1-c}$, then the demand $S_t^d \to (1+\delta)/\delta$. Thus we can only have demand equal to the supply (of one) for $\delta \to \infty$. We can also explicitly calculate the slope of the pricing function at the SCE. We have:

$$\delta'(v) = \frac{1-c}{1-(1-c)} + \frac{(1-c)^2 v}{(1-(1-c)v)^2},$$

so that

$$\delta'(1) = \frac{1-c}{1-(1-c)} + \frac{(1-c)^2}{(1-(1-c))^2} = \delta(1) + \delta(1)^2.$$

Thus the slope of the pricing function at the SCE scales with the square of its value in the SCE: $\delta'(1) = (1+\delta(1))\delta(1)$, which explains the sensitivity of prices near the SCE.

I now detail the portfolio constraint. Using similar calculations as above, we can write:

$$\frac{B_t}{D_t} = (1-c)(1-v)\frac{X_t}{D_t} = (1+\delta)(1-c)(1-v)\frac{RB_{t-1}+D_t}{D_t} = (1+\delta)(1-c)(1-v)\omega_t,$$

where the second and third equalities use (6). Therefore the debt restriction can be written:

$$(1+\delta)(1-c)(1-v)\omega_t \equiv f(v)\omega_t \geq -\frac{\exp(d-\sigma)}{R}.$$

Since δ and c both depend on v, I write the left expression as a function $f(v)$ given beliefs θ. This function f is monotonically decreasing in v, so the restriction can be written as $v \leq \overline{v}_t$ for some $\overline{v}_t$, where the time dependence comes from the net asset position ω_t. Starting from a position of zero debt, $\omega_t = 1$ so the constraint is simply $v \leq \overline{v}$, a constant. But as the economy evolves, the net asset position changes. In turn, this affects how much leverage the economy can support. So the portfolio constraint must evolve over time.

B Proofs of the Results

B.1 Convergence Results

Marcet and Sargent (1989), Evans and Honkapohja (2001), and others have adapted the theory of stochastic approximation to show that in the limit as the gain gets small, the evolution of beliefs can be approximated by a deterministic ordinary differential equation. The basic idea is that we can rewrite the updating equation (7) as:

$$\frac{m_{t+1} - m_t}{\varepsilon} = \log Z(\theta_t, \theta_{t-1}, \omega_{t-1}, \chi_t) - m_t.$$

The left side of the equation can be interpreted as a finite-difference approximation of a time derivative, where ε is the notional "time" between observations. Then as $\varepsilon \to 0$ the left side of the equation converges to that time derivative. Also in this limit, more and more observations are packed into any finite interval of time, so the dynamics of outcomes are fast

relative to beliefs. Thus the belief dynamics effectively average over the shock realizations and we can apply a law of large numbers.

Following Theorem 8.5.1 in Kushner and Yin (1997), as $\varepsilon \to 0$, under some regularity conditions, the beliefs converge weakly to the trajectories of the mean dynamics ODEs:

$$\dot{m} = E[\log Z(\theta, \theta, \omega, \chi)] - m, \tag{17}$$

$$\dot{s}^2 = E[\log Z(\theta, \theta, \omega, \chi) - m]^2 - s^2. \tag{18}$$

Here the expectation is with respect to the unconditional distribution of ω_t for fixed θ, as in (16), and the distribution of the dividend growth shocks χ.

The self-confirming equilibrium is clearly an equilibrium point of the mean dynamics, as seen by setting $\dot{m} = \dot{s}^2 = 0$:

$$\begin{aligned} m = \mu &= E[\log Z(\theta, \theta, \omega, \chi)] \\ s^2 = \sigma^2 &= E[(\log Z(\theta, \theta, \omega, \chi) - \mu)^2] \end{aligned}$$

I now show that the SCE is a locally stable equilibrium point of the mean dynamics, and therefore of the learning rule. As in Evans and Honkapohja (2001), the SCE is locally expectationally stable if all of the eigenvalues of the Jacobian matrix of the mean dynamics have negative real parts when evaluated at the SCE. So for small enough gain ε, and (at least) for beliefs that start near the SCE, over time beliefs will converge to the SCE. The condition on the risk aversion coefficient is sufficient but stronger than necessary, as discussed in the proof. The baseline calibration has larger risk aversion, but I verified numerically (evaluating the eigenvalues of the Jacobian with the specified parameters) that stability still holds.

Theorem 2 *When $\gamma \leq 2$, the self-confirming equilibrium beliefs $\overline{\theta} = (\mu, \sigma)$ are a locally stable fixed point of the learning dynamics (7)-(8). Therefore as $\varepsilon \to 0$ and $t \to \infty$, the stochastic belief process $\{\theta_t\}$ will converge weakly to $\overline{\theta}$.*

Proof. The convergence result follows from Kushner and Yin (1997). Apart from the stability conditions established below, the regularity conditions are satisfied as the shocks are bounded, and the portfolio constraints imply that beliefs and outcomes are bounded.

