Finance and Economics Discussion Series
Divisions of Research & Statistics and Monetary Affairs
Federal Reserve Board, Washington, D.C.

The Role of Learning for Asset Prices, Business Cycles, and Monetary Policy

Fabian Winkler

2016-019

Please cite this paper as:
Winkler, Fabian (2016). "The Role of Learning for Asset Prices, Business Cycles, and Monetary Policy," Finance and Economics Discussion Series 2016-019. Washington: Board of Governors of the Federal Reserve System, http://dx.doi.org/10.17016/FEDS.2016.019.

The Role of Learning for Asset Prices, Business Cycles, and Monetary Policy

Fabian Winkler*

20 January 2016

Abstract

The importance of financial frictions for the business cycle is widely recognized, but it is less recognized that their effects depend heavily on the underlying asset pricing theory. This paper examines the implications of learning-based asset pricing. I construct a model in which firms' ability to access credit depends on their market value, and investors rely on past observation to predict future stock prices. Agents' expectations remain model-consistent conditional on their beliefs about stock prices, which disciplines the expectation formation process. The model matches several asset price properties such as return volatility and predictability and also leads to a powerful feedback loop between asset prices and real activity, substantially amplifying business cycle shocks. Agents' expectational errors on asset prices spill over to forecasts of economic activity, resulting in forecast error predictability that closely matches survey data. A reaction of monetary policy to asset prices is welfare-improving under learning but not under rational expectations.

Keywords: Learning, Asset Pricing, Credit Constraints, Monetary Policy, Survey Data
JEL Classification: D83, E32, E44, E52, G12

*Board of Governors of the Federal Reserve System, 20th St and Constitution Ave NW, Washington DC 20551, fabian.winkler@frb.gov. I thank Wouter den Haan for his guidance and support; Klaus Adam, Andrea Ajello, Johannes Boehm, Peter Karadi, Albert Marcet, Stéphane Moyen, Rachel Ngai, Markus Riegler, and participants at the ECB-CFS-Bundesbank lunchtime seminar, Singapore Management University, Econometric Society Winter Meeting in Madrid, EEA Annual Meeting in Toulouse, and Econometric Society World Congress in Montreal, for useful comments and suggestions. The views herein are those of the author and do not represent the views of the Board of Governors of the Federal Reserve System or the Federal Reserve System.

1 Introduction

I think financial factors in general, and asset prices in particular, play a more central role in explaining the dynamics of the economy than is typically reflected in macro-economic models, even after the experience of the crisis.
— Andrew Haldane, Chief Economist of the Bank of England, 30 April 2014

The above statement may provoke disbelief among macroeconomists. After all, a wealth of research in the last 15 years has been dedicated precisely to the links between the financial sector and the real economy. Financial frictions are now seen as a central mechanism by which asset prices interact with macroeconomic dynamics. Yet our understanding of this interaction remains incomplete, in part due to the inherent difficulty of modelling asset prices.

Typical business cycle models still rely on an asset pricing theory based on rational expectations, time-separable preferences and moderate degrees of risk aversion. It is well known that such an asset pricing theory is inadequate for many empirical asset price regularities, which have therefore been named "puzzles". Some of the most famous such puzzles are that prices are volatile (Shiller, 1981), returns are predictable at business cycle frequency (Fama and French, 1988), and exhibit negative skewness and excess kurtosis (Campbell and Hentschel, 1992).[1] At the same time, asset prices play a central role in the real economy in the presence of financial frictions. Conclusions drawn from models with financial frictions but without a good asset pricing theory are therefore questionable.

There is not a shortage of theories that aim to explain asset prices. Most of them keep the rational expectations assumption and engineer preferences that deliver highly volatile discount factors. The dominant approaches are based on non-linear habit formation (Campbell and Cochrane, 1999) or Epstein-Zin preferences together with "long-run risk" (Bansal and Yaron, 2004). There are however compelling arguments for relaxing the rational expectations assumption instead. Measurements of expectations in surveys do not support the notion that agents have rational expectations. Rational expectations imply, for example, that investors are fully aware of return predictability in the stock market, expecting lower returns when prices are high and vice versa. Instead, measured expectations imply they expect higher returns. This pattern has been documented extensively by Greenwood and Shleifer (2014) and is illustrated in Figure 1. The left panel plots the mean 12-month return expectation of the S&P500, as measured in the Graham-Harvey survey of American CFOs, against the value of the P/D ratio in the month preceding the survey. The correlation is strongly positive: Return expectations are more optimistic when stock valuations are high. However, high stock valuations actually predict low future returns of the S&P500, as documented above and illustrated again in the right panel of the figure. Unless one rejects surveys as an unbiased measure of expectations, such a pattern cannot be reconciled with rational expectations.

[1] Another such puzzle is the size of the equity premium. Adam, Marcet, and Nicolini (2015) show that learning models similar to the one in this paper are able to generate a sizable equity premium, but it is not the subject of this paper.

Figure 1: Return expectations and expected returns.

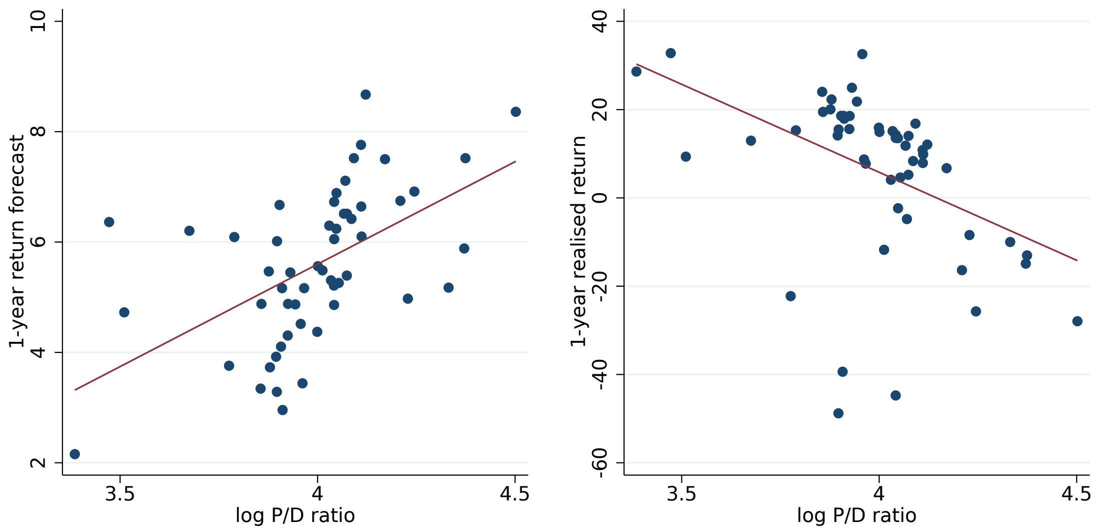

Expected nominal returns (left) are the mean response in the Graham-Harvey survey, realized nominal returns (right) and P/D ratio are from the S&P 500. Data period 2000Q32012Q4. Correlation coefficient for return forecasts $\rho = .54$, for realized returns $\rho = -.44$.

Based on such observations, Adam, Marcet, and Nicolini (2015) have developed an asset pricing theory based on learning. The interpretation of price dynamics under learning is quite different from rational expectations. Stock prices fluctuate not because of variations in the discounting of prices and returns, but because of variations in subjective beliefs about these prices and returns themselves. The deviation of subjective beliefs from rational expectations is a natural measure of "price misalignments." In an endowment economy, this approach is able to explain the most common asset price puzzles remarkably well. For example, return predictability arises because high asset valuations result from over-optimistic expectations. As expectations are corrected downwards, prices fall and returns are low.

In this paper, I examine the implications of a learning-based asset pricing theory for the business cycle. I construct a model of firm credit frictions in which agents are learning about price growth in the stock market, as in Adam, Marcet, and Nicolini (2015). At the same time, the model has a "financial accelerator" mechanism in which asset prices play a key role. Firms are subject to credit constraints, the tightness of which depends on its market value. This constraint emerges from a limited commitment problem in which defaulting firms can be restructured and resold as opposed to being liquidated. It provides a mechanism by which high stock market valuations translate into easier access to credit.

Deviating from the rational expectations hypothesis in a business cycle model is not without problems. One needs to explicitly spell out the entire belief formation process, filling many degrees of freedom. The existing learning literature often suffers from a lack of transparency in this respect, or abandons expected utility maximization altogether in favor of more reduced-form equilibrium conditions. To address this problem, I develop an expectation concept that I call "conditionally model-consistent expectations." This can be seen as a refinement of the "internal rationality" requirement developed by Adam and Marcet (2011). Agents continue to maximize a well-defined stable objective function with coherent and time-consistent beliefs about the variables affecting their decisions. They can entertain arbitrary beliefs about one relative price in the economy, which will be the price of stocks in this paper. But their beliefs about any other variable must be consistent with the equilibrium conditions of the model, except for market clearing in the stock market (and one more market, owing to Walras' law). This implies that when agents endowed with these expectations evaluate their forecast errors, they find that their forecasting rules cannot be improved upon conditional on their subjective belief about stock prices. In this sense this is a minimal departure from rational expectations. What's more, spelling out a belief system for stock prices and then imposing conditionally model-consistent expectations is all that is needed to obtain a unique dynamic equilibrium. This allows me to transparently incorporate asset price learning into any business cycle model while introducing only one additional parameter and one state variable.

The analysis of the model yields three results. First, a positive feedback loop emerges between asset prices and the production side of the economy, which leads to considerable amplification and propagation of business cycle shocks. When investor beliefs are more optimistic, their demand for stocks increases. This raises firm valuations and relaxes credit condition, in turn allowing firms to move closer to their profit optimum. If countervailing general equilibrium forces are not too strong, firms will be able to pay higher dividends to their shareholders, raising stock prices further and propagating investor optimism. The financial accelerator mechanism becomes quantitatively more important than under rational expectations. At the same time, learning greatly improves asset price properties such as price and return volatility and predictability. This result suggests that the relatively weak quantitative strength of the financial accelerator effect in many existing models—as discussed in (Cordoba and Ripoll, 2004)—is at least in part due to low endogenous asset price volatility.

Second, I compare the forecasts of agents in the model with actual forecasts in survey data, for both asset prices and real quantities. Even though agents only learn about stock prices in the model, their expectational errors spill over into their forecasts of other variables. For example, when agents are too optimistic about asset prices, they also become too optimistic about the tightness of credit constraints and therefore over-predict future investment. I find that the replicates remarkably well the predictability of forecast errors by forecast revisions as in Coibion and Gorodnichenko (2015), as well as by the level and growth rate of the price-dividend ratio. This lends credibility to the choice of the expectation formation process.

Third, I show that the model has important normative implications. A recurring question in monetary economics is whether policy should react to asset price "misalignments"; or "lean against the wind." Gali (2014) writes that justifying such a reaction requires "the presumption that an increase in interest rates will reduce the size of an asset price bubble," for which "no empirical or theoretical support seems to have been provided." This paper is a first step towards filling this gap. Indeed, I find that under learning, the welfare-maximizing monetary interest rate rule reacts strongly to asset price growth. By raising interest rates when stock prices are rising, policy is able to curb the endogenous build-up of over-optimistic investor beliefs, reducing both asset price and business cycle volatility. In contrast, under rational expectations, adding a policy reaction to asset prices carries no welfare benefits, in line with earlier findings in the literature (Bernanke and Gertler, 2001; Faia and Monacelli, 2007). The result highlights the importance of expectations in financial markets for normative questions.

There is a large literature studying the effect of asset prices on business cycle fluctuations, which relates to this paper. In particular, Xu, Wang, and Miao (2013) develop a model in which borrowing limits depend on stock market valuations through a credit friction similar to that in my model. They prove the existence of rational liquidity bubbles and introduce a shock that governs the size of this bubble, thus allowing them to exactly match the stock prices seen in the data. Other studies (e.g. Iacoviello, 2005; Liu, Wang, and Zha, 2013) similarly "explain" asset price fluctuations by direct shocks to prices in order to study the effects of financial frictions. In this paper, asset price volatility is not driven by an additional shock but instead an endogenous outcome of the learning dynamics.

There are some studies that endogenize asset price dynamics in production economies with rational expectations. They make use of habit formation (Boldrin, Christiano, and Fisher, 2001) or long-run risk (Tallarini Jr., 2000; Croce, 2014). The goal there is to replicate both standard asset price and business cycle moments within a close variant of the real business cycle model. This turns out to be a difficult task because the complex preferences in these models have some counterfactual implications (Lettau and Uhlig, 2000; Epstein, Farhi, and Strzalecki, 2013); to my knowledge, these asset price theories have not been applied to more elaborate business cycle models. The model in this paper can be solved conveniently in the presence of such frictions using a perturbation approach. Moreover, the inefficient nature of price fluctuations leads to policy implications that are quite different from an efficient markets world.

This paper also contributes to the literature on learning in business cycle models. A number of papers in this area have studied learning in combination with financial frictions (Caputo, Medina, and Soto, 2010; Milani, 2011; Gelain, Lansing, and Mendicino, 2013). The approach consists of two steps: first, derive the linearized equilibrium conditions of the economy under rational expectations; second, replace all terms involving expectations with parameterized forecast functions, and update the parameters using recursive least squares every period. Such models certainly produce very rich dynamics, but they are problematic on several grounds. First, such parameterized expectation equations often do not correspond anymore to intertemporal optimization problems. Second, analysis of these models is often complex and lack transparency due to the large number of parameters and state variables involved. Here, I develop a more transparent and parsimonious approach. Beliefs are restricted to be conditionally model-consistent and agents make optimal choices given a consistent set of beliefs, preserving much of the intuition of a rational expectations model.

The remainder of the paper is structured as follows. Section 2 presents a simplified version of the model that permits an analytic solution. It shows that credit frictions or asset price learning alone does not generate either amplification of shocks or interesting asset price dynamics, although their combination does. The full model is presented in Section 3. Section 4 contains the quantitative results. Section 5 contains sensitivity checks. Section 6 discusses monetary policy implications. Section 7 concludes.

2 Simplified model

In this section, I construct a simplified version of the model which illustrates the interaction between credit frictions and learning about asset prices. The model allows for a closed-form solution, but quantitative analysis will require extending it in the next section. The key insight here is that financial frictions alone do not generate sizable amplification of business cycle shocks or asset price volatility, while in combination with learning they do.

2.1 Model setup

Time is discrete at $t = 0, 1, 2, \ldots$. The model economy consists of a representative household and a representative firm. The household is risk-neutral and inelastically supplies one unit of labor. Its utility maximization program is as follows:

$$\max_{(C_t, S_t, B_t)_{t=0}^{\infty}} \mathbb{E}^{\mathcal{P}} \sum_{t=0}^{\infty} \beta^t C_t$$

$$\begin{aligned} \text{s.t. } C_t + S_t P_t + B_t &= w_t + S_{t-1}(P_t + D_t) + R_{t-1} B_{t-1} \\ S_t &\in \left[0, \bar{S}\right], S_{-1}, B_{-1} \end{aligned}$$

C_t is the amount of nondurable consumption goods purchased by the household in period t. The consumption good serves as the numéraire. w_t is the real wage rate. Moreover, the household can trade two financial assets: one-period bonds, denoted by B_t and paying gross real interest R_t in the next period; and stocks, S_t, which trade at price P_t and entitle their holder to dividend payments D_t. The household cannot short-sell stocks and his maximum stock holdings are capped at some $\bar{S} > 1$.[2] All markets are competitive.

The household maximizes the expectation of discounted future consumption under the probability measure $\mathcal{P}$. This measure is the subjective belief system held by agents in the model economy and might differ from rational expectations.

[2] The constraint on S_t is necessary to guarantee existence of the learning equilibrium, but never binds along the equilibrium path.

The firm engages in the production of the consumption good, which can also be used for investment. It is produced using capital K_{t-1}, owned by the firm and depreciating at the rate δ, and labor L_t according to the constant returns to scale technology

$$Y_t = K_{t-1}^{\alpha} \left(A_t L_t\right)^{1-\alpha}, \tag{2.1}$$

where A_t is its productivity. To allow for a closed-form solution, shocks to productivity are permanent:

$$\log A_t = \log A_{t-1} + \varepsilon_t,\, \varepsilon_t \sim iid\mathcal{N}\left(-\frac{\sigma^2}{2}, \sigma^2\right). \tag{2.2}$$

There are two financial claims on the firm: shares and noncontingent bonds. The firm's period budget constraint reads as follows:

$$Y_t + (1-\delta) K_{t-1} + B_t + S_t P_t = w_t L_t + K_t + S_{t-1} \left(P_t + D_t\right) + R B_{t-1} \tag{2.3}$$

I impose constraints on the issuance of financial instruments. On the equity side, the firm is not allowed to change the quantity of shares outstanding, fixed at $S_t = 1$. Further, it is not allowed to use retained earnings to finance investment. Instead, all earnings have to be paid out to shareholders:

$$D_t = Y_t - w_t L_t - \delta K_{t-1} - (R-1) B_{t-1}. \tag{2.4}$$

These assumptions imply that the firm's capital stock must be entirely debt-financed: Dividends are paid out until $K_t = B_t$ at the end of every period.[3] This is done purely to eliminate a state variable, and I relax these assumptions in the next section.

The firm's level of debt is limited to a fraction $\xi \in [0,1]$ of its total market value (i.e., the sum of debt and equity):

$$\begin{aligned} B_t &\leq \xi \left(B_t + P_t\right) \\ \Leftrightarrow K_t &\leq \frac{\xi}{1-\xi} P_t \end{aligned} \tag{2.5}$$

Equation (2.5) is a simple constraint on leverage: i.e., debt divided by the value of total assets. In the full version of the model, I will formally derive (2.5) from a limited commitment problem in which creditors can sell a firm as a going concern in the event of default, so that the market value of the firm enters the natural borrowing limit.

The firm maximizes the presented discounted sum of future dividends, using the household discount factor:

$$\max_{(K_t, L_t, D_t)_{t=0}^{\infty}} \mathbb{E}^{\mathcal{P}} \sum_{t=0}^{\infty} \beta^t D_t \text{ s.t. } (2.4),\ (2.5),\ K_{-1}$$

[3] For very low realizations of the productivity shock, the dividend payment will be negative, which is allowed. The value of the firm will be determined by *expected* dividends, which are always positive.

In particular, it makes its decisions under the same belief system $\mathcal{P}$ as the household—expectations are homogenous.

The model is closed by specifying market clearing conditions for the goods, labor and equity markets:

$$Y_t + K_t = C + (1-\delta) K_{t-1} \tag{2.6}$$

$$L_t = 1 \tag{2.7}$$

$$S_t = 1. \tag{2.8}$$

An equilibrium for an arbitrary subjective probability measure $\mathcal{P}$ is defined as a mapping from realizations of the exogenous variable $(A_t)_{t=0}^{\infty}$ and initial conditions (B_{-1}, K_{-1}, R_{-1}) to the endogenous variables $(B_t, K_t, L_t, D_t, P_t, R_t, w_t, C_t, S_t)_{t=0}^{\infty}$ such that markets clear and agents' choices solve their optimization problem under the probability measure $\mathcal{P}$. In particular, the first order conditions of the household and firm must hold at all times.[4]

The first-order conditions describing the household's optimal plan are $R_t = R = \beta^{-1}$ for bonds and

$$S_t \begin{cases} = 0 & \text{if } P_t > \beta \mathbb{E}_t^{\mathcal{P}} \left[P_{t+1} + D_{t+1}\right] \\ \in \left[0, \bar{S}\right] & \text{if } P_t = \beta \mathbb{E}_t^{\mathcal{P}} \left[P_{t+1} + D_{t+1}\right] \\ = \bar{S} & \text{if } P_t < \beta \mathbb{E}_t^{\mathcal{P}} \left[P_{t+1} + D_{t+1}\right] \end{cases} \tag{2.9}$$

for stocks. In a rational expectations setup, one would quickly substitute the market clearing condition $S_t = 1$ and write (2.9) as an equality. However, this already assumes that agents know how many outstanding shares they will hold in equilibrium. Under learning, they will not be endowed with this knowledge.

The firm's optimal labor demand and constraints imply the following expression for dividends:

$$D_t = \left(R_t^k - R_{t-1}\right) K_{t-1}, \tag{2.10}$$

where the marginal return on capital is

$$R_t^k = \alpha \left((1-\alpha) \frac{A_t}{w_t}\right)^{\frac{1-\alpha}{\alpha}} + 1 - \delta. \tag{2.11}$$

When choosing the optimal amount of capital then, the firm exhausts its borrowing limit as long

[4] This definition satisfies the "internal rationality" requirement of Adam and Marcet (2011).

as the expected internal return on capital exceeds the external return paid to creditors:

$$K_t \begin{cases} = 0 & \text{if } \mathbb{E}_t^{\mathcal{P}} R_{t+1}^k < R_t \\ \in \left[0, \frac{\xi}{1-\xi} P_t\right] & \text{if } \mathbb{E}_t^{\mathcal{P}} R_{t+1}^k = R_t \\ = \frac{\xi}{1-\xi} P_t & \text{if } \mathbb{E}_t^{\mathcal{P}} R_{t+1}^k > R_t. \end{cases} \tag{2.12}$$

These optimality conditions must hold under any belief $\mathcal{P}$ used to form expectations.

2.2 Rational expectations equilibrium

I first describe the equilibrium under rational expectations. Rational expectations amount to a particular choice of the measure $\mathcal{P}$. This measure has to coincide with the measure induced by the equilibrium allocations. In that case, one writes $\mathbb{E}^{\mathcal{P}}[\cdot] = \mathbb{E}[\cdot]$.

The equilibrium under rational expectations admits a closed-form solution. First, let us consider the case $\xi = 1$. In this case, the borrowing constraint (2.5) can never bind. Capital is therefore at its efficient level, which is simply proportional to productivity: $K_t = K^* A_t$ such that $\mathbb{E}_t R_{t+1}^k = R$. Expected dividends and the market value of equity are zero: $\mathbb{E}_t D_{t+1} = 0$ and $P_t = 0$. Intuitively, when the firm is unconstrained, the expected return on capital must equal the interest rate on borrowing, and since all capital is financed by debt and the production function has constant returns to scale, in expectation all profits are paid out as interest payments to debt holders. The residual equity claims trade at a price of zero.

Once we introduce financial frictions by setting $\xi < 1$, how much amplification do we get? The answer is: none. Now the equilibrium is characterized by two equations:

$$\frac{P_t}{A_t} = \bar{P} = \frac{\exp\left(-\frac{\alpha(1-\alpha)}{2}\sigma^2\right)\alpha\bar{K}^{\alpha} - (R-1+\delta)\bar{K}}{R} \tag{2.13}$$

$$\frac{K_t}{A_t} = \bar{K} = \frac{\xi}{1-\xi}\bar{P} \tag{2.14}$$

The first equation pins down the stock market value of the firm as a function of the capital stock. The second equation determines the capital stock that can be reached by exhausting the borrowing constraint that depends on the stock market value. In particular, the borrowing constraint is always binding. In equilibrium, the equilibrium capital stock and stock price comove perfectly with productivity, just as in the case of $\xi = 1$. Financial frictions do not lead to any amplification or propagation of shocks in the rational expectations equilibrium. They have a *level* effect on output, capital, etc., but the *dynamics* of the model are identical for any value of ξ. The variances of log stock price and output growth do not depend on ξ and are bounded by the variance of the exogenous shock σ^2:

$$\begin{aligned}\text{Var}\left[\Delta \log P_t\right] &= \sigma^2 && (2.15)\\ \text{Var}\left[\Delta \log Y_t\right] &= \left(1-2\alpha+2\alpha^2\right)\sigma^2 && (2.16)\end{aligned}$$

Intuitively, with financial frictions, a shock to productivity raises asset prices just as much as to allow the firm to instantly adjust the capital stock proportionately. At the same time, stock returns are not volatile and unpredictable at all horizons.

Before moving on to the learning equilibrium, it is worth noting that the stock price and the dividend payment of the firm are non-monotonic in the level of financial frictions ξ:

$$\mathbb{E}_t\left[D_{t+1}\right] = D\left(k_t, K_t, A_t\right) = \left(\mathbb{E}_t \alpha \left(\frac{A_{t+1}}{K_t}\right)^{1-\alpha} + 1 - \delta - R\right) k_t. \tag{2.17}$$

Here, I have made a distinction between the capital choice k_t of the firm that takes future wages as given, and the aggregate capital stock K_t, which determines wages and the return on capital in general equilibrium. Of course, in equilibrium the two are equal. When $\xi = 0$, the firm cannot borrow at all and $k_t = 0$, and when $\xi = 1$ we have $\mathbb{E}_t R^k_{t+1} = R$. In both cases, expected dividends and the value of the firm are zero. In between these extreme cases, there are two opposing forces affecting expected dividends:

$$\frac{d}{d\xi} D\left(\bar{K}, \bar{K}, A_t\right) = \left[\underbrace{\frac{\partial D}{\partial k}\left(k_t, K_t, A_t\right)}_{>0} + \underbrace{\frac{\partial D}{\partial K}\left(k_t, K_t, A_t\right)}_{<0}\right] \underbrace{\frac{d\bar{K}}{d\xi}}_{>0}. \tag{2.18}$$

The first term in brackets captures a partial equilibrium effect of leverage, which is internalized by the firm. When a firm is financially constrained, its internal rate of return is higher than the return it has to pay to debt holders. It will then want to increase its capital stock by borrowing until it reachtes the borrowing constraint, thereby increasing expected dividends. The second term captures a general equilibrium effect: Higher investment lowers the marginal product of capital, which in practice is realized through an increase in the equilibrium wage w_{t+1}, reducing expected dividends. When ξ is small (financial frictions are severe) the partial equilibrium effect dominates, while for a large ξ, the general equilibrium effect dominates.

2.3 Learning equilibrium

I now introduce learning about stock prices by changing the subjective probability measure $\mathcal{P}$. Learning will increase the volatility of stock prices, account for return and forecast error predictability, and, most importantly, induce endogenous amplification and propagation on the production side of the model in combination with financial frictions.

The only departure from rational expectations is that agents do not understand the pricing function that maps fundamentals into an equilibrium stock price. Agents are not endowed with the knowledge that prices obey the market-clearing condition:

$$P_t = \beta \mathbb{E}_t^{\mathcal{P}} \left[P_{t+1} + D_{t+1} \right]. \tag{2.19}$$

Instead, they form subjective beliefs about the law of motion of prices and update them using realized price observations. I impose the following restrictions on these beliefs. Under the subjective measure $\mathcal{P}$,

1. agents have the correct belief about the fundamental A_t;
2. agents believe that the stock price P_t evolves according to

$$\log P_t - \log P_{t-1} = \mu_t + \eta_t \tag{2.20}$$

$$\mu_t = \mu_{t-1} + \nu_t \tag{2.21}$$

$$\text{where} \begin{pmatrix} \eta_t \\ \nu_t \end{pmatrix} \sim \mathcal{N}\left(-\frac{1}{2} \begin{pmatrix} \sigma_\eta^2 \\ \sigma_\nu^2 \end{pmatrix}, \begin{pmatrix} \sigma_\eta^2 & 0 \\ 0 & \sigma_\nu^2 \end{pmatrix} \right) iid, \tag{2.22}$$

   the variable μ_t and the disturbances η_t and ν_t are unobserved and the prior about μ_t in period 0 is given by

$$\mu_0 \mid \mathcal{F}_0 \sim \mathcal{N}\left(\hat{\mu}_0, \sigma_\mu^2 \right) \text{ where } \sigma_\mu^2 = \frac{-\sigma_\nu^2 + \sqrt{\sigma_\nu^4 + 4\sigma_\nu^2 \sigma_\eta^2}}{2}; \tag{2.23}$$

3. agents update their beliefs about μ_t after making their choices and observing equilibrium prices in period t;
4. for any variable x_t relevant for agents' decision problems and $t \geq 0$, agents' beliefs coincide with model outcomes on the equilibrium path, *conditional* on the realization of stock prices and fundamentals:

$$\mathbb{E}^{\mathcal{P}} \left[x_t \mid A_0, \ldots, A_t, P_0, \ldots, P_t \right] = x_t.$$

The first assumption implies that agents have as much information about the fundamental shocks of the economy as under rational expectations. The second assumption amounts to saying that agents believe stock prices to be a random walk. This random walk is believed to have a small, unobservable, and time-varying drift μ_t. Learning about this drift is going to be the key driver of asset price dynamics. The third assumption imposes that forecasts of stock prices are updated after equilibrium prices are determined.[5] Together, (2.20)–(2.23) form a linear state-space system. Under $\mathcal{P}$, agents' beliefs about μ_t at time t are normally distributed with stationary variance σ_μ^2

[5] This "lagged belief updating" is common in the learning literature. It makes all feedback between forecasts and prices inter- rather than intratemporal. For further discussion see Adam, Beutel, and Marcet (2014).

and mean $\hat{\mu}_{t-1}$, which evolves according to the Kalman filtering equation:

$$\hat{\mu}_t = \hat{\mu}_{t-1} - \frac{\sigma_\nu^2}{2} + g\left(\log P_t - \log P_{t-1} + \frac{\sigma_\eta^2 + \sigma_\nu^2}{2} - \hat{\mu}_{t-1}\right) \tag{2.24}$$

In this equation, P_t and P_{t-1} are observed, realized stock prices. These are determined in equilibrium under the actual law of motion of the economy and will generally not follow the perceived law of motion described by (2.20)–(2.23).[6] The parameter g is the learning gain. It governs the speed with which agents move their prior in the direction of the last forecast error.[7] Note that the learning gain is constant rather than diminishing over time, and so beliefs never converge. Intuitively, even after observing a long time series, agents discount data from the distant past as they see believe the growth rate of asset prices to vary continuously over time.

To the best of my knowledge, the fourth assumption is novel to the learning literature. To make their choices, agents must form expectations about more than just the stock price and the exogenous processes of the economy. For example, investors need to form a belief about dividends, which are endogenous and depend on wages and the capital choice of the firm in equilibrium. What I assume here is that conditional on a realization of stock prices on the equilibrium path, agents forecasts are consistent with equilibrium outcomes. This retains the logic of rational expectations while still allowing for learning about asset prices.

More specifically, let y_t be the collection of all endogenous model variables and ε_t the collection of iid exogenous shocks distributed with cdf F. Agents form their beliefs about the future using a probability measure $\mathcal{P}$. Assume that there exists a recursive equilibrium defined by the policy function $y_t = g(y_{t-1}, \varepsilon_t)$. In addition, the stock price in equilibrium follows the law of motion, $P_t = g_P(y_{t-1}, \varepsilon_t)$. Under learning, agents are not endowed with knowledge of the mappings g and g_P from fundamental shocks and state variables to equilibrium outcomes and prices. I assume that instead, agents' subjective beliefs $\mathcal{P}$ are defined by the distribution of the shocks ε_t, and recursive function that define the subjective law of motion for the price $P_t = h_P(y_{t-1}, \varepsilon_t)$ and for the other endogenous variables conditional on the price: $y_t = h(y_{t-1}, \varepsilon_t, P_t)$. In other words,

$$\mathbb{E}_t^{\mathcal{P}}\left[y_{t+1} \mid \varepsilon_{t+1}, P_{t+1}\right] = h\left(y_t, \varepsilon_{t+1}, P_{t+1}\right). \tag{2.25}$$

$$\mathbb{E}_t^{\mathcal{P}}\left[P_{t+1} \mid \varepsilon_{t+1}\right] = h_P\left(y_t, \varepsilon_{t+1}\right). \tag{2.26}$$

Imposing conditionally model-consistent expectations amounts to the following requirement for subjective beliefs:

$$h\left(y_t, \varepsilon_{t+1}, g_P\left(y_{t,}, \varepsilon_{t+1}\right)\right) = g\left(y_t, \varepsilon_{t+1}\right). \tag{2.27}$$

[6]The exception is when $g = 0$ and $\hat{\mu}_0 = -\sigma^2/2$, in which case agents the subjective belief system coincides with rational expectations.

[7]The gain is related to the variances of the disturbances by the formula $g = \left(1 + 2\left(\frac{\sigma_\nu^2}{\sigma_\eta^2} + \sqrt{\frac{\sigma_\nu^4}{\sigma_\eta^4} + 4\frac{\sigma_\nu^2}{\sigma_\eta^2}}\right)^{-1}\right)^{-1}$, and is strictly increasing in the signal-to-noise ratio $\sigma_\nu^2/\sigma_\eta^2$.

Conditional on the price P_t coinciding with its equilibrium realization $g_P(y_{t,}, \varepsilon_{t+1})$, agents' other forecasts are correct. I solve for these beliefs in much the same way as one would under rational expectations; the procedure is discussed in Section 3 and Appendix C.