Tedious but straightforward calculations verify that $v(\overline{\theta}) = 1$ and $c(\overline{\theta}) = 1/(1+\delta)$. Therefore $\delta(\overline{\theta}) = \delta^*$, $q(\overline{\theta}) = 0$. In particular, using the relations above we have:

$$\mu = \log(1+\delta^*) - \log(\delta^*) + d = -\log\beta - \log\overline{\phi} + \gamma d = \log R + \log(\phi/\overline{\phi}).$$

Therefore, defining $\psi = \phi/\overline{\phi}$, we have:

$$Z_H - R = (e^{\sigma}\psi - 1)\, R, \quad R - Z_L = \left(1 - e^{-\sigma}\psi\right) R.$$

Then we can write:

$$v(\overline{\theta}) = \frac{1-\Gamma}{1-\Gamma+\psi(\Gamma e^{\sigma}-e^{-\sigma})}$$

where:

$$\Gamma = \left(\frac{Z_h - R}{R - Z_L}\right)^{-\frac{1}{\gamma}} = \left(\frac{e^{(1+\gamma)\sigma} - e^{-(1-\gamma)\sigma}}{e^{(1-\gamma)\sigma} - e^{-(1+\gamma)\sigma}}\right)^{-\frac{1}{\gamma}} = e^{-2\sigma}.$$

Therefore:

$$\Gamma e^{\sigma} - e^{-\sigma} = e^{-2\sigma}e^{\sigma} - e^{-\sigma} = 0,$$

which then gives:

$$v(\overline{\theta}) = \frac{1-\Gamma}{1-\Gamma} = 1.$$

Furthermore, this implies

$$c(\overline{\theta}) = 1 - \frac{\beta^{\frac{1}{\gamma}}}{2}\left(Z_L^{1-\gamma} + Z_H^{1-\gamma}\right) = 1 - \beta^{\frac{1}{\gamma}} e^{\frac{1-\gamma}{\gamma}\mu}\overline{\phi}^{\frac{1}{\gamma}} = 1 - \beta\overline{\phi}e^{(1-\gamma)d} = \frac{1}{1+\delta^*}.$$

Thus we have $q(\overline{\theta}) = 0$ so $\omega_t \equiv 1$. This means that we only need evaluate the dynamics locally around $\omega_t = \omega_{t-1} = 1$. From (9) we can evaluate:

$$\log Z(\theta, \theta, 1, \chi) = \log(1 + \delta(\theta)) - \log(\delta(\theta)) - \log \chi_t.$$

Then we have:

$$E \log Z(\overline{\theta}, \overline{\theta}, 1, \chi) = \log(1 + \delta^*) - \log(\delta^*) - d = \mu,$$

and therefore:

$$E[\big(\log Z(\overline{\theta}, \overline{\theta}, 1, \chi) - \mu\big)^2] = E(\sigma W_t)^2 = \sigma^2.$$

Thus $\overline{\theta}$ is a fixed point of the ODEs (17), and therefore of the learning dynamics.

To establish local stability, we need to evaluate the Jacobian of the mean dynamics at the SCE $\overline{\theta}$. Since the variance estimate s^2 is of second order, it does matter for the local stability. Thus it is enough to simply verify the local stability of the mean estimate. To see this explicitly, we can stack (17) into a vector, and write the dynamics as:

$$\dot{\theta} = b(\theta),$$

where I abuse notation set and $\theta = (m, s^2)$. Simple calculations establish:

$$\frac{\partial b}{\partial \theta}(\overline{\theta}) = \begin{bmatrix} \frac{\partial E \log Z}{\partial m} - 1 & \frac{\partial E \log Z}{\partial s^2} \\ 0 & -1 \end{bmatrix}.$$

So for local stability, we simply require:

$$\frac{\partial}{\partial m} E \log Z(\theta, \theta, 1, \chi))\Big|_{\theta=\overline{\theta}} < 1.$$

From above we have:

$$E \log Z(\theta, \theta, 1, \chi) = \log(1 + \delta(\theta)) - \log(\delta(\theta)) - d,$$

so that:

$$\frac{\partial E \log Z}{\partial m} = \left(\frac{1}{1+\delta^*} - \frac{1}{\delta^*} \right) \frac{\partial \delta}{\partial m}(\overline{\theta}) = -\frac{1}{\delta^*(1+\delta^*)} \frac{\partial \delta}{\partial m}(\overline{\theta}).$$

Then letting $\kappa = 1 - c$, we have:

$$\begin{aligned} \frac{\partial \delta}{\partial m}(\overline{\theta}) &= \left(\frac{\kappa}{1-\kappa} + \frac{\kappa^2}{(1-\kappa)^2} \right) \frac{\partial v}{\partial m}(\overline{\theta}) + \left(\frac{1}{1-\kappa} - 1 \right) \frac{\partial \kappa}{\partial m}(\overline{\theta}) \\ &= \frac{\kappa}{(1-\kappa)^2} \frac{\partial v}{\partial m}(\overline{\theta}) + \frac{\kappa}{1-\kappa} \frac{\partial \kappa}{\partial m}(\overline{\theta}) \\ &= \delta^* \left[(1+\delta^*) \frac{\partial v}{\partial m}(\overline{\theta}) + \frac{\partial \kappa}{\partial m}(\overline{\theta}) \right]. \end{aligned}$$

Then since $\frac{\partial Z_H}{\partial m} = Z_H$ and $\frac{\partial Z_L}{\partial m} = Z_L$, we can write:

$$\frac{\partial v}{\partial m}(\overline{\theta}) = \frac{Z_L \left[R(R - Z_L)^{\frac{-1-\gamma}{\gamma}} + (\gamma - 1)(R - Z_L)^{\frac{-1}{\gamma}} \right] + Z_H \left[R(Z_H - R)^{\frac{-1-\gamma}{\gamma}} - (\gamma - 1)(Z_H - R)^{\frac{-1}{\gamma}} \right]}{\gamma \left[(R - Z_L)^{\frac{\gamma-1}{\gamma}} + (Z_H - R)^{\frac{\gamma-1}{\gamma}} \right]}.$$

But above we established that at $\overline{\theta}$:

$$(Z_H - R)^{\frac{-1}{\gamma}} = e^{-2\sigma} (R - Z_L)^{\frac{-1}{\gamma}},$$

so therefore at the SCE:

$$Z_L (R - Z_L)^{\frac{-1}{\gamma}} - Z_H (Z_H - R)^{\frac{-1}{\gamma}} = (R - Z_L)^{\frac{-1}{\gamma}} (Z_L - e^{-2\sigma} Z_H) = 0.$$

Thus we can simplify the above to:

$$\frac{\partial v}{\partial m}(\overline{\theta}) = \frac{R \left[Z_L (R - Z_L)^{\frac{-1-\gamma}{\gamma}} + Z_H (Z_H - R)^{\frac{-1-\gamma}{\gamma}} \right]}{\gamma \left[(R - Z_L)^{\frac{\gamma-1}{\gamma}} + (Z_H - R)^{\frac{\gamma-1}{\gamma}} \right]} > 0.$$

Then using the definition of $c(\theta)$ we have:

$$\frac{\partial \kappa}{\partial m}(\overline{\theta}) = (1 - \gamma) \frac{\beta^{\frac{1}{\gamma}}}{2} \left(Z_L^{-\gamma} \left((Z_L - R) \frac{\partial v}{\partial m} + Z_L \right) + Z_H^{-\gamma} \left((Z_H - R) \frac{\partial v}{\partial m} + Z_H \right) \right).$$

Similar to above, at the SCE we have:

$$Z_H^{-\gamma} (Z_H - R) - Z_L^{-\gamma} (R - Z_L) = \left(Z_H^{-\gamma} e^{2\gamma\sigma} - Z_L^{-\gamma} \right) (R - Z_L) = 0,$$

so that we can simplify to:

$$\frac{\partial \kappa}{\partial m}(\overline{\theta}) = (1-\gamma)\frac{\beta^{\frac{1}{\gamma}}}{2}\left(Z_L^{1-\gamma} + Z_H^{1-\gamma}\right) = (1-\gamma)\kappa(\overline{\theta}) = (1-\gamma)\frac{\delta^*}{1+\delta^*}.$$

Thus we have

$$\frac{\partial E \log Z}{\partial m}(\overline{\theta}) = -\frac{\partial v}{\partial m}(\overline{\theta}) - (1-\gamma)\frac{\delta^*}{(1+\delta^*)^2}.$$

When $\gamma < 1$ this is clearly negative. More generally, we require:

$$\frac{\partial v}{\partial m}(\overline{\theta}) > (\gamma - 1)\frac{\delta^*}{(1+\delta^*)^2} - 1.$$

For $\gamma \leq 2$, since $\delta^* > 0$ the right side of the above inequality is negative. Since $\frac{\partial v}{\partial m}(\overline{\theta}) > 0$ the condition is guaranteed to hold.