It is important to note that this is still different from rational expectations, since agents' subjective beliefs about the price P_t differ from those under rational expectations: Because $h_P \neq g_P$, it is generally the case that

$$\mathbb{E}_t^{\mathcal{P}}[y_{t+1}] = \int h(y_t, \varepsilon, h_P(y_t, \varepsilon_t)) F(d\varepsilon) \neq \int g(y_t, \varepsilon_t) F(d\varepsilon) = \mathbb{E}_t[y_{t+1} \mid \varepsilon_{t+1}]. \tag{2.28}$$

In other words, forecast errors about stock prices spill over into forecasts about other variables as well. Since agents' beliefs about stock prices produce biased forecast errors, their *unconditional* forecasts about other endogenous variables are also biased, even if the conditional forecasts are not.

I now compute the equilibrium with learning and conditionally model-consistent expectations. Start with the stock market clearing condition , which has to hold even though agents think that prices just follow the system (2.20)–(2.23)). The market clearing condition (2.19) now reads as follows:

$$P_t = \frac{\mathbb{E}_t^{\mathcal{P}} P_{t+1} + \mathbb{E}_t^{\mathcal{P}} D_{t+1}}{R} \tag{2.29}$$

$$= \frac{P_t \exp\left(\hat{\mu}_{t-1} + \frac{1}{2}\sigma_\mu^2\right) + D(K_t, K_t, A_t)}{R} \tag{2.30}$$

$$= A_t \frac{D(K_t, K_t, 1)}{R - \exp\left(\hat{\mu}_{t-1} + \frac{1}{2}\sigma_\mu^2\right)} \tag{2.31}$$

In most of the learning literature, one would need to separately specify beliefs about expected prices and expected dividends. With conditionally model-consistent expectations, agents are endowed with the knowledge of equilibrium wages and dividends conditional on a realization of future stock prices. Therefore, their expectations about dividends depend on the current capital stock in the same way as under rational expectations, so that one can substitute in the function D as before.

The remaining equilibrium conditions are static and therefore independent of expectations. The model is summed up by the following three equations:

$$P_t = \frac{D(K_t, K_t, A_t)}{R - \exp\left(\hat{\mu}_{t-1} + \frac{1}{2}\sigma_\mu^2\right)} \tag{2.32}$$

$$K_t = \frac{\xi}{1-\xi} P_t \tag{2.33}$$

$$\hat{\mu}_t = \hat{\mu}_{t-1} - \frac{\sigma_\nu^2}{2} + g\left(\log \frac{P_t}{P_{t-1}} - \hat{\mu}_{t-1} + \frac{\sigma_\eta^2 + \sigma_\nu^2}{2}\right) \tag{2.34}$$

Clearly, the stock price and capital stock are not proportional to productivity under learning. Figure 2 depicts the dynamics of stock prices after a positive innovation at $t = 1$. The initial shock

Figure 2: Stock price dynamics under learning.

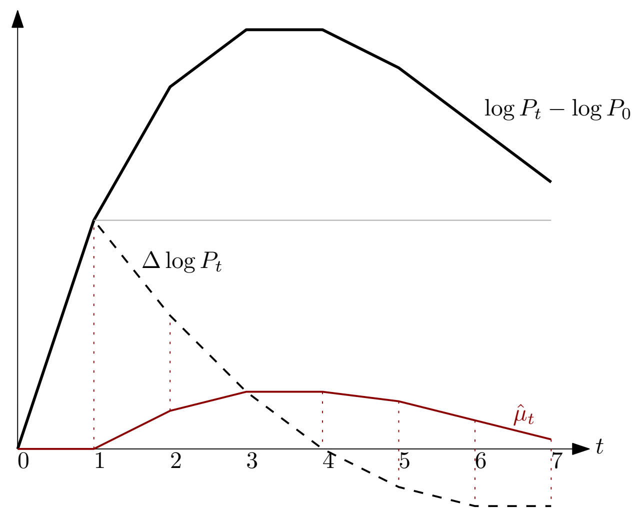

at $t = 1$ raises stock prices proportionally to productivity, just as under rational expectations. But learning investors are unsure whether the rise in P_1 is indicative of a transitive or permanent increase in the growth rate of stock prices. They therefore revise their beliefs upward. In $t = 2$ then, demand for stocks is higher and stock prices need to rise further to clear the market. Beliefs continue to rise as long as observed asset price growth (dashed black line in Figure 2) is higher than the current belief $\hat{\mu}_t$ (solid red line). The differences between observed and expected price growth are the forecast errors (dotted red lines). In the figure, the increase in prices and beliefs ends at $t = 3$, when the forecast error is zero. There is no need for a further belief revision. Now, by equations (2.32)–(2.33), in the absence of subsequent shocks, P_t just co-moves with beliefs $\hat{\mu}_{t-1}$, so when there is no belief revision at $t = 4$, realized asset price growth is also zero. This triggers an endogenous reversal in prices, as investors observe stalling asset prices at the peak of their optimism. They subsequently revise their beliefs $\hat{\mu}_t$ downward, pushing the stock price down until it returns to its steady-state level.

Learning leads to volatility and return predictability. To see this, it is convenient to look at the forward P/D ratio:

$$\frac{P_t}{\mathbb{E}_t^{\mathcal{P}} D_{t+1}} = \frac{1}{R - \exp\left(\hat{\mu}_{t-1} + \frac{1}{2}\sigma_\mu^2\right)}.$$

The forward P/D ratio is directly related to the belief $\hat{\mu}_t$, and even small changes in this belief can have a large impact on the P/D ratio as the denominator is close to zero. Furthermore, since the system (2.32)–(2.33) is stationary, a high P/D ratio predicts a future decline in $\hat{\mu}_t$ and therefore falling prices and low returns.

The aforementioned asset pricing implications are present even when dividends are completely exogenous, as in Adam, Marcet, and Nicolini (2015). But the model considered here also contains a link among asset prices, output and dividends. The capital stock K_t is directly related to equity valuations P_t through Equation (2.33). Thus, the fluctuations in the stock market translate into

Figure 3: Endogenous response of dividends.

(a) Amplification. (b) Dampening.

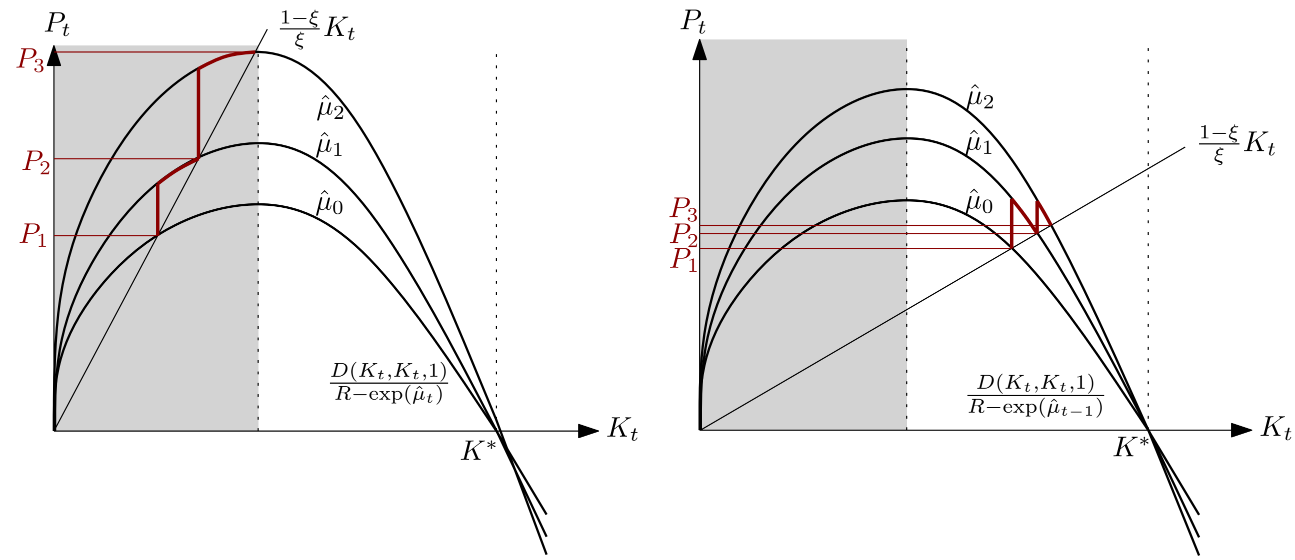

corresponding fluctuations in investment, the capital stock and hence output.

It is also possible that this amplification mechanism is *further* enhanced by positive feedback from capital to expected dividends. This additional feedback, however, depends on the slope of the dividend function. The expected dividend $D(k_t, K_t, A_t)$ is increasing in the firm's capital choice k_t but decreasing in the aggregate capital stock K_t. In equilibrium ($k_t = K_t$) it is increasing in aggregate capital only if financial frictions are sufficiently severe. This case is depicted in Panel (a) of Figure 3. When the degree of financial frictions is high, the credit constraint line is steep. Assume that the initial equilibrium in period 0 is at P_0 and $\hat{\mu}_0$ and consider the effect of a positive productivity shock. The immediate effect will be a proportionate rise in stock prices and capital, together with a rise in beliefs from $\hat{\mu}_0$ to $\hat{\mu}_1$. This leads to higher stock prices at $t = 2$ and allows the firm to invest more and increase its expected profits—the $\partial D/\partial k$ effect dominates. This adds to the rise in realized stock prices, further relaxing the borrowing constraint and increasing next period's beliefs. Stock prices, investment, and output all rise more than proportionally to productivity.[8]

However, this additional amplification channel only works when ξ is sufficiently low. In Panel (b), ξ is large and the firm is operating in the downward-sloping part of the dividend curve. A relaxation of the borrowing constraint due to a rise in $\hat{\mu}$ still allows the firm to invest and produce more, but dividends fall in equilibrium—the $\partial D/\partial K$ effect dominates. Decreasing returns to capital at the aggregate level manifest through a rise in the wage, effectively reducing the firm's profits. The response of dividends dampens the dynamics of investment and asset prices.

[8] To my knowledge, this paper is the first to establish a positive feedback from fundamentals to beliefs under learning. Adam, Kuang, and Marcet (2012) also model economies with endogenous fundamentals. Their learning specification is similar, but the "dividend" in their asset pricing equation is simply the marginal utility of housing, which is strictly decreasing in the level of the housing stock.

I end this section with the observation that the equilibrium dynamics are driven by the interaction of financial frictions and learning. Neither is able to achieve amplification alone. When taking the limit $\xi \to 1$ of vanishing financial frictions, the entire amplification mechanism disappears:

$$\frac{d \log P_t}{d\varepsilon_{t-s}} \xrightarrow{\xi \to 1} 1. \tag{2.35}$$

Intuitively, as financial frictions disappear, the economy moves into a region where the general equilibrium effects become so strong that any potential rise in price growth beliefs is countered by a fall in expected dividends. The degree of amplification then drops back to zero.

3 Full model for quantitative analysis

This section embeds the mechanism discussed so far into a quantitative New-Keynesian model with a financial accelerator. Compared with the simple model in the previous section, there are a number of additional elements here. First, firms are allowed to finance capital out of retained earnings. Second, the borrowing constraint is generalized and microfounded by a limited commitment problem. Third, I add nominal rigidities and investment adjustment costs to improve the quantitative fit. I characterize in turn the rational expectations and the learning equilibrium, and then choose parameters for the model based partly on calibration and partly on estimation by simulated method of moments.

3.1 Model setup

The economy is closed and operates in discrete time. It is populated by two types of households.

1. *Lending households* consume final goods and supply labor. They trade debt claims on intermediate goods producers and receive interest from them.

2. *Firm owners* only consume final goods. They trade equity claims on intermediate goods producers and receive dividends from them.

The two households own four types of firms. Only the first type is substantial to the model analysis; the other three serve to introduce nominal rigidities and adjustment costs to the model.

1. *Intermediate goods producers* (or simply *firms*) are at the heart of the model. They combine capital and differentiated labor to produce a homogeneous intermediate good. They are financially constrained and borrow funds from households.

2. *Labor agencies* transform homogeneous household labor into differentiated labor services, which they sell to intermediate goods producers. They are owned by households.

3. *Final good producers* transform intermediate goods into differentiated final goods. They are owned by households.

4. *Capital goods producers* produce new capital goods from final consumption goods subject to an investment adjustment cost. They are owned by households.

Finally, there is a *fiscal authority* setting tax rates to offset steady-state distortions from monopolistic competition, and a *central bank* setting nominal interest rates.

Most elements of the model are standardand their desription is relegated to Appendix A. Here I will describe mainly households, intermediate goods producers, and firm owners. I also discuss the microfoundation of the borrowing constraint.

3.1.1 Households

A representative household with time-separable preferences maximizes utility as follows:

$$\max_{(C_t,L_t,B_{jt},B_t^g)_{t=0}^{\infty}} \mathbb{E}_0^{\mathcal{P}} \sum_{t=0}^{\infty} \beta^t \log\left(C_t\right) - \eta \frac{L_t^{1+\phi}}{1+\phi}$$

$$\text{s.t. } C_t \quad = \quad \tilde{w}_t L_t + B_t^g - (1+i_{t-1}) \frac{p_{t-1}}{p_t} B_{t-1}^g + \int_0^1 \left(B_{jt} - R_{jt-1} B_{jt-1}\right) dj + \Pi_t$$

The utility function u satisfies standard concavity and Inada conditions and $\beta \in (0,1)$. Further, $\tilde{w}_t$ is the real wage received by the household and L_t is the amount of labor supplied. B_t^g are real quantities of nominal one-period government bonds (in zero net supply) that pay a nominal interest rate i_t and p_t is the price level, defined below. Households also lend funds B_{jt} to intermediate goods producers indexed by $j \in [0,1]$ at the real interest rate R_{jt}. These loans are the outcome of a contracting problem described later on. Π_t represents lump-sum profits and taxes. Finally, consumption C_t is itself a utility flow from a variety of differentiated goods that takes the familiar constant elasticity of substitution (CES) form:

$$\begin{aligned} C_t &= \max_{C_{it}} \left(\int_0^1 \left(C_{it}\right)^{\frac{\sigma-1}{\sigma}} di \right)^{\frac{\sigma}{\sigma-1}} \\ \text{s.t. } p_t C_t &= \int_0^1 p_{it} C_{it} di \end{aligned}$$

As usual, the price index p_t of composite consumption consistent with utility maximization and the demand function for good i is given by

$$p_t = \left(\int_0^1 \left(p_{it}\right)^{1-\sigma} di \right)^{\frac{1}{1-\sigma}} ; \; C_{it} = \left(\frac{p_{it}}{p_t}\right)^{-\sigma} C_t. \tag{3.1}$$

Consequently, the inflation rate is given by $\pi_t = p_t/p_{t-1}$. The first-order conditions of the household are also standard and given by

$$\tilde{w}_t = \eta L_t^{\phi} C_t \tag{3.2}$$

$$1 = \beta \mathbb{E}_t^{\mathcal{P}} \frac{C_t}{C_{t+1}} \frac{1+i_t}{\pi_t}. \tag{3.3}$$

We can define the stochastic discount factor of the households as $\Lambda_{t+1} = \beta C_t / C_{t+1}$.

3.1.2 Intermediate good producers (firms)

The production of intermediate goods is carried out by a continuum of firms, indexed $j \in [0,1]$. Firm j enters period t with capital K_{jt-1} and a stock of debt B_{jt-1} which needs to be repaid at the gross real interest rate R_{jt-1}. First, capital is combined with a labor index L_{jt} to produce output:

$$Y_{jt} = (K_{jt-1})^{\alpha} (A_t L_{jt})^{1-\alpha}, \tag{3.4}$$

where A_t is aggregate productivity. The labor index is a CES combination of differentiated labor services parallel to the differentiated final goods bought by the household:

$$L_{jt} = \max_{L_{jht}} \left(\int_0^1 (L_{jht})^{\frac{\sigma_w - 1}{\sigma_w}} dh \right)^{\frac{\sigma_w}{\sigma_w - 1}} \tag{3.5}$$

$$\text{s.t. } w_t p_t L_{jt} = \int_0^1 W_{jht} L_{jht} dh \tag{3.6}$$

The firm's problem can then be treated as if the labor index was acquired in a competitive market at the real wage index w_t.[9] Output is sold competitively to final good producers at price q_t. During production, the capital stock depreciates at rate δ. This depreciated capital can be traded by the firm at the price Q_t.