These conditions are stronger than necessary, but the relatively complicated formulas make weaker conditions difficult to establish. For example, with $\gamma = 3$ a sufficient condition would be $\frac{\partial v}{\partial m}(\overline{\theta}) > 1$. Such a condition is easy to verify with given parameters, as in the calibrated model. But it is harder to provide conditions on primitives. In the Merton continuous time analogue of the model, the portfolio decision rule is $v = (m-r)/(\gamma s^2)$, so $\frac{\partial v}{\partial m}(\overline{\theta}) = \frac{1}{\gamma \sigma^2}$. Then as long as $\sigma^2 < 1/\gamma$ we know that $\frac{\partial v}{\partial m}(\overline{\theta}) < 1$. ■

B.2 Averaging and Proof of the Bifurcation Theorem

First I discuss the approximation of the full stochastic system (11) by the averaged system (12). Note that for any fixed X we can write:

$$G(X, \chi_t) = \bar{G}(X) + \sigma V_t.$$

Given X, σV_t is an i.i.d. random variable whose variance is of order σ^2. Starting at an initial value $X_0 = X$, we have:

$$X_1 = G(X, \chi_1) = \bar{G}(X) + \sigma V_1,$$

and clearly $E_0 X_1 = \bar{G}(X)$. But continuing, note that we can write:

$$\begin{aligned} X_2 &= G(X_1, \chi_2) = \bar{G}(X_1) + \sigma V_2 \\ &= \bar{G}(\bar{G}(X) + \sigma V_1) + \sigma V_2 \\ &\approx \bar{G}(\bar{G}(X)) + \bar{G}'(\bar{G}(X))\sigma V_1 + \sigma V_2, \end{aligned}$$

where the last approximation is valid for small σ. Thus we have $E_0 X_2 \approx \bar{G}(\bar{G}(X))$, and we can continue in this fashion. Therefore for small shocks, we can approximate the dynamics of the mean of the stochastic system by the averaged dynamics.

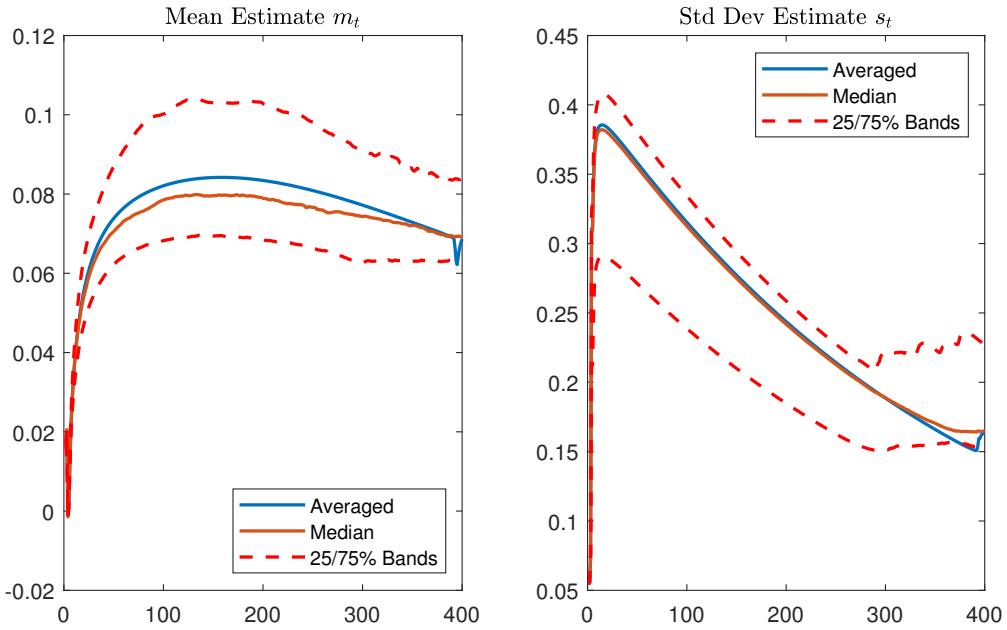


Figure 10: Beliefs m_t and s_t from the averaged dynamics along with quantiles from the simulated distribution.

Figure 10 shows the averaged dynamics for m_t and s_t, along with quantiles of the belief distribution from 1,000 simulations. In both cases the time paths are initialized near the self-confirming equilibrium values. The figure clearly illustrates that the averaging in the small noise approximation provides a good characterization of the belief distribution for the baseline parameterization.

Furthermore, the averaging approximation justifies using the averaged dynamics to study local properties of the dynamics in the neighborhood of the steady state. Using that $\bar{X} = \bar{G}(\bar{X})$ and starting at $X_0 = \bar{X}$, the above approximation gives:

$$X_t = \bar{X} + \sigma \left[V_t + \bar{G}'(\bar{X}) V_{t-1} + \bar{G}'(\bar{X})^2 V_{t-2} + \ldots \right].$$

Therefore local to the self-confirming equilibrium $\bar{X}$, the Jacobian matrix of the averaged system $\bar{G}'(\bar{X})$ governs the system dynamics.

Now I turn to the proof of Theorem 1.

Proof. I first collect the results from Kuznetsov (1998) which underly the theorem, then illustrate how to evaluate the conditions. The result of part 1 of the theorem is a standard stability linearization condition, such as Theorem 1.2 in Kuznetsov. Part 2 of the theorem is simply a result that defines some notation. Part 3 of the theorem condenses Theorems 4.5 and 4.6 in Kuznetsov. Conditions (a) and (b) correspond to (C.1) and (C.2) of Theorem 4.5. Kuznetsov's Theorem 4.6 requires a genericity condition (C.3), which is simply that $a(\bar{\varepsilon}) \neq 0$. The stability condition $a(\bar{\varepsilon}) < 0$ results from the nonlinear stability condition

discussed prior to Lemma 4.3. In particular, under a change of variables Kuznetsov shows that a comes from a third-order approximation of the system (12) and that the terms of fourth order or higher do not affect the stability of the invariant curve.

In later results, Kuznetsov provides a direct expression for the key variable $a(\bar{\varepsilon})$ in terms of the original system. This comes from Theorem 5.3, which applied to my model establishes the existence of a center manifold as a two-dimensional map which defines locally topologically equivalent dynamics of (12) near the SCE $\bar{X}$. While computation of the center manifold is in general difficult, Kuznetsov works out formulas for the third-order approximation that in turn delivers the key term $a(\bar{\varepsilon})$. This requires some notation. First, let $\tilde{X} = X - \bar{X}$ and define the Taylor approximation to the dynamics (12) around the SCE $\bar{X}$ at $\bar{\varepsilon}$ as:

$$\bar{G}(X) = \bar{G}(\bar{X}) + A\tilde{X} + \frac{1}{2}B(\tilde{X}, \tilde{X}) + \frac{1}{6}C(\tilde{X}, \tilde{X}, \tilde{X}) + O(||\tilde{X}||^4),$$

where $A = J(\bar{\varepsilon})$, and $B(x, y)$ and $C(x, y, z)$ are multi-linear functions defined as:

$$B_i(x, y) = \sum_{j=1}^{n}\sum_{k=1}^{n} \frac{\partial^2 \bar{G}_i(\bar{X})}{\partial X_j \partial X_k} x_j y_k, \quad C_i(x, y, z) = \sum_{j=1}^{n}\sum_{k=1}^{n}\sum_{l=1}^{n} \frac{\partial^3 \bar{G}_i(\bar{X})}{\partial X_j \partial X_k \partial X_l} x_j y_k z_l.$$

By assumption, the matrix A has a pair of complex eigenvalues on the unit circle. Denote the corresponding (complex) eigenvectors as $Aq = e^{i\psi_0}q$ and $A\bar{q} = e^{-i\psi_0}\bar{q}$, and define the adjoint eigenvector p as $A'p = e^{-i\psi_0}p$. Also let $\langle p, q\rangle = \sum_i \bar{p}_i q_i$ denote the standard scalar product for complex vectors. Then in equation (5.74) Kuznetsov gives an explicit expression for $a(\bar{\varepsilon})$:

$$a(\bar{\varepsilon}) = \frac{1}{2}\mathrm{Re}\left\{e^{-i\psi_0}\left[\langle p, C(q, q, \bar{q})\rangle + 2\langle p, B(q, (I - A)^{-1}B(q, \bar{q})) + \langle p, B(\bar{q}, \left(e^{2i\psi_0}I - A\right)^{-1}B(q, q))\rangle\right]\right\}. \tag{19}$$

I now show how to verify the conditions in the theorem. The key is to calculate the Jacobian $J(\varepsilon)$, whose eigenvalues determine the stability of the SCE $\bar{X}$. The higher-order derivatives needed for the higher-order approximations B and C follow similarly. Explicitly, we have $X_t = [m_t, s_t^2, m_{t-1}, s_{t-1}^2, \omega_{t-1}]'$ so that the averaged dynamics are:

$$\bar{G}(X_t) = \begin{bmatrix} m_t + \varepsilon(E\log Z(\theta_t, \theta_{t-1}, \omega_{t-1}, \chi_t) - m_t) \\ s_t^2 + \varepsilon(E(\log Z(\theta_t, \theta_{t-1}, \omega_{t-1}, \chi_t) - m_t)^2 - s_t^2) \\ m_t \\ s_t^2 \\ 1 + q(\theta_{t-1})\omega_{t-1}E(\chi_t) \end{bmatrix}.$$

In the following, all expressions and functions are evaluated at the SCE, so I suppress the arguments. First, since Z_t is binomial, we can write $E\log Z = 0.5(\log Z_L + \log Z_H)$ where Z_L is associated with shock realization $\chi_L = \exp(\mu - \sigma)$ and Z_H with $\chi_H = \exp(\mu + \sigma)$.