The firm's net worth is the difference between the value of its assets and its outstanding debt:

$$N_{jt} = q_t Y_{jt} - w_t L_{jt} + Q_t (1-\delta) K_{jt-1} - R_{jt-1} B_{jt-1}. \tag{3.7}$$

I assume that the firm exits with probability γ. This probability is exogenous and independent across time and firms. As in Bernanke, Gertler, and Gilchrist (1999), exit prevents firms from becoming financially unconstrained. If a firm does not exit, it needs to pay out a fraction $\zeta \in (0,1)$ of its earnings as dividends (where earnings are given by $E_{jt} = N_{jt} - Q_t K_{jt-1} + B_{jt-1}$). The number ζ therefore represents the dividend payout ratio for continuing firms.[10] If a firm does exit, it must

[9] This real wage index does not necessarily equal the wage $\tilde{w}_t$ received by households due to wage dispersion.

[10] The optimal dividend payout ratio in this model would be $\zeta = 0$, as firms would always prefer to build up net worth to escape the borrowing constraint over paying out dividends. However, this would imply that aggregate dividends would be proportional to aggregate net worth, which is rather slow-moving. The resulting dividend process

pay out its entire net worth as dividends. It is subsequently replaced by a new firm, which receives the index j. I assume that this new firm gets endowed with a fixed number of shares, normalized to one, and is able to raise an initial amount of net worth. This amount equals $\omega\left(N_t - \zeta E_t\right)$, where $\omega \in (0,1)$ and N_t and E_t are aggregate net worth and earnings, respectively.[11]

The net worth of firm j after equity changes, entry and exit is given by

$$\tilde{N}_{jt} = \begin{cases} N_{jt} - \zeta E_{jt} & \text{for continuing firms,} \\ \omega\left(N_t - \zeta E_t\right) & \text{for new firms.} \end{cases}$$

This firm then decides on the new stock of debt B_{jt} and the new capital stock K_{jt}, maximizing the present discounted value of dividend payments using the discount factor of its owners. Its balance sheet must satisfy:

$$Q_t K_{jt} = B_t^j + \tilde{N}_{tj}. \tag{3.8}$$

3.1.3 Firm owners

Firm owners differ from households in their capacity to own intermediate firms. The representative firm owner is risk-neutral. It can buy shares in firms indexed by $j \in [0,1]$. As described above, when a firm exits, it pays out its net worth N_{jt} as dividends, and is replaced by a new firm, which raises equity $\omega\left(N_t - \zeta E_t\right)$. Let the set of exiting firms in each period t be denoted by $\Gamma_t \subset [0,1]$. Then, the firm owner's utility maximization problem is given by:

$$\max_{\left(C_t^f, S_t^j\right)_{t=0}^{\infty}} \mathbb{E}_0^{\mathcal{P}} \sum_{t=0}^{\infty} \beta^t C_t^f$$

$$\text{s.t. } C_t^f + \int_0^1 S_{jt} P_{jt} dj = \int_{j \notin \Gamma_t} S_{jt-1}\left(P_{jt} + D_{jt}\right) dj \tag{3.9}$$

$$+ \int_{j \in \Gamma_t} \left[S_{jt-1} D_{jt} - \omega\left(N_t - \zeta E_t\right) + P_{jt}\right] dj \tag{3.10}$$

$$S_t^j \in \left[0, \bar{S}\right] \tag{3.11}$$

for some $\bar{S} > 1$. Here, firm owners' consumption C_t^f is the same aggregator of differentiated final goods as for households.

The first term on the right-hand side of the budget constraint deals with continuing firms and is standard: Each share in firm j pays dividends D_{jt} and continues to trade, at price P_{jt}. The second term deals with firm entry and exit. If the household owns a share in the exiting firm j, it receives

would not be nearly as volatile as in the data. Imposing $\zeta > 0$ allows to better match the volatility of dividends and therefore obtain better asset price properties.

[11] The simplified firm problem of Section 2 is nested as the case$\zeta = 1$ and $\gamma = 0$.

a terminal dividend. At the same time, a new firm j appears that is able to raise a limited amount of equity $\omega\left(N_t - \zeta E_t\right)$ from the firm owner in exchange for a unit amount of shares that can be traded at price P_{jt}. In addition, upper and lower bounds on traded stock holdings are introduced to make firm owners' demand for stocks finite under arbitrary beliefs, as in the stylized model of the previous section. In equilibrium, they are never binding.

The first-order conditions of the firm owner are

$$S_{jt} \begin{cases} = 0 & \text{if } P_{jt} > \beta \mathbb{E}_t^{\mathcal{P}}\left[D_{jt+1} + P_{jt+1}\mathbb{1}_{\{j \notin \Gamma_{t+1}\}}\right] \\ \in \left[0, \bar{S}\right] & \text{if } P_{jt} = \beta \mathbb{E}_t^{\mathcal{P}}\left[D_{jt+1} + P_{jt+1}\mathbb{1}_{\{j \notin \Gamma_{t+1}\}}\right] \\ = \bar{S} & \text{if } P_{jt} < \beta \mathbb{E}_t^{\mathcal{P}}\left[D_{jt+1} + P_{jt+1}\mathbb{1}_{\{j \notin \Gamma_{t+1}\}}\right] \end{cases} \tag{3.12}$$

3.1.4 Borrowing constraint

In choosing their debt holdings, firms are subject to a borrowing constraint. The constraint is the solution to a particular limited commitment problem in which the outside option for the lender in the event of default depends on equity valuations.

Each period, lenders (households) and borrowers (firms) meet to decide on the lending of funds. Pairings are anonymous. Contracts are incomplete because the repayment of loans cannot be made contingent. Only the size B_{jt} and the interest rate R_{jt} of the loan can be contracted in period t. Both the lender (a household) and the firm have to agree on a contract (B_{jt}, R_{jt}). Moreover, there is limited commitment in the sense that at the end of the period, but before the realization of next period's shocks, firm j can always choose to enter a state of default. In this case, the value of the debt repayment must be renegotiated. If the negotiations are successful, then wealth is effectively shifted from creditors to debtors. The outside option of this renegotiation process is bankruptcy of the firm and seizure by the lender.

Bankruptcy carries a cost of a fraction $1 - \xi$ of the firm's capital being destroyed. The lender, a household, does not have the ability to operate the firm. It can liquidate the firm's assets, selling the remaining capital in the next period. This results in a recovery value of $\xi Q_{t+1} K_{jt}$. With some probability x (independent across time and firms), the lender receives the opportunity to "restructure" the firm if it wants. Restructuring means that, similar to Chapter 11 bankruptcy proceedings, the firm gets partial debt relief but remains operational. I assume that the lender has to sell the firm to another firm owner, retaining a fraction ξ of the initial debt. In equilibrium, the recovery value in this case will be $\xi\left(P_{jt} + B_{jt}\right)$ and this will always be higher than the recovery value after liquidation. Thus, the debt contract takes the form of a leverage constraint in which total firm value is a weighted average of liquidation and market value:

$$B_{jt} \leq \xi \left(x \underbrace{\mathbb{E}_t^{\mathcal{P}} \Lambda_{t+1} Q_{t+1} \xi K_{jt}}_{\text{liquidation value}} + (1-x) \underbrace{\left(P_{jt} + B_{jt}\right)}_{\text{market value}} \right) \tag{3.13}$$

3.1.5 Central bank

The model is cashless, with the central bank setting the nominal interest rate according to a Taylor-type interest rate rule:

$$i_t = \rho_i i_{t-1} + (1-\rho_i)\left(\beta^{-1} + \phi_\pi \pi_t + \varepsilon_{it}\right), \tag{3.14}$$

where ϕ_π is the reaction coefficient on inflation, ρ_i is the degree of interest rate smoothing, and ε_{it} is an interest rate shock.

3.1.6 Further model elements and market clearing

Final good producers, indexed by $i \in [0,1]$, combine the homogeneous intermediate good into a differentiated final good using a one-for-one technology. Their revenue is subsidized by the government at the rate τ.[12] Per-period profits of producer i are $\Pi^Y_{it} = (1+\tau)(p_{it}/p_t) Y_{it} - q_t Y_{it}$. They are subject to a Calvo price-setting friction: Every period, each final good producer can change his price only with probability $1-\kappa$, independent across time and producers. Similarly, labor agencies (indexed by $h \in [0,1]$) combine the homogeneous labor provided by households into differentiated labor goods which they sell on to intermediate good producers. labor agencies' revenue is subsidized at the rate τ_w, the per-period profit of agency h is $\Pi^L_{ht} = (1+\tau)(W_{ht}/p_t) L_{ht} - \tilde{w}_t L_{ht}$, and each agency can change its nominal wage W_{ht} only with probability $1-\kappa_w$. The government collects subsidies as lump sum taxes from households and runs a balanced budget each period. The government sets the subsidy rates such that under flexible prices, the markup over marginal cost is zero in both the labor and output markets.

Capital goods producers produce new capital goods subject to standard investment adjustment costs and have profits Π^I_t. Thus, the total amount of lump-sum payments Π_t received by the household is the sum of the profits of all final good producers, labor agencies, and capital goods producers, minus the sum of all subsidies.

Finally, the exogenous stochastic processes are productivity and the monetary policy shock:

$$\log A_t = (1-\rho)\log \bar{A} + \rho \log A_{t-1} + \log \varepsilon_{At} \tag{3.15}$$

$$\varepsilon_{At} \sim \mathcal{N}\left(0, \sigma_A^2\right) \tag{3.16}$$

$$\varepsilon_{it} \sim \mathcal{N}\left(0, \sigma_i^2\right) \tag{3.17}$$

All market clearing conditions are listed in Appendix A.

[12] This assumption is standard in the New Keynesian literature. It eliminates distortions from monopolistic competition where firms price above marginal cost. The only distortion is then due to sticky prices, which simplifies the solution by perturbation methods.

3.2 Rational expectations equilibrium

I first describe the equilibrium under rational expectations. An equilibrium is a set of stochastic processes for prices and allocations, a set of strategies in the limited commitment game, and an expectation measure $\mathcal{P}$ such that the following holds for all states and time periods: Markets clear; allocations solve the optimization programs of all agents given prices and expectations $\mathcal{P}$; the strategies in the limited commitment game are a subgame-perfect Nash equilibrium for all lender-borrower pairs; and the measure $\mathcal{P}$ satisfies rational expectations.

Under appropriate parameter restrictions, there exists a rational expectations equilibrium characterized by the following properties. Proofs are relegated to Appendix B.

1. All firms choose the same capital-labor ratio K_{jt}/L_{jt}. This allows one to define an aggregate production function and an internal rate of return on capital:

$$Y_t = \alpha K_{t-1}^{\alpha}\left(A_t \tilde{L}_t\right)^{1-\alpha} \tag{3.18}$$

$$R_t^k = q_t \alpha \frac{Y_t}{K_{t-1}} + Q_t (1-\delta) K_{t-1} \tag{3.19}$$

2. The expected return on capital is higher than the internal return on debt: $\mathbb{E}_t R_{t+1}^k > R_{jt}$.

3. At any time t, the stock market valuation P_{jt} of a firm j is proportional to its net worth after entry and exit $\tilde{N}_{jt}$. This permits one to write an aggregate stock market index as

$$P_t = \int_0^1 P_{jt} = \beta \mathbb{E}_t \left[D_{t+1} + \frac{1-\gamma}{1-\gamma+\gamma\omega} P_{t+1}\right]. \tag{3.20}$$

4. Borrowers never default on the equilibrium path and borrow at the risk-free rate

$$R_{jt} = R_t = \frac{1}{\mathbb{E}_t \Lambda_{t+1}}. \tag{3.21}$$

   The lender only accepts debt payments up to the limit given by (3.13), which is proportional to the firm's net worth $\tilde{N}_{jt}$, and the firm always exhausts this limit.

5. As a consequence of the previous properties of the equilibrium, all firms can be aggregated. Aggregate debt, capital, and net worth are sufficient to describe the intermediate goods sector:

$$N_t = R_t^k K_{t-1} - R_{t-1} B_{t-1} \tag{3.22}$$

$$Q_t K_t = (1-\gamma+\gamma\omega)\left((1-\zeta) N_t + \zeta \left(B_{t-1} - Q_t K_{t-1}\right)\right) + B_t \tag{3.23}$$

$$B_t = x \mathbb{E}_t \Lambda_{t+1} Q_{t+1} \xi K_t + (1-x)\xi\left(P_t + B_t\right). \tag{3.24}$$

I solve for a second-order approximation of this rational expectations equilibrium around its non-stochastic steady state.

3.3 Learning equilibrium

I introduce learning about stock market valuations, as in the simple model of Section 2. There is now a continuum of firms to be priced in the market, and I retain the belief that the stock price of an individual firm is proportional to firm net worth, as is the case under rational expectations. As such, under $\mathcal{P}$,

$$P_{jt} = \frac{N_{jt}}{N_t} P_t. \tag{3.25}$$

But while investors know how to price individual stocks by observing the valuation of the market, they are uncertain about the evolution of the market itself. As in the simple model of the previous section, I impose the same beliefs about aggregate stock prices as in the last section along with the other assumptions (equations (2.20)–(2.23)), including expectations on other variables that are conditionally consistent with outcomes on the equilibrium path: For any variable x_t, any date t, and any sequence $P_0, \ldots, P_t$ that is on the equilibrium path, agents' beliefs coincide with the best statistical prediction of x_t *conditional* on the realization of stock prices.

In practice, I solve the model using a two-stage procedure. The first stage is to solve for the policy functions and beliefs under $\mathcal{P}$. The Kalman filtering equations that describe beliefs about stock prices are as follows:

$$\log P_t = \log P_{t-1} + \hat{\mu}_{t-1} - \frac{\sigma_\nu^2 + \sigma_\eta^2}{2} + z_t \tag{3.26}$$

$$\hat{\mu}_t = \hat{\mu}_{t-1} - \frac{\sigma_\nu^2}{2} + g z_t, \tag{3.27}$$

where $\hat{\mu}_t$ is the mean belief about the trend in stock price growth, and z_t is the *forecast error.* Under the subjective beliefs $\mathcal{P}$, it is normally distributed white noise. I impose that beliefs about any other endogenous variable are consistent with model outcomes conditional on the evolution of stock prices, and so beliefs and policy functions can be calculated much in the same way as under rational expectations, taking z_t as an exogenous shock process. The market clearing condition for stocks does not enter this first stage of the problem. Adding it would effectively require that beliefs about stock prices, too, be consistent with equilibrium outcomes—and the solution would collapse to the rational expectations equilibrium. Now, if x_t is the set of model variables and u_t the set of exogenous shocks, solving this first stage leads to a subjective policy function $x_t = h(x_{t-1}, u_t, z_t)$.

The second stage of the model consists in finding the value for z_t which leads to market clearing in the stock market and thereby establishes equilibrium. This results in a mapping from the state variables and exogenous shocks to the perceived forecast error $r : (x_{t-1}, u_t) \mapsto z_t$. Clearly, this function generally does not make z_t an iid disturbance in equilibrium. This is why agents make systematic forecast errors. The complete solution of the model is given by $x_t = g(x_{t-1}, u_t) = h(x_{t-1}, u_t, r(x_{t-1}, u_t))$. A complete description of a second-order perturbation of this solution is contained in Appendix C.

3.4 Choice of parameters

I partition the set of parameters into two groups. The first set of parameters is calibrated to first-order moments, and the second set is estimated by simulated method of moments (SMM) on second-order moments of US quarterly data.

3.4.1 Calibration

The capital share in production is set to $\alpha = 0.33$, implying a labor share in output of two thirds. The depreciation rate $\delta = 0.025$ corresponds to 10 percent annual depreciation. The persistence of the temporary component of productivity is set to 0.95.

The discount factor is set such that the steady-state interest rate matches the average annual real return on Treasury bills of 2.5 percent, implying a discount factor $\beta = 0.9938$. The elasticity of substitution between varieties of the final consumption good, as well as that among varieties of labor used in production, is set to $\sigma = \sigma_w = 4$. The Frisch elasticity of labor supply is set to 3, implying $\phi = 0.33$.

The strength of monetary policy reaction to inflation is set to $\phi_\pi = 1.5$, and the degree of nominal rate smoothing is set to $\rho_i = 0.85$.

Four parameters describe the structure of financial constraints: x, the probability of restructuring after default; ξ, the tightness of the borrowing constraint; ω, the equity received by new firms relative to average equity; and γ, the rate of firm exit and entry. I calibrate the restructuring rate to $x = 0.093$. This is the fraction of US business bankruptcy filings in 2006 that filed for Chapter 11 instead of Chapter 7, and that subsequently emerged from bankruptcy with an approved restructuring plan (a sensitivity check is included in Section 5.2).[13] The remaining three parameters are chosen such that the non-stochastic steady state of the model jointly matches the US average investment share in output of 20 percent, average debt-to-equity ratio of 1:1 (as recorded in the Fed flow of funds), and average quarterly P/D ratio of 139 (taken from the S&P500). The parameter values thus are $\gamma = 0.0165$, $\xi = 0.4152$, and $\omega = 0.018$.

3.4.2 Estimation

The remaining parameters are the standard deviations of the technology and monetary shocks (σ_A, σ_i), the degree of nominal price and wage rigidities (κ, κ_w), the size of investment adjustment costs (ψ), the fraction of dividends paid out as earnings by continuing firms (ζ), and the learning gain (g). I estimate these six parameters to minimize the distance to a set of eight moments

[13] 2006 is the only year for which this number can be constructed from publicly available data. Data on bankruptcies by chapter are available at http://www.uscourts.gov/Statistics/BankruptcyStatistics.aspx. Data on Chapter 11 outcomes are analyzed in various samples by Flynn and Crewson (2009), Warren and Westbrook (2009), Lawton (2012), and Altman (2014).

Table 1: Estimated parameters.

param.	σ_a	σ_i	κ	κ_w	ψ	ζ	g
learning	.00884	.000423	.546	.932	13.7	.632	.00563
	(.000967)	(.00195)	(.089)	(.132)	(3.85)	(.0935)	(.000334)
RE	.0114	.000895	.691	.572	.618	.0490	-
	(.00212)	(.00173)	(.168)	(2.73)	(10.7)	(11.3)	
fric.less	.0116	.00121	.671	.847	.558	-	-
	(.00716)	(.000701)	(.261)	(.398)	(.124)		

Parameters as estimated by simulated method of moments. Asymptotic standard errors in parentheses.

pertaining to both business cycle and asset price statistics: The standard deviation of output; the standard deviations of consumption, investment hours worked, and stock prices relative to output; and the standard deviations of inflation, the nominal interest rate, and stock returns (see also Table2). The set of estimated parameters θ solves

$$\min_{\vartheta \in \mathcal{A}} \left(m\left(\theta\right) - \hat{m}\right)' W \left(m\left(\theta\right) - \hat{m}\right)',$$

where $m\left(\theta\right)$ are moments obtained from model simulation paths with 50,000 periods, $\hat{m}$ are the estimated moments in the data, and W is a weighting matrix.[14] I also impose that θ has to lie in a subset $\mathcal{A}$ of the parameter space which rules out deterministic oscillations of stock prices.[15] Such oscillations are not observed in the data, but can be consistent with equilibrium when asset price volatility is high and subjective beliefs are far away from rational expectations. In a sense, this restriction therefore constrains the departure of subjective beliefs from rational stock price expectations. Table 1 summarizes the SMM estimates for both the learning and rational expectations version of the model, as well as for a comparison (rational expectations) model in which all financial frictions are eliminated. The first row presents the results under learning. Exogenous shocks come mainly from productivity shocks, since σ_i is estimated to be relatively small. The Calvo price adjustment parameter is set to $\kappa = 0.546$, implying retailers adjust their prices every two quarters. The SMM procedure selects a high degree of nominal wage rigidities κ_w and of adjustment costs ψ. The estimates are substantially larger than what is commonly found in the literature. The fraction of earnings paid out as dividends is fitted to $\zeta = 0.632$, which is in line with the historical average for the S&P500 at about 50 percent. Finally, the learning gain $g = 0.00563$ implies that agents believe the degree of predictability in the stock market to be very small.

[14] I choose $W = \text{diag}\left(\hat{\Sigma}\right)^{-1}$ where $\hat{\Sigma}$ is the covariance matrix of the data moments, estimated using a Newey-West kernel with optimal lag order. This choice of W leads to a consistent estimator that places more weight on moments which are more precisely estimated in the data.

[15] $\theta \notin \mathcal{A}$ iff there exists an impulse response of stock prices with positive peak value also having a negative value of more than 20% of the peak value.

The second row contains the parameters estimated under rational expectations. The fit to the asset price moments included in the estimation is worse and the asymptotic standard errors are large, implying that the distance of the moments to the data at the point estimate is relatively flat. Nevertheless, at this point estimate, the size of the shocks σ_a and σ_i is substantially larger under learning. This implies that learning about stock prices leads to substantial amplification of shocks: The increased endogenous volatility of asset prices greatly magnifies the financial accelerator effect, just as in the simple model of Section 2. The degree of wage rigidities and investment adjustment costs required to fit the data is smaller than under learning.

The third row contains parameter estimates under the model without financial frictions (and rational expectations). Since the financial structure is eliminated from that model, the dividend payout ratio ζ is not present. The size of the shocks is larger than under the rational expectations model with the financial accelerator present. This points to moderate amplification effects of financial frictions under rational expectations.

4 Results

I now present the quantitative results of this paper. First, I review standard business cycle statistics. Learning and asset price volatility account for more than a third of the volatility of output, pointing to the strength of the endogenous amplification mechanism. In contrast, the financial accelerator mechanism under rational expectations is relatively weak. I then look at asset pricing moments and find that the model with learning closely matches not only the volatility of stock prices (which is targeted by the estimation), but also the predictability of stock returns as well as negative skewness and excess kurtosis. Under rational expectations, all these statistics are close to zero even though the estimation tries to target asset price volatility. Next, I present impulse response functions to both supply and demand shocks, confirming the strong amplification mechanism in all main macroeconomic aggregates. The main element is the endogenous volatility of asset prices induced by learning, leading to very procyclical credit constraints. But I also show that this is not the whole story: The fact that agents are too optimistic about future asset prices during expansions and vice-versa affects aggregate demand, adding additional amplification that is unique to a model with learning. Finally, I compare forecast errors made by agents in the model with those observed in actual survey data. The patterns of predictability are remarkably similar, lending credibility to the chosen expectation formation process.

4.1 Business cycle and asset price moments

To get a better understanding of the quantitative properties of the model, I review key moments in the data and across model specifications. Table 2 starts with business cycle statistics. The data moments are in Column (1). Moments for the estimated learning model are in Column

Table 2: Business cycle statistics in the data and across model specifications.

	moment	(1) data	(2) learning	(3) RE	(4) fric.less	(5) RE re-estimated
output volatility	$\sigma_{hp}(Y_t)$	1.43%* (0.14%)	1.56%	1.00	.79	1.51%
volatility rel. to output	$\sigma_{hp}(C_t)/\sigma_{hp}(Y_t)$	.60* (.035)	.60	.94	1.34	.58
	$\sigma_{hp}(I_t)/\sigma_{hp}(Y_t)$	2.90* (.12)	2.78	.48	.31	2.79
	$\sigma_{hp}(L_t)/\sigma_{hp}(Y_t)$	1.13* (.061)	1.18	.84	.40	1.10
correlation with output	$\sigma_{hp}(C_t, Y_t)$	.94 (.0087)	.59	.86	1.00	.84
	$\sigma_{hp}(I_t, Y_t)$	.95 (.0087)	.87	.89	.25	.90
	$\sigma_{hp}(L_t, Y_t)$	.85 (.035)	.88	.65	.24	.76
inflation	$\sigma_{hp}(\pi_t)$	.27%* (.047%)	.34%	.28%	.28%	.25%
nominal rate	$\sigma_{hp}(i_t)$	.37%* (.046%)	.11%	.11%	.12%	.07%

Quarterly U.S. data 1962Q1–2012Q4. Standard errors in parentheses. π_t is quarterly CPI inflation. i_t is the federal funds rate. L_t is total non-farm payroll employment. Consumption C_t consists of services and non-durable private consumption. Investment I_t consists of private non-residential fixed investment and durable consumption. Output Y_t is the sum of consumption and investment. $\sigma_{hp}(\cdot)$ is the standard deviation and $\rho_{hp}(\cdot,\cdot)$ is the correlation coefficient of HP-filtered data (smoothing coefficient 1600). Moments used in the SMM estimation are marked with an asterisk.

(2), while Columns (3) and (4) contain the corresponding moments for the model under rational expectations and the frictionless benchmark, respectively. Here, the parameters are held constant at the estimated values as for the learning model. By nature of the estimation, the learning model has the best fit across Columns (2) to (4). The comparison serves to single out the contribution of learning and financial frictions to the fit. Column (5) presents the moments under rational expectations when the parameters are re-estimated to fit the data.

The first row reports the standard deviation of de-trended output. By construction, this is matched well by the learning model in Column (2). When learning is shut off in Column (2), the standard deviation drops one-third. This shows the great degree of amplification that learning adds to the model. Of course, it is possible to match output volatility with rational expectations, using larger shock sizes, as in Column (5). But the comparison between Columns (2) and (3) singles out the contribution of learning to the internal amplification mechanism. The standard (rational expectations) financial accelerator mechanism is present in the model as well, since the volatility of output drops further in Column (4) when financial frictions are shut off.

The next three rows report the standard deviation of consumption, investment, and hours worked

Table 3: Asset price statistics in the data and across model specifications.

	moment	(1) data	(2) learning	(3) RE	(4) fric.less	(5) RE re-estimated
excess volatility	$\sigma_{hp}\left(P_t\right)/\sigma_{hp}\left(Y_t\right)$	7.86*	8.96	.26	-	.16
		(.61)				
	$\sigma\left(\frac{P_t}{D_t}\right)$	41.08%	22.62%	4.12%	-	3.59%
		(6.11%)				
	$\sigma\left(R^e_{t,t+1}\right)$	8.14%*	7.12%	.19%	-	.19%
		(.61%)				
return predictability	$\rho\left(\frac{P_t}{D_t}, R^e_{t,t+4}\right)$	-.297	-.376	-.040	-	-.035
		(.092)				
	$\rho\left(\frac{P_t}{D_t}, R^e_{t,t+20}\right)$	-.585	-.732	-.006	-	0.011
		(.132)				
	$\rho\left(\frac{P_t}{D_t}, \frac{P_{t+4}}{D_{t+4}}\right)$	.904	.637	.303	-	.564
		(.056)				
negative skewness	$\text{skew}\left(R^e_{t,t+1}\right)$	-.897	-.404	.022	-	.005
		(.154)				
heavy tails	$\text{kurt}\left(R^e_{t,t+1}\right)$	1.57	.92	.04	-	-.03
		(.62)				

Quarterly U.S. data 1962Q1–2012Q4. Standard errors in parentheses. Dividends D_t are four-quarter moving averages of S&P 500 dividends. The stock price index P_t is the S&P 500. $\sigma(\cdot)$ is the standard deviation; $\sigma_{hp}(\cdot)$ is the standard deviation of HP-filtered data (smoothing coefficient 1600); $\rho(\cdot,\cdot)$ is the correlation coefficient; skew $(\cdot)$ is skewness;kurt $(\cdot)$ is excess kurtosis. Moments used in the SMM estimation are marked with an asterisk.

relative to output. Moving from Column (2) to (3), it can be seen that the removal of learning leads to a sharp drop in the relative volatility of both investment and hours worked. This is because the estimated learning model features a high level of investment adjustment costs to match investment volatility. Without large asset price fluctuations generated by learning, investment becomes too smooth, as does the marginal product of capital and hence labor demand. The next rows report the volatility of inflation and the nominal interest rate. Inflation volatility is roughly in line with the data, but the nominal interest rate is less volatile across all model specifications. This might be due to the fact that the data sample includes the volatile '70s and the following Volcker disinflation period.

Next, I present asset price statistics in Table 3. The statistics correspond to some well-known asset price puzzles. The learning model fits them remarkably well, despite being solved only with a second-order perturbation approach. Starting with excess volatility in Column (2), the model with learning produces standard deviations of prices, P/D ratio, and returns that are close to the data. By contrast, the model with rational expectations in Column (5) cannot produce a similar amount ov volatility, despite the fact that price and return volatility are explicitly targeted by the SMM estimation.

Stock returns also exhibit considerable predictability by the P/D ratio at business-cycle frequency.

Figure 4: Impulse responses to a persistent productivity shock.

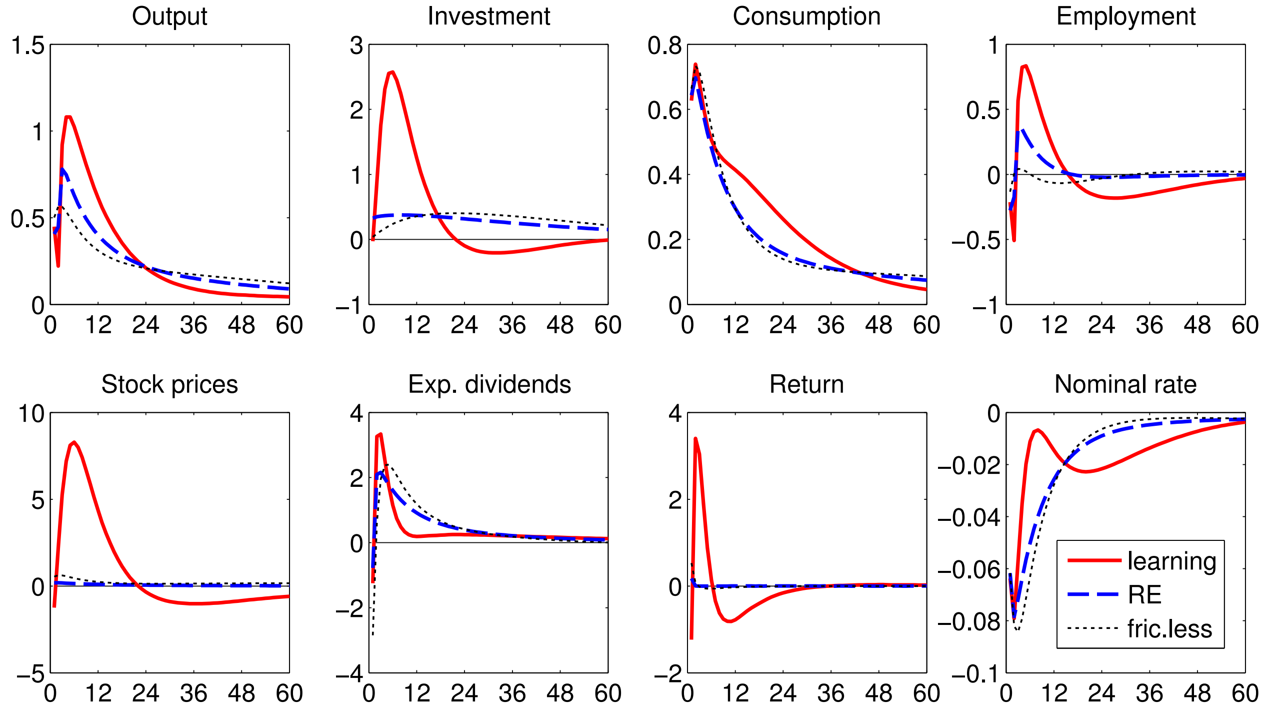

Impulse responses to a one-standard deviation innovation in ε_{At}, averaged over 5,000 random shock paths with a burn-in of 1,000 periods. Stock prices, dividends, output, investment, consumption, and employment are in 100*log deviations. Stock returns and the nominal interest rate are in percentage point deviations.

The same is true in the model with learning. Predictability is not targeted by the SMM procedure, and in fact it is somewhat stronger than in the data, reflected in a persistence of the P/D ratio somewhat lower than in the data. But again, the rational expectations model is not able to produce return predictability.

Finally, the learning model also produces a distribution of returns that is negatively skewed and heavy-tailed to a similar degree as in the data. This points to the importance of non-linearities in the asset price dynamics under learning.

4.2 Impulse response functions

Impulse response functions reveal some of the workings of the amplification mechanism at play. Figure 4 plots the impulse responses to a persistent productivity shock. Red solid lines represent the learning equilibrium, blue dashed lines represent the rational expectations version, and black thin lines represent the comparison model without financial frictions. The impulse responses are averaged across states and therefore mask the tail dynamics present under learning, but they are nevertheless instructive. Looking at the first row of impulse responses, output rises persistently after the shock due to both the increased productivity and the relaxation of credit constraints from higher asset prices. The increase in output is larger under rational expectations than under

Figure 5: Impulse responses to a monetary shock.