Similarly, $E(\log Z - m)^2 = .5(\log Z_L - m)^2 + .5(\log Z_H - m)^2$. Therefore for any $x \in X$, $\frac{\partial E \log Z}{\partial x} = 0.5(\frac{\partial \log Z_L}{\partial x} + \frac{\partial \log Z_H}{\partial x})$ and so on. Also let $\bar{\chi} = E\chi_t = .5(\chi_L + \chi_H)$. Then we have:

$$\begin{aligned} J_1(\varepsilon) &= \frac{\partial \bar{G}_1(\bar{X})}{\partial X} \\ &= \varepsilon \left[\frac{1}{\varepsilon} + \frac{\partial E \log Z}{\partial m_t} - 1, \frac{\partial E \log Z}{\partial s_t^2}, \frac{\partial E \log Z}{\partial m_{t-1}}, \frac{\partial E \log Z}{\partial s_{t-1}^2}, \frac{\partial E \log Z}{\partial \omega_{t-1}} \right] \\ &= \varepsilon \left[\frac{1}{\varepsilon} + \frac{1}{1+\delta}\frac{\partial \delta}{\partial m}, \frac{1}{1+\delta}\frac{\partial \delta}{\partial s^2}, \bar{\chi}\frac{\delta}{1+\delta}\frac{\partial q}{\partial m} - \frac{1}{\delta}\frac{\partial \delta}{\partial m}, \bar{\chi}\frac{\delta}{1+\delta}\frac{\partial q}{\partial s^2} - \frac{1}{\delta}\frac{\partial \delta}{\partial s^2}, -1 \right]. \end{aligned}$$

Then we can also write:

$$J_2(\varepsilon) = \varepsilon \begin{bmatrix} (\log Z_L - m)\left(\frac{\partial \log Z_L}{\partial m_t} - 1\right) + (\log Z_H - m)\left(\frac{\partial \log Z_H}{\partial m_t} - 1\right) \\ \frac{1}{\varepsilon} + (\log Z_L - m)\frac{\partial \log Z_L}{\partial s_t^2} + (\log Z_H - m)\frac{\partial \log Z_H}{\partial s_t^2} - 1 \\ (\log Z_L - m)\frac{\partial \log Z_L}{\partial m_{t-1}} + (\log Z_H - m)\frac{\partial \log Z_H}{\partial m_{t-1}} \\ (\log Z_L - m)\frac{\partial \log Z_L}{\partial s_{t-1}^2} + (\log Z_H - m)\frac{\partial \log Z_H}{\partial s_{t-1}^2} \\ (\log Z_L - m)\frac{\partial \log Z_L}{\partial \omega_{t-1}} + (\log Z_H - m)\frac{\partial \log Z_H}{\partial \omega_{t-1}} \end{bmatrix}'$$

where the $\log Z$ derivative expressions can be evaluated as in $J_1(\varepsilon)$. Similarly, we have:

$$J_3(\varepsilon) = [1, 0, 0, 0, 0], \quad J_4(\varepsilon) = [0, 1, 0, 0, 0],$$

and finally:

$$J_5(\varepsilon) = \left[0, 0, \bar{\chi}\frac{\partial q}{\partial m}, \bar{\chi}\frac{\partial q}{\partial s^2}, 0 \right].$$

It remains to evaluate the additional derivatives in the expressions $\frac{\partial q}{\partial \theta_i}$ and $\frac{\partial \delta}{\partial \theta_i}$, for $\theta_i = m, s^2$. For δ and the case of $\theta_i = m$, these explicit calculations are given in the proof of Theorem 2, where embedded in these expressions are further derivatives $\frac{\partial v}{\partial m}$ and $\frac{\partial \kappa}{\partial m}$. The derivatives for $\theta_i = s^2$ follow similarly. For q, the derivatives at the SCE are given by:

$$\frac{\partial q}{\partial \theta_i} = -R\kappa\delta\frac{\partial v}{\partial \theta_i},$$

which can be evaluated given the expressions above. ■

C Additional Quantitative Results

C.1 Predictability of Returns

In this section, I discuss the time series predictability of returns. In the model, returns are strongly positively autocorrelated over one- and two-year horizons. In addition, a strong positive return in a current period makes a crash more likely several periods in the future.