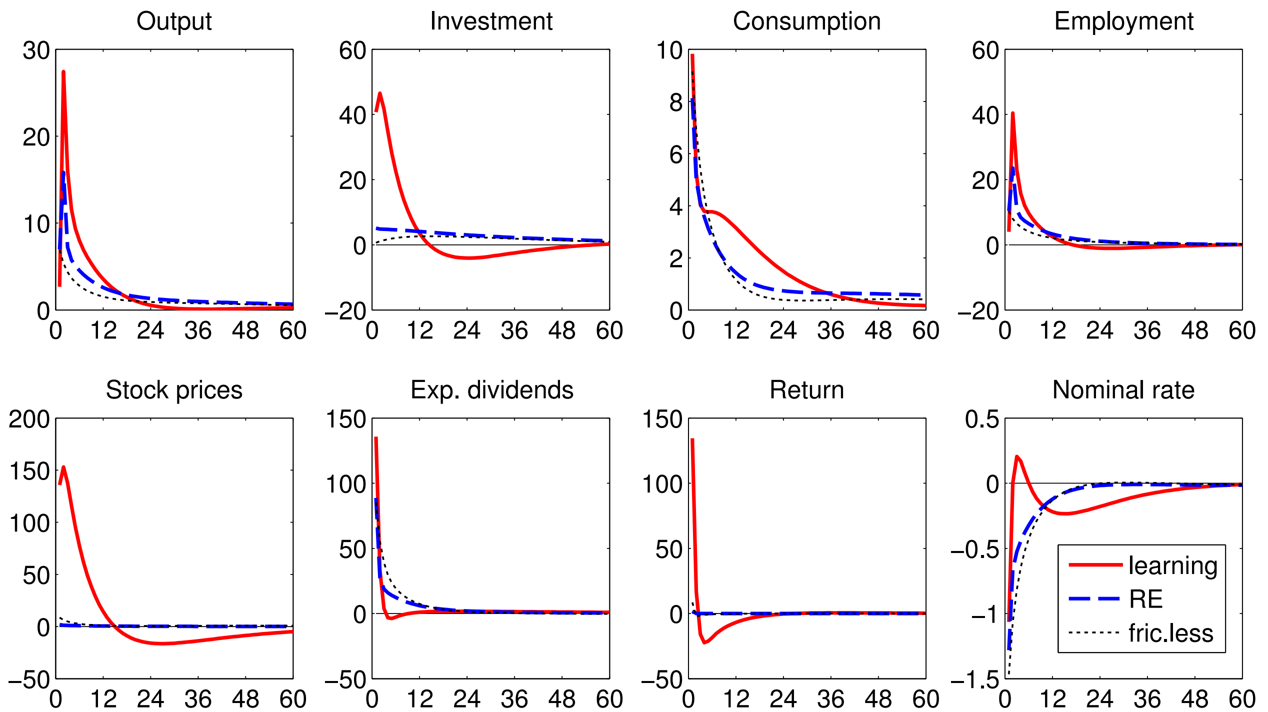

Impulse responses to a innovation in ε_{mt}, averaged over 5,000 random shock paths with a burn-in of 1,000 periods. The size of the innovation is chosen to produce a 10 basis point fall in the equilibrium nominal rate. Stock prices, dividends, output, investment, consumption and employment are in 100*log deviations. Stock returns and the nominal interest rate are in percentage point deviations.

the frictionless comparison; this is the standard financial accelerator effect. When learning is introduced, the response to the shock is amplified further. This also translates into amplification of the responses of investment, consumption, and employment. The amplification is due to two channels: First, learning leads to higher stock prices. The increase in firms' market value allows them to borrow more and invest and produce more. Second, agents under learning are not aware of the mean reversion in stock prices and predict the stock price boom to last for a long time. Consequently, they overestimate the availability of credit and therefore production in the future, leading to an aggregate demand effect that increases output today (see also 4.3). The rise in stock prices in the second row of Figure 4 is large under learning and accompanied by an initial spike in dividend payments, although dividends subsequently fall below their counterpart under rational expectations. The nominal interest rate falls less under learning as the central bank reacts to the inflationary pressures stemming from the relaxation in credit constraints.

Figure 5 plots the response to a temporary reduction in the nominal interest rate. Again, all macroeconomic aggregates rise substantially more under learning than under both rational expectations and the frictionless benchmark. The monetary stimulus increases stock prices and thus relaxes credit constraints. The consequent increase in aggregate demand raises inflationary pressure, so that the systematic reaction of the interest rate rule raises the interest rate sharply again after the shock.

4.3 Does learning matter?

The discussion so far has mainly focused on how large swings in asset prices lead to large swings in real activity through their effect on credit constraints. But is learning necessary for this story at all? Maybe all that matters for amplification is that asset price volatility has to be increased, by some mechanism or other. In this section, I show that learning has an effect on amplification over and above its effect on asset prices.

I replace the stock market value P_t in the borrowing constraint (3.24) with an exogenous process V_t that has the same law of motion as the stock price under learning. More precisely, I fit an ARMA(10,5) process for V_t such that its impulse responses are as close as possible to those of P_t under learning (the exogenous shock in the ARMA process are the productivity and monetary shocks). I then solve this model, but with rational expectations. If learning only matters because it affects stock price dynamics, then this hypothetical model should have exactly identical dynamics to the model under learning.[16]

Figure 6 shows that this is not the case. The ARMA process fits stock prices well: The impulse response of P_t under learning and V_t in the counterfactual experiment are indistinguishable. But after a positive productivity shock, output, investment, and consumption rise more under learning, even though the counterfactual model has the same stock price dynamics by construction. The reason is that expectations of future asset prices matter beyond their direct impact on current prices. Under learning, agents do not fully internalize mean reversion in stock prices and therefore predict that credit constraints are loose for longer than they turn out to be. This leads to a wealth effect on households that increases their consumption, raising aggregate demand, and it leads to higher future expected prices of capital goods $\mathbb{E}_t Q_{t+1}$, which enters the liquidation value of firms and hence relaxes borrowing constraints, even if stock prices are the same as under rational expectations. These effects are powerful enough to create significant endogenous amplification through the departure of subjective beliefs from rational expectations.

4.4 Relation to survey evidence on expectations

The rational expectations hypothesis asserts that "outcomes do not differ systematically [...] from what people expect them to be" (Sargent, 2008). Put differently, a forecast error should not be systematically predictable by information available at the time of the forecast. The absence of predictability is almost always rejected in the data.

Similarly, agents in the model under learning also make systematic, predictable forecast errors. This holds not only for stock prices but also other endogenous model variables, despite the fact

[16] For this exercise I only compute a first-order approximation to the model equations.

Figure 6: Does learning matter?

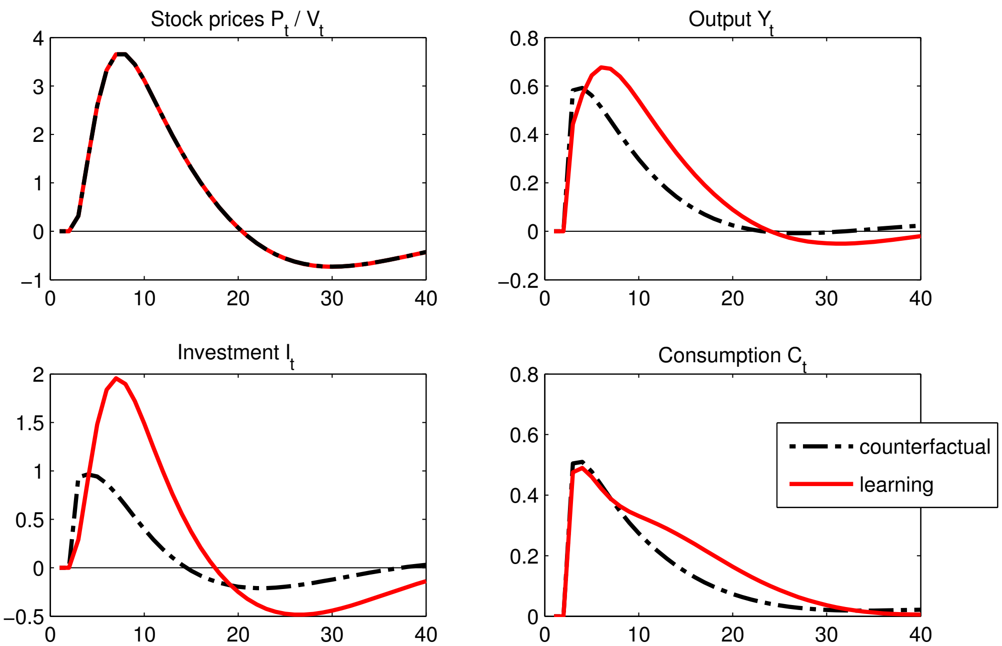

Solid red line: Impulse response to a one standard deviation positive productivity shock under learning. Black dash dotted line: Impulse response to a hypothetical rational expectations model with stock price dynamics identical to those under learning (see text). The impulse responses in the figure are produced using a first-order approximation to the model equations.

that, *conditional* on stock prices, agents' beliefs are model-consistent. A systematic mistake in predicting stock prices will still spill over into a corresponding mistake in predicting the tightness of credit constraints, and hence investment, output, and so forth. Owing to the internal consistency of beliefs, I can compute well-defined forecast errors made by agents in the model at any horizon and for any model variable.

Figure 7 repeats the scatter plot at the beginning of the paper, contrasting expected and realized one year-ahead returns in a model simulation. The same pattern as in the data emerges: When the P/D ratio is high, return expectations are most optimistic. In the learning model, this has a causal interpretation: High return expectations drive up stock prices. At the same time, realized future returns are, on average, low when the P/D ratio is high. This is because the P/D ratio is mean-reverting (which agents do not realize, instead extrapolating past price growth into the future): At the peak of investor optimism, realized price growth is already reversing and expectations are due to be revised downward, pushing down prices toward their long-run mean.

Table 4 describes tests using the Federal Reserve's Survey of Professional Forecasters (SPF) as well as the CFO survey data and compares the statistics to those obtained from simulated model data. Each entry corresponds to a correlation of the error of the mean survey forecast with a variable that is observable by respondents at the time of the survey. Under the null of rational expectations,

Figure 7: Return expectations and expected returns in a model simulation.

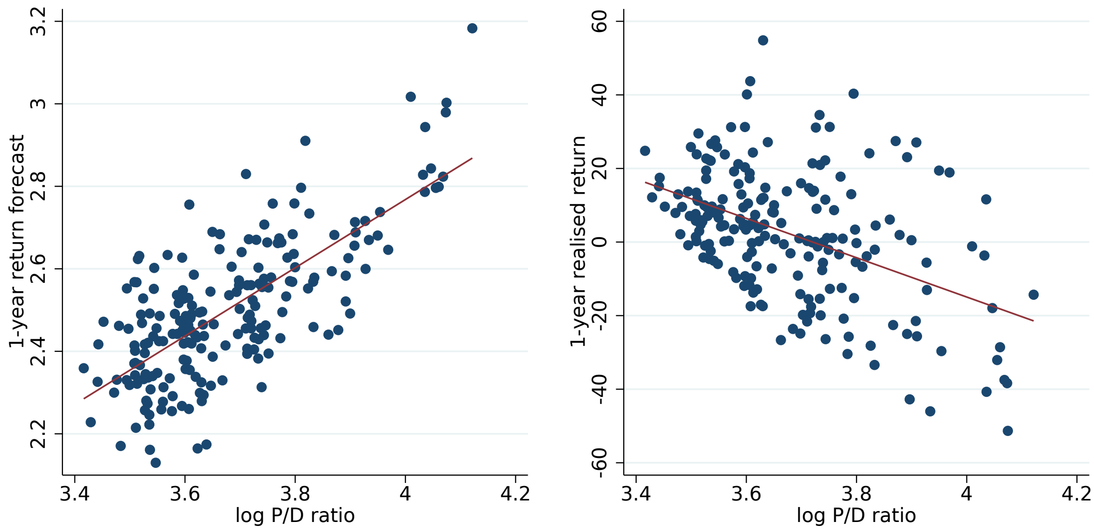

Expected and realized nominal returns along a simulated path of model with learning. Simulation length 200 periods. Theoretical correlation coefficient for subjective expected returns $\rho = .47$, for future realized returns $\rho = -.38$.

all entries should be zero.

Column (1) shows that the P/D ratio negatively predicts forecast errors. When stock prices are high, people systematically under-predict economic outcomes. This holds in particular for stock returns, as was already shown in the scatter plot above. But it also holds true for macroeconomic aggregates, albeit at lower levels of significance. The same holds true in Column (2), which shows the correlation coefficients obtained from simulated model data.

Column (3) repeats the exercise for the growth rate of the P/D ratio. This measure positively predicts forecast errors, suggesting that agents' expectations are too cautious and under-predict an expansion in its beginning but then overshoot and over-predict it when it is about to end. In the model (Column 4), this pattern also emerges because expectations about asset prices (and hence lending conditions) adjust only slowly. The similarity of the correlations in the data and in the model is striking, with the exception of aggregate consumption. The reason is that consumption forecasts in the model are only biased at longer horizons: A relaxation of borrowing constraints first leads to an increase in investment and only later to an increase in consumption. Agents are aware of this relationship, so that their three-quarter forecasts, as in Table 4 do not become much more optimistic when the P/D ratio increases. At longer forecast horizons, one would observe more predictability for consumption as well.

Column (5) reports the results of a particular test of rational expectations devised by Coibion and Gorodnichenko (2015). Since for any variable x_t, the SPF asks for forecasts at one- through four-quarter horizons, it is possible to construct a measure of agents' revision of the change in x_t as $\hat{\mathbb{E}}_t\left[x_{t+3} - x_t\right] - \hat{\mathbb{E}}_{t-1}\left[x_{t+3} - x_t\right]$. Forecast errors are positively predicted by this revision

Table 4: Forecast errors under learning and in the data.

	(1)	(2)	(3)	(4)	(5)	(6)
	$\log PD_t$		$\Delta \log PD_t$		forecast revision	
forecast variable	data	model	data	model	data	model
$R^{stock}_{t,t+4}$	-.44	-.38	.06	.30	-	-.30
	(-3.42)		(.41)			
$Y_{t,t+3}$	-.21	-.16	.22	.16	.29	.28
	(-1.78)		(2.42)		(3.83)	
$I_{t,t+3}$	-.20	-.37	.25	.27	.31	.35
	(-1.74)		(2.88)		(3.79)	
$C_{t,t+3}$	-.19	-.04	.21	.01	.23	.02
	(-1.85)		(2.37)		(2.67)	
$u_{t,t+3}$	.05	.20	-.27	-.20	.43	.32
	(.12)		(-3.07)		(6.07)	

Correlation coefficients for mean forecast errors on one year-ahead nominal stock returns (Graham-Harvey survey) and three quarters-ahead real output growth, investment growth, consumption growth and the unemployment rate (SPF). t-statistics in parentheses. Regressors: Column (1) is the S&P 500 P/D ratio and Column (2) is its first difference. Column (3) is the forecast revision, as in Coibion and Gorodnichenko (2015). Data from Graham-Harvey covers 2000Q3–2012Q4. Data for the SPF covers 1981Q1–2012Q4. For the model, correlations are computed using a simulation of length 50,000, where subjective forecasts are computed using a second-order approximation to the subjective belief system on a path in which no more future shocks occur, starting at the current state in each period. Unemployment in the model is taken to be $u_t = 1 - L_t$.

measure. Coibion and Gorodnichenko take this as evidence for sticky information models in which information sets are gradually updated over time. But it is also consistent with the learning model: The correlation coefficients in Column (6) are very similar to those in the data.[17]

The spill-over of expectational errors is illustrated graphically in Figure 8. The solid red line is a standard impulse response function to a technology shock in the learning model. The green dashed line of Figure 4 depicts agents' average subjective forecast at the peak of the stock market boom following the shock. Agents do not foresee the decline in the stock market and instead extrapolate high stock price growth into the future. Because stock market valuations matter for access to credit, agents also forecast loose borrowing conditions and are too optimistic about output, investment, and (to a lesser extent) consumption, as in the data. The green dashed line (expectations) is above the red solid line (realization) when the P/D ratio is high: It negatively predicts forecast errors. The blue dotted line is the forecast made at a time in which the P/D ratio is rising. In this situation, agents under-predict the size of the coming boom in the stock market and real activity. Therefore, the growth rate of the P/D ratio positively predicts forecast errors.

[17] The model predicts a negative correlation of forecast errors on stock returns with their forecast revisions. The CFO survey does not allow for the construction of the corresponding statistic in the data, but it is an interesting implication since a negative correlation cannot be produced by rigid information models as in Coibion and Gorodnichenko.

Figure 8: Actual versus expected impulse response.