	Data		Simulations	
Years	Coefficient	t-Stat	Coefficient	t-Stat
1	0.0181	0.22	0.6509	27.75
2	-0.1889	-2.32	0.1345	4.32
3	0.0934	1.13	-0.1309	-4.19
4	-0.0613	-0.73	-0.2086	-6.75
5	-0.1302	-1.56	-0.2056	-6.65

Table 2: Coefficients from return predictability regressions in the data along with average regression coefficients and t-statistics from 1000 simulations of the calibrated model.

In the data, there is little or no evidence of autocorrelations in annual market returns, and weak evidence of time series reversals.[17] In Table 2, I report time regressions of the form:

$$R_{t+f-1,t+f} = a_f + b_f R_{t-1,t} + \epsilon_{t+f-1,t+f}$$

where t indexes the current date and $f \in \{1, 2, 3, 4, 5\}$ is the number of years ahead. The table reports the results from the Shiller data, along with simulations of my baseline model. The simulated return data use the same 1,000 simulated time series as in the main text, and I report the mean coefficient and t-statistic from the regressions in each simulation run.

The model generates return predictability from persistent changes in beliefs θ_t and net asset positions ω_t. Moreover, the dynamics of the cycles in my model mean that a market boom today predicts a crash in the future. Both the autocorrelation and the reversal are much stronger in the model than in the data. The data shows some evidence of reversal at two years and very weak evidence of longer-term reversal.

C.2 Intermediate Between Open and Closed Economies

The baseline small open economy model has very different properties from a standard closed economy version. Here I report some simulations from alternative intermediate specifications, where interest rates adjust each period but enough to ensure that there is zero debt. I find that the qualitative features of the closed economy model persist in many intermediate cases. This suggests that the debt dynamics and asset allocations features of the open economy are more crucial for asset pricing implications than the lack of interest rate variation.

Using the results above we can write the demand for bonds (scaled by dividends) as:

$$\frac{B_t}{D_t} = (1 + \delta(R))(1 - c(R))(1 - v(R)) \left(1 + R\frac{B_{t-1}}{D_{t-1}}\chi_t\right),$$

[17] These facts are distinct from the well-documented cross-sectional momentum and reversal.

Statistic	Data	$\alpha = 1$	$\alpha = 0.8$	$\alpha = 0.6$	$\alpha = 0.4$	$\alpha = 0.2$	$\alpha = 0$	$\alpha = 0.5$ (Opt)
$E(\log P/D)$	3.25	3.16	3.31	3.46	3.59	3.68	3.70	3.20
$\sigma(\log P/D)$	0.46	0.46	0.40	0.33	0.24	0.15	0.15	0.38
$100 \times E(\log Z)$	6.51	6.44	5.56	4.91	4.40	4.11	4.10	5.95
$100 \times \sigma(\log Z)$	16.90	16.93	15.30	13.97	12.71	11.70	11.17	15.47
Crash Frequency	25.25	38.30	41.16	46.30	59.82	266.18	–	36.19

Table 3: Summary statistics from data and simulations of the baseline model ($\alpha = 1$), the closed economy version with the baseline parameters ($\alpha = 0$), and a range of intermediate cases ($0 < \alpha < 1$).

where I emphasize the dependence on the interest rate R. In the small open economy case R is a fixed constant, while in the closed economy case R adjusts to level, say R_t^*, so that $B_t = 0$. For an intermediate case, I suppose that bonds are supplied with an interest elasticity which is finite (unlike the small open economy) and positive (unlike the open economy).

Rather than specifying an explicit supply curve, I analyze cases where the equilibrium interest rate R_t is a weighted average of the open and closed economy interest rates:

$$R_t = \alpha R + (1 - \alpha) R_t^*.$$

Varying the weight α between zero and one effectively traces out the implications of different supply elasticities.

The results from the simulations for different α values are reported in Table 3. The setting is the same as in Table 1, where I fix the baseline parameters and consider the same shock realizations, iterations, and sample length as above. Thus $\alpha = 1$ repeats the data from Table 1 for the baseline case while $\alpha = 0$ is the closed economy with the baseline parameterization. As expected, the statistics for the asset pricing data vary continuously with the weighting α, with the mean P/D ratio increasing and its volatility falling with α. However the crash frequency is relatively insensitive to α, increasing only slightly until the weight on the closed economy interest rate drops to 0.2. The last column shows the results for the case of $\alpha = 0.5$ where now I re-optimize over the parameter vector (which yielded a larger gain of 0.01 and slightly larger risk aversion of nearly 4). The results come close to the baseline small open economy parameterization both in terms of the fit to the asset price statistics (the error is 0.204) and roughly match the timing of market crashes.