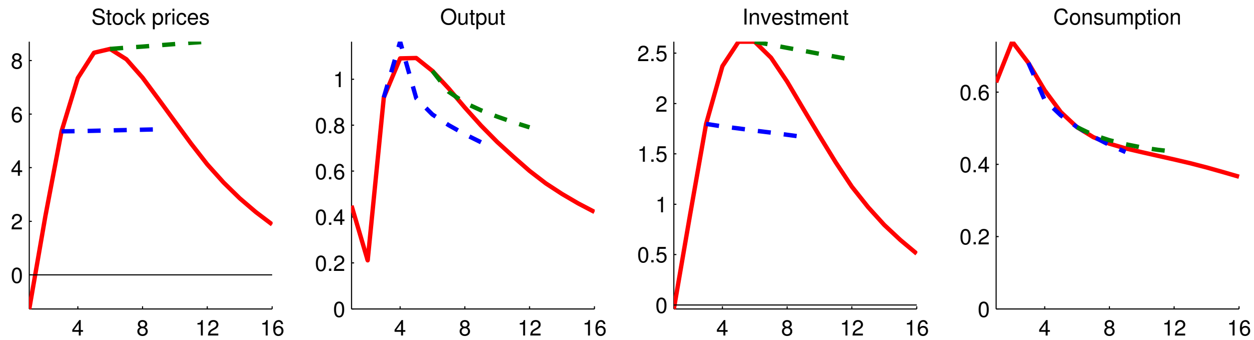

Solid red line: Impulse response to a one-standard deviation positive productivity shock, averaged over 5,000 random shock paths with a burn-in of 1,000 periods. Green dashed line: mean subjective forecasts taken in period 6. Blue dashed line: mean subjective forecasts taken in period 3. Subjective forecasts are computed using a second-order approximation to the subjective policy function on a path in which no more future shocks occur.

5 Sensitivity checks

5.1 Nominal rigidities

The quantitative model includes price- and wage-setting frictions that complicate the model dynamics. They are nevertheless important for the quantitative fit of the model, as I will argue here. Recall that in the simple model of Section 2, the amplifying effect of asset price learning depended crucially on the behavior of the real wage (affecting the ϵ_D term in Equation **??**). As credit constraints relax and investment picks up, wages rise which work to diminish firms' profits and expected dividend payments. This drives down stock prices and dampens the learning dynamics. The same mechanism is at play in the quantitative model. Introducing nominal rigidities greatly helps to obtain amplification.

Specifically, expected dividends in this model are given by:

$$\begin{aligned}\mathbb{E}_t^{\mathcal{P}} D_{t+1} &= (\gamma+(1-\gamma)\zeta)\left(\alpha q_t^{\frac{1}{\alpha}}\mathbb{E}_t^{\mathcal{P}}\left((1-\alpha)\frac{A_{t+1}}{w_{t+1}}\right)^{\frac{1-\alpha}{\alpha}}+(1-\delta)\,\mathbb{E}_t^{\mathcal{P}}Q_{t+1}-R_t\frac{B_t}{K_t}\right)K_t\\ &- (1-\gamma)\,\zeta\left(1-\frac{B_t}{K_t}\right)K_t. \qquad (5.1)\end{aligned}$$

There are four relative prices that enter this equation: The price of intermediates q_t, the real wage w_{t+1}, the price of capital goods Q_{t+1}, and the borrowing rate R_t. Suppose now that asset prices rise because of optimistic investor beliefs, relaxing credit constraints. This directly leads to an increase in the capital stock K_t and also allows for higher leverage B_t/K_t, and the rise in investment is also expected to persist in the future. The next period's price of capital goods Q_{t+1} increases, and this raises expected dividends, helping amplification through higher asset prices. Likewise, increased aggregate demand can raise q_t if prices are sticky. But the increased labor demand, together with a positive wealth effect that households expect from the relaxation of credit constraints, will drive up the real wage w_t. This will tend to reduce D_{t+1} and dampen the expansion. Also, to the extent that higher investment comes at the expense of lower consumption in the economy, real borrowing rates R_t will rise, also dampening the expansion. This latter effect is stronger the more leverage there is in the economy.

Nominal rigidities have effects on real wages, the price of intermediates, and real rates. Wage rigidities will counteract the dampening effects of real wage responses to shocks, allowing for greater dividend, and therefore asset price, volatility. They also lead to amplification in the response of employment to movements in financial market sentiment. Price rigidities, together with a relatively loose monetary policy rule, imply that the prices of intermediates q_t are pro-cyclical, and lead to smaller real interest rates movements in response to changes in investor sentiment, also helping amplification. The mirror image of this result, through the consumption Euler equation, is that consumption is not pushed down as much by increases in investment, so its response, too, is amplified. In sum, nominal price and wage rigidities allow for comovement of all macroeconomic aggregates in response to changes in subjective beliefs. This co-movement property obtains more generally and has been documented in the context of news shocks (Kobayashi and Nutahara, 2010) and financial shocks Ajello (2016).

To illustrate this point, I re-compute impulse responses of the model with learning, but without nominal rigidities—i.e., setting $\kappa = \kappa_w = 0$. I also reduce the size of investment adjustment costs to $\psi = 0.125$. With the high degree of adjustment costs in the baseline version, the model would include an explosive two-period oscillation and reducing ψ ensures stability. The low costs of adjusting investment also gives the real version a better chance at delivering strong impulse responses. Even then, the nominal version delivers far greater amplification. Figure 9 plots impulse responses to a positive productivity shock for the baseline and the real versions of the model. Owing to lower adjustment costs, the initial response of investment is expectedly stronger in the real version. However, the real wage w_t rises by much more, and also the price of intermediates q_t is fixed, so dividends do not rise as much after the shock. This considerably dampens the learning dynamics and mutes the response of stock prices. By consequence, the response of output, investment, consumption and hours worked (not shown) is overall weaker than in the baseline version of the model.

5.2 Borrowing constraint parameters

I will now briefly discuss the sensitivity of the results to the two main parameters affecting the borrowing constraint (3.24): The probability x that a firm can be sold as a going concern after filing for bankruptcy, governing the dependency of the constraint on stock prices; and the fraction of assets ξ preserved in bankruptcy, governing the overall tightness of the constraint. Figure 10 plots the standard deviation of output and stock prices as a function of these two parameters, respectively.

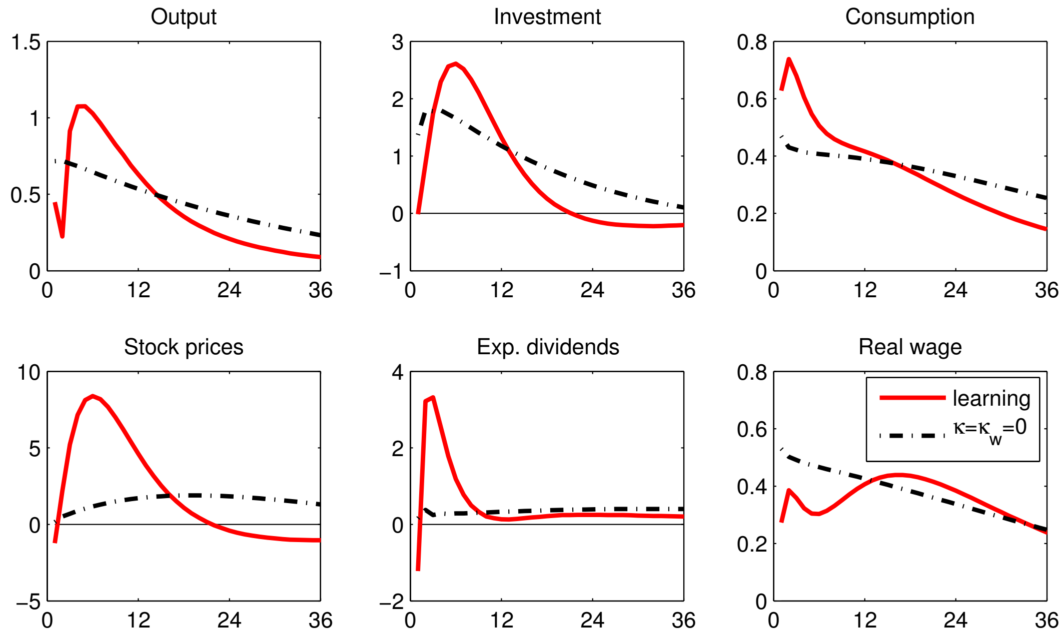

Figure 9: Role of nominal rigidities.

Solid red line: Impulse response to a one-standard deviation positive productivity shock for the model with learning and price and wage rigidities ("nominal" baseline). Black dash-dotted line: Impulse response to a productivity shock for the model with learning but without nominal rigidities, re-estimated as in Section 3.4.2 to fit the data ("real" comparison). The size of the shock shown is the same as in the nominal model.

Panel (a) shows the role of the x parameter, which crucially affects the amplification mechanism. The value $x = 0$ is a special case. At this point, stock prices do not enter the borrowing constraint and serve no role for allocations in the economy. Allocations under learning and rational expectations coincide perfectly, even though stock price dynamics are still amplified under learning. As x increases, the higher volatility of stock prices under learning translates into higher volatility in real activity as well. Since swings in real activity feed back into asset prices through their effect on dividends, the amplification becomes very strong for high values of x until the dynamics become explosive. Beyond a value of x of about 0.26, no stable learning equilibrium exists. By contrast, the rational expectations equilibrium barely depends on the parameter x. This might be a reason why the distinction between market and liquidation value has not featured prominently in the existing literature on firm credit frictions.

Panel (b) shows the role of the ξ parameter. Amplification is hump-shaped with respect to ξ. At $\xi = 0$, no collateral is pledgeable and firms have to finance their capital stock entirely out of equity. In this case, fluctuations in stock prices again do not impact the economy and allocations coincide under learning and rational expectations; there is no amplification from learning. However, as pledgeability increases to its maximum value (beyond which a steady state with permanently binding borrowing constraint does not exist), amplification also disappears. This mirrors the analysis of the simplified model.

Figure 10: Sensitivity to borrowing constraint parameters.

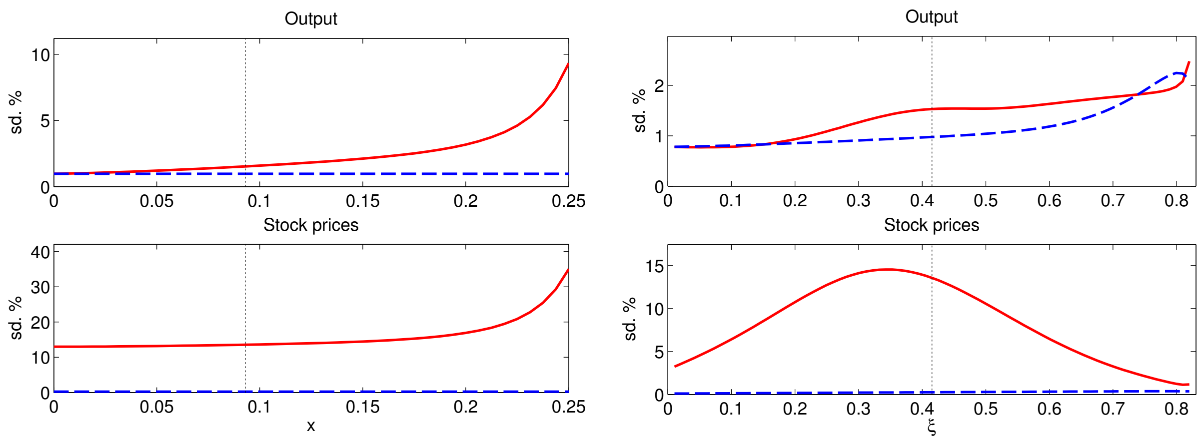

(a) Market value dependency x. (b) Recovery rate ξ.

Simulated model data HP-filtered with smoothing parameter 1600, sample length 50,000 periods. The dashed black lines indicate parameter value in the estimated learning model.

6 Implications for monetary policy

Changes in subjective expectations in financial markets can lead to large and inefficient asset price and business cycle fluctuations. A natural question is then whether policy should step in to stabilize asset prices. In the context of monetary policy, the case of "leaning against the wind"—raising rates when asset price "bubbles" are forming—has been floating in policy circles for more than 15 years, but it is usually not supported in formal models (e.g. Bernanke and Gertler, 2001). To my knowledge, most models employed to study the merits of letting interest rates react to asset prices are based on rational expectations.[18] Here, I will revisit the question for the model with learning developed in this paper.

The model does not permit to solve analytically for optimal monetary policy. But it is possible to numerically evaluate interest rate rules augmented with a reaction to asset prices. Consider extending the interest rate rule (3.14) as follows:

$$i_t = \rho_i i_{t-1} + (1-\rho_i)\begin{pmatrix} 1/\beta + \phi_\pi \pi_t \\ +\phi_{\Delta Y}\left(\log Y_t - \log Y_{t-1}\right) + \phi_{\Delta P}\left(\log P_t - \log P_{t-1}\right) \end{pmatrix} \tag{6.1}$$

In addition to raising interest rates when inflation is above its target level (taken to be zero), the central bank can raise interest rates $\phi_{\Delta Y}$ percentage points when real GDP growth increases 1 percentage point and $\phi_{\Delta P}$ percentage points when stock price growth increases 1 percentage point. I deliberately exclude levels of output or asset prices or output gap measures. Doing so would imply that the central bank has more knowledge than the private sector, as the perceived equilibrium level of asset prices and the output gap depends on agents' subjective beliefs. From

[18] A notable exception is Caputo, Medina, and Soto (2010).

Table 5: Optimal policy rules.

	learning			RE (re-estimated)		
	(1)	(2)	(3)	(4)	(5)	(6)
	baseline	w/o ΔP	w/ ΔP	baseline	w/o ΔP	w/ ΔP
ϕ_π	1.50	1.37	1.20	1.50	10.00	9.30
$\phi_{\Delta Y}$		.61	.95		2.67	2.35
$\phi_{\Delta P}$			.12			.47
$\sigma(Y)$	3.27%	2.77%	2.04%	3.79%	3.41%	3.41%
$\sigma(P)$	22.8%	15.6%	9.35%	.96%	.93%	.93%
$\sigma(\pi)$	.35%	.38%	.35%	.30%	.15%	.15%
$\sigma(i)$	.17%	.17%	.12%	.18%	.08%	.08%
welfare cost χ	**0.0346%**	**.0275%**	**.0218%**	**.0537%**	**.0396%**	**.0396%**

Standard deviations of output, stock prices, inflation, and interest rates (unfiltered) in percent. The interest rate smoothing coefficient is kept at $\rho_i = 0.85$ for all rules considered.

a practical perspective as well, a target for the level of asset prices is certainly a more audacious policy objective than a target for price growth.

I compute the parameters $(\phi_\pi, \phi_{\Delta Y}, \phi_{\Delta P})$ that minimize the welfare cost of business cycles. This cost χ is defined as

$$u\left((1-\chi)\,\mathbb{E}\left[\tilde{C}_t\right], \mathbb{E}\left[\tilde{L}_t\right]\right) = \mathbb{E}\left[u\left(C_t, L_t\right)\right].$$

That is, χ is the fraction of steady-state consumption the household would need to give up in order to have its period utility at the same level as the average stochastic period utility, in a steady state in which consumption and labor are constant and equal to their average stochastic level, and price and wage dispersion is nil (see Appendix A for the precise definition of $\tilde{C}_t$ and $\tilde{L}_t$).[19] The cost is calculated using averages over time $\mathbb{E}[\cdot]$ rather than subjective expectations $\mathbb{E}^{\mathcal{P}}[\cdot]$. Under learning, this criterion is therefore paternalistic because it does not coincide with the policy that would maximize agents' subjective welfare. In this sense, I assume that the central bank can free itself from systematic forecast errors and instead commit to a rule that minimizes the average cost of business cycles over time, as opposed to the private sector's subjectively preferred policy.

Table 5 summarizes the key findings. Column (1) reports the baseline model under learning. The bottom row shows the welfare cost. In the baseline this is $\chi = 0.0346$ percent. Column (2) then calculates the rule that minimizes χ when interest rates do not systematically react to asset prices ($\phi_{\Delta P} = 0$). This optimal rule reduces the volatility of output and asset prices and achieves a reduction of welfare costs of 21 percent, to $\chi = 0.0275$ percent. Column (3) then allows for a reaction to asset prices as well. This time, the reduction in welfare costs is 37 percent to $\chi = 0.0218$ percent. This larger reduction comes with a further substantial reduction in the volatility of output and asset prices that was not achievable previously. Under this rule, the central

[19] I simulate a model time series for $(C_t, L_t, \pi_t, \pi_t^w)$ of 10,000 periods using the second-order approximation method above, and I compute series for $\tilde{C}_t$ and $\tilde{L}_t$ using the exact formulae given in Appendix A. I then evaluate period utility using its exact formula as well to calculate the welfare loss.

bank raises its interest rate target 12 basis points when annualised quarterly stock price growth increases 1 percentage point. With the interest rate smoothing parameter $\rho_i = 0.85$ this implies an immediate increase of 1.8 basis points. These numbers are not as small as it might look given the high volatility of stock prices. When (non-annualized) quarterly growth of the S&P500 dropped from -10 percent in 2008Q3 to -25.7 percent 2008Q4 at the height of the financial crisis, the rule in Column (3) would have called for an immediate reduction in interest rates of 113 basis points from the stock market movement alone.