Data Sources	Correlation
Nagel-Xu (2022): excess return, Shiller $\log P/D$	0.476
Baseline model: m, $\log P/D$	-0.027
Baseline model: $(m - \log R)/s$, $\log P/D$	0.472
Model with $(\varepsilon_m, \varepsilon_s)$: m, $\log P/D$	0.257

Table 4: Correlation coefficients between different measures of subjective beliefs and the log price-dividend ratio. The first line uses the Nagel-Xu (2022) subjective expected excess returns and the price-dividend ratio from the Shiller data set. The others use simulated data from 1,000 simulations of the baseline model or variations of it.

D Beliefs and Survey Expectations

As discussed in the paper, several papers have shown that survey measures of expected returns are positively correlated with price-dividend ratios. The top row of Table 4 reports a such a finding, using the subjective expectations data series from Nagel and Xu (2022).[18] Consistent with the previous literature, the correlation between the survey expected excess return and the empirical price-dividend ratio is positive, at 0.476.

However in my model, the subjective excess return is essentially uncorrelated with the price-dividend ratio. Table 4 reports the correlations from 1,000 simulations of 2,000 periods, discarding the first 1,000 periods. For the baseline calibration of the model from above, the correlation between expected returns m_t and the log price-dividend ratio is -0.027. Thus the baseline model misses the positive correlation between survey expected returns and the price dividend ratio. What really matters for agents in the model is the subjective risk-return tradeoff, which determines their portfolio position, and hence drives asset flows and returns. The third row of Table 4 reports the correlation between the subjective Sharpe ratio and the log-price dividend ratio in the model. The value of 0.472 nearly matches the procyclical correlation observed in the survey expectations.

In fact, many of the surveys ask for positive or negative outlooks rather than quantitative predictions, and so could reflect investors' subjective risk-adjusted assessments rather than simply their expected return forecast. For example, four of the six sources of survey beliefs analyzed by Greenwood and Shleifer (2014) report subjective measures of stock market optimism, rather than quantitative predictions of market returns. In addition, one of the main Gallup surveys asks respondents about expected returns on their own portfolios, which would combine the endogenous portfolio weight (v_t in the model) with beliefs.

The largest movements in beliefs in the model come as a consequence of the collapse of

[18]This series combines and imputes subjective survey expectations from multiple sources, and is available quarterly from 1987:2 to 2021:4.

prices during a market crash. In the simulations, the largest realized returns occur in the immediate aftermath of a crash, as the stock price first falls to a very low level during the crash and then rebounds to a higher (but still historically low) value shortly thereafter. The volatility in returns before and after the crash leads to a spike in the estimated volatility s_t, while the large return after the collapse leads a large upward revision in both the expected return m_t. Both of these changes in beliefs take a long time to fade away. A variation in the model that would dampen the collapse of prices in a crash, such as having pool of outside investors willing to buy at low prices, would lead to less dramatic changes in beliefs after a crash and increase the correlation between expected returns and prices.

Variations in the belief specification alter the model dynamics and hence its predictions. In the baseline model, the gain ε governs the speed of adjustment of both the expected return m_t and the subjective variance s_t^2. While this is a useful benchmark, the belief components could drift at different rates, with gain ε_m for the mean and ε_s for the variance.[19] The fourth row of Table 4 considers a version of the model that separates $(\varepsilon_m, \varepsilon_s)$, where the calibration includes the previous asset pricing moments but also aims to match the correlation between m_t and $\log P_t/D_t$, with a target of 0.476. The other statistics and the parameter values for this specification are given in the last column of Table 1. With this minor change in the belief dynamics, the model is able to generate a larger correlation of 0.257, and the duration between crashes (which was un-targeted) decreases to 27.3 years, which is closer to the empirical value. But the fit of the model on the asset pricing moments worsens, with the relative error increasing from 0.031 to 0.234. This model variation that separates learning speeds improves along some dimensions, but is not a complete fix.

My main focus in the paper was to analyze the dynamics of a learning model developed as an interpretation of Minsky's theory. As such, it was kept simple in order to aid the analysis. If we interpret survey measures as capturing risk-adjusted subjective beliefs, then my model also fits with this survey evidence. But to directly fit the survey measures of expected returns, variations on the baseline model seem necessary. The more general belief specification gets part of the way there, but other sources of shocks or belief specifications may be required. For example, Adam, Marcet, and Beutel (2017) introduce labor income shocks and subjective beliefs with persistent and transitory components in order to fit moments of asset prices and survey expectations.

[19]The limit theory would go with a simple scale factor adjustment. For example, we could specify $\varepsilon_m = \varepsilon$ and $\varepsilon_s = k_s\varepsilon$ and let $\varepsilon \to 0$ as in the theory. Then the limit dynamics for s_t^2 would be scaled by k_s.