One critique of "leaning against the wind" policies is that they would lead to excessively volatile interest rate policies. Here, the opposite is true. Despite the reaction of interest rates to stock prices in Column (3), the volatility of interest rates drops from 0.17 percent to 0.12 percent. The reason is that the asset price reaction greatly stabilizes equilibrium asset prices. Since investors' subjective expectations are formed on the basis of observations of past price growth, raising interest rates when price growth is high stabilizes these expectations at the right moment and reduces volatility.

While a reaction to asset prices is desirable under learning, the same is not true under rational expectations. This is shown on the right side of Table 5 which evaluates policies in the model with rational expectations, re-estimated to fit the data as discussed above. At the baseline policy in Column (4), the welfare cost is $\chi = 0.0537$ percent.[20] This cost can be reduced by 26 percent to $\chi = 0.0396$ percent in Column (5) under the optimal rule without a reaction to asset prices. The optimal rule reacts much more strongly to inflation than under learning and manages to considerably reduce inflation volatility while having little effect on the volatility of output and asset prices. When I calculate the optimal rule allowing for a non-zero $\phi_{\Delta P}$ in Column (6), I do find a positive coefficient. But the allocations obtained under this rule are almost identical to those in Column (5) and the welfare cost is virtually unchanged. In other words, the objective function is flat at the optimum and a rule that leans against the wind is as good as one that doesn't. This result mirrors the findings of the existing literature. As it turns out, it crucially depends on the assumption of rational expectations.

7 Conclusion

In this paper, I analyzed the implications of a learning-based asset pricing theory in a business cycle model with financial frictions. When firms borrow against the market value of their assets, learning in the stock market interacts with credit frictions to form a two-sided feedback loop between stock prices and firm profits that amplifies the learning dynamics encountered in Adam, Beutel, and Marcet (2014). At the same time, it leads to a large amount of propagation and amplification of both supply and demand shocks, while matching standard business cycle statistics and a number of asset price properties such as the volatility, predictability, skewness and excess kurtosis of stock returns.

[20] Since the model is re-estimated and does not match the data moments in the same way, the baseline welfare cost under learning and rational expectations cannot be easily compared.

An important innovation in developing the model was to introduce a belief system that combines learning about asset prices with a high degree of rationality and internal consistency. Beliefs about variables other than asset prices are restricted to be such that forecast errors conditional on future prices and fundamentals are zero. This is different from rational expectations where forecast errors conditional on fundamentals alone are zero. But it is also different from the existing adaptive learning literature which usually parametrizes every forward-looking equation separately, resulting in a large number of degrees of freedom in specifying beliefs. The method can be widely applied in other models of the business cycle. Here, it led to forecast error predictability that closely matches survey data on expectations on a range of variables, despite the fact that learning only takes place about asset prices.

The model was also used to study normative implications of learning. In particular, I have revisited the question of whether monetary policy should react systematically to asset prices. I found that a reaction to stock price growth is desirable from a welfare perspective when investors in financial markets are learning. In contrast, under rational expectations, such a reaction does not improve welfare, in line with previous findings in the literature. These normative findings are certainly only a first step, but they illustrate that the choice of an asset pricing theory can have profound policy implications.

References

Adam, K., J. Beutel, and A. Marcet (2014): "Stock Price Booms and Expected Capital Gains," Working paper.

Adam, K., P. Kuang, and A. Marcet (2012): "House Price Booms and the Current Account," *NBER Macroeconomics Annual*, 26(1), 77–122.

Adam, K., and A. Marcet (2011): "Internal rationality, imperfect market knowledge and asset prices," *Journal of Economic Theory*, 146(3), 1224–1252.

Adam, K., A. Marcet, and J. P. Nicolini (2015): "Stock Market Volatility and Learning," *Journal of Finance*.

Ajello, A. (2016): "Financial intermediation, investment dynamics and business cycle fluctuations," *American Economic Review*, (forthcoming).

Altman, E. I. (2014): "The Role of Distressed Debt Markets, Hedge Funds and Recent Trends in Bankruptcy on the Outcomes of Chapter 11 Reorganizations," *American Bankruptcy Institute Law Review*, 22(1), 75–267.

Bansal, R., and A. Yaron (2004): "Risks for the Long Run: A Potential Resolution of Asset Pricing Puzzles," *Journal of Finance*, 59(4), 1481–1509.

Bernanke, B. S., and M. Gertler (2001): "Should Central Banks Respond to Movements in Asset Prices?," *American Economic Review*, 91(2), 253–257.

Bernanke, B. S., M. Gertler, and S. Gilchrist (1999): "The financial accelerator in a quantitative business cycle framework," in *Handbook of Macroeconomics*, ed. by J. B. Taylor, and M. Woodford, vol. 1, chap. 21, pp. 1341–1393. Elsevier.

Boldrin, M., L. J. Christiano, and J. D. M. Fisher (2001): "Habit Persistence, Asset Returns, and the Business Cycle," *American Economic Review*, 91(1), pp.149–166.

Campbell, J. Y., and J. H. Cochrane (1999): "By Force of Habit: A Consumption-Based Explanation of Aggregate Stock Market Behavior," *Journal of Political Economy*, 107(2), 205–251.

Campbell, J. Y., and L. Hentschel (1992): "No news is good news," *Journal of Financial Economics*, 31(3), 281 – 318.

Caputo, R., J. P. Medina, and C. Soto (2010): "The Financial Accelerator Under Learning and The Role of Monetary Policy," Working Papers of the Central Bank of Chile 590, Central Bank of Chile.

Coibion, O., and Y. Gorodnichenko (2015): "Information Rigidity and the Expectations Formation Process: A Simple Framework and New Facts," *American Economic Review*, 105(8), 2644–2678.

Cordoba, J.-C., and M. Ripoll (2004): *International Economic Review*45(4), 1011–1046.

Croce, M. M. (2014): "Long-run productivity risk: A new hope for production-based asset pricing?," *Journal of Monetary Economics*, 66(0), 13–31.

Epstein, L. G., E. Farhi, and T. Strzalecki (2013): "How Much Would You Pay to Resolve Long-Run Risk?," Working Paper 19541, NBER.

Faia, E., and T. Monacelli (2007): "Optimal interest rate rules, asset prices, and credit frictions," *Journal of Economic Dynamics and Control*, 31(10), 3228–3254.

Fama, E. F., and K. R. French (1988): "Dividend yields and expected stock returns," *Journal of Financial Economics*, 22(1), 3–25.

Flynn, E., and P. Crewson (2009): "Chapter 11 Filing Trends in History and Today," *American Bankruptcy Institute Journal*, 14, 65 ff.

Gali, J. (2014): "Monetary Policy and Rational Asset Price Bubbles," *American Economic Review*, 104(3), 721–752.

Gelain, P., K. J. Lansing, and C. Mendicino (2013): "House Prices, Credit Growth, and Excess Volatility: Implications for Monetary and Macroprudential Policy," *International Journal of Central Banking*, 9(2), 219–276.

Greenwood, R., and A. Shleifer (2014): "Expectations of Returns and Expected Returns," *Review of Financial Studies*, 27(3), 714–746.

Iacoviello, M. (2005): "House Prices, Borrowing Constraints, and Monetary Policy in the Business Cycle," *American Economic Review*, 95(3), 739–764.

Kobayashi, K., and K. Nutahara (2010): "Nominal Rigidities, News-Driven Business Cycles, and Monetary Policy," *B.E. Journal of Macroeconomics*, 10(1), 1–26.

Lawton, A. (2012): "Chapter 11 Triage: Diagnosing a Debtor's Prospects for Success," *Arizona Law Review*, 54(4), 985–1028.

Lettau, M., and H. Uhlig (2000): "Can Habit Formation Be Reconciled with Business Cycle Facts?," *Review of Economic Dynamics*, 3(1), 79–99.

Liu, Z., P. Wang, and T. Zha (2013): "Land-Price Dynamics and Macroeconomic Fluctuations," *Econometrica*, 81(3), 1147–1184.

Maussner, A. (2010): "The Analytics of New Keynesian Phillips Curves," Discussion Paper Series 313, Universitaet Augsburg, Institute for Economics.

Milani, F. (2011): "Expectation Shocks and Learning as Drivers of the Business Cycle," *Economic Journal*, 121(552), 379–401.

Sargent, T. J. (2008): "Rational Expectations," in *The Concise Encyclopedia of Economics*, ed. by D. R. Henderson. Library of Economics and Liberty, 2nd edn.

Schmitt-Grohe, S., and M. Uribe (2004): "Solving dynamic general equilibrium models using a second-order approximation to the policy function," *Journal of Economic Dynamics and Control*, 28(4), 755–775.

Shiller, R. (1981): "Do Stock Prices Move Too Much to be Justified by Subsequent Changes in Dividends?," *American Economic Review*, 71(3), 421–436.

Tallarini Jr., T. D. (2000): "Risk-sensitive real business cycles," *Journal of Monetary Economics*, 45(3), 507–532.

Warren, E., and J. L. Westbrook (2009): "The Success of Chapter 11: A Challenge to the Critics," *Michigan Law Review*, 107(4), 604–641.

Xu, Z., P. Wang, and J. Miao (2013): "A Bayesian DSGE Model of Stock Market Bubbles and Business Cycles," Meeting Paper 167, Society for Economic Dynamics.

A Further details on the full model

Retailers

Retailers transform a homogeneous intermediate good into differentiated final consumption goods using a one-for-one technology. The intermediate good trades in a competitive market at the real price q_t (expressed in units of the composite final good). Each retailer enjoys market power in her output market, though, and sets a nominal price p_t^i for her production. A standard price adjustment friction à la Calvo means that a retailer cannot adjust her price with probability κ, which is independent across retailers and across time. Hence, the retailer solves the following optimization:

$$\max_{P_{it}} \sum_{s=0}^{\infty} \left(\prod_{\tau=1}^{s} \kappa \Lambda_{t+\tau} \right) \left((1+\tau) P_{it} - q_{t+s} P_{t+s} \right) Y_{it+s}$$

$$\text{s.t. } Y_{it} = \left(\frac{P_{it}}{P_t} \right)^{-\sigma} \tilde{Y}_t$$

where $Q_{t,t+s}$ is the nominal discount factor of households between time t and $t+s$, and $\tilde{Y}_t$ is aggregate demand for the composite final good. Since all retailers that can re-optimize at t are identical, they all choose the same price $P_{it} = P_t^*$. Since I want to evaluate welfare in the model, I cannot log-linearize the first-order conditions of this problem. Their derivation is nevertheless standard (for example, Maussner, 2010) and I only report the final equations here:

$$\frac{P_t^*}{P_t} = \frac{1}{1+\tau} \frac{\sigma}{\sigma - 1} \frac{\Gamma_{1t}}{\Gamma_{2t}} \tag{A.1}$$

$$\Gamma_{1t} = q_t + \kappa \mathbb{E}_t^{\mathcal{P}} \Lambda_{t+1} \frac{\tilde{Y}_{t+1}}{\tilde{Y}_t} \pi_{t+1}^{\sigma} \tag{A.2}$$

$$\Gamma_{2t} = 1 + \kappa \mathbb{E}_t^{\mathcal{P}} \Lambda_{t+1} \frac{\tilde{Y}_{t+1}}{\tilde{Y}_t} \pi_{t+1}^{\sigma-1} \tag{A.3}$$

I assume that the government sets subsidies such that $\tau = 1/(\sigma - 1)$ so that the steady-state markup over marginal cost is zero. Inflation and the reset price are linked through the price aggregation equation which can be written as

$$1 = (1-\kappa) \left(\frac{P_t^*}{P_t} \right)^{1-\sigma} + \kappa \pi_t^{\sigma-1} \tag{A.4}$$

and the Tak-Yun distortion term is

$$\Delta_t = (1-\kappa) \left(\frac{\Gamma_{1t}}{\Gamma_{2t}} \right)^{-\sigma} + \kappa \pi_t^{\sigma} \Delta_{t-1}. \tag{A.5}$$

This term $\Delta_t \geq 1$ is the wedge due to price distortions between the amount of intermediate goods produced and the amount of the final good consumed. The amount of final goods available for consumption and investment is $\tilde{Y}_t = Y_t / \Delta_t$. Similarly, one can define $\tilde{C}_t = C_t / \Delta_t$ as the level of consumption the household could obtain if price distortions were zero.

Labor agencies

Similarly to retailers, labor agencies transform the homogeneous household labor input into differentiated labor goods at the nominal price $\tilde{w}_t P_t$ and sell them to intermediate firms at the price W_{ht}, which cannot be adjusted with probability κ_w. Labor agency h solves the following optimization:

$$\max_{W_{ht}} \mathbb{E}_t^{\mathcal{P}} \sum_{s=0}^{\infty} \left(\prod_{\tau=1}^{s} \kappa_w \Lambda_{t+\tau} \right) \left(\left(1+\tau_w\right) W_{ht} - \tilde{w}_{t+s} P_{t+s} \right) L_{ht+s}$$

$$\text{s.t. } L_{ht} = \left(\frac{W_{ht}}{W_t} \right)^{-\sigma_w} \tilde{L}_t$$

Since all labor agencies that can re-optimize at t are identical, they all choose the same price $W_{ht} = W_t^*$. The first-order conditions are analogous to those for retailers. Again, I assume that the government sets taxes such that $\tau = 1/(\sigma_w - 1)$ so that the steady-state markup over marginal cost is zero. Wage inflation π_{wt} and the Tak-Yun distortion Δ_{wt} are defined in the same way as for retailers. Finally, the real wage that intermediate producers effectively pay is

$$w_t = \frac{W_t}{P_t} = w_{t-1} \frac{\pi_{wt}}{\pi_t}. \tag{A.6}$$

Capital good producers

Capital good producers operate competitively in input and output markets, producing new capital goods using old final consumption goods. For the latter, they have a CES aggregator just like households. Their maximization program is entirely intratemporal:

$$\max_{I_t} Q_t I_t - \left(I_t + \frac{\psi}{2} \left(\frac{I_t}{I_{t-1}} - 1 \right)^2 \right)$$

In particular, they take past investment levels I_{t-1} as given when choosing current investment output. Their first-order condition defines the price for capital goods:

$$Q_t \quad = \quad 1 + \psi \left(\frac{I_t}{I_{t-1}} - 1 \right) \tag{A.7}$$

Market clearing

The market clearing conditions are summarized below. Supply stands on the left-hand side, demand on the right-hand side.

$$Y_t = \int_0^1 Y_{jt} dj = \int_0^1 Y_{it} di \tag{A.8}$$

$$\tilde{Y}_t = \frac{Y_t}{\Delta_t} = C_t + I_t + -\frac{\psi}{2}\left(\frac{I_t}{I_{t-1}} - 1\right)^2 + C_t^e \tag{A.9}$$

$$L_t = \int_0^1 L_{ht} dh \tag{A.10}$$

$$\tilde{L}_t = \frac{L_t}{\Delta_{wt}} = \int_0^1 L_{jt} dj \tag{A.11}$$

$$K_t = \int_0^1 K_{jt} dj = (1-\delta) K_{t-1} + I_t \tag{A.12}$$

$$1 = S_{jt},\ j \in [0,1] \tag{A.13}$$

$$0 = B_t^g \tag{A.14}$$

B Properties of the rational expectations equilibrium

The rational expectations equilibrium considered here has the following properties that need to be verified. All statements are local in the sense that for each of them, there exists a neighborhood of the non-stochastic steady-state in which the statement holds.

1. All firms choose the same capital-labor ratio K_{jt}/L_{jt} .
2. The expected return on capital is higher than the internal return on debt: $\mathbb{E}_t R^k_{t+1} > R_t$.
3. At any time t, the stock market valuation P_{jt} of a firm j is proportional to its net worth after entry and exit $\tilde{N}_{jt}$ with a slope that is strictly greater than one.
4. Borrowers never default on the equilibrium path and borrow at the risk-free rate, and the lender only accepts debt payments up to a certain limit.
5. If the firm defaults and the lender seizes the firm, it always prefers restructuring to liquidation.
6. The firm always exhausts the borrowing limit.
7. All firms can be aggregated. Aggregate debt, capital and net worth are sufficient to describe the intermediate goods sector.

I take the following steps to prove the existence of this equilibrium. After setting up the firm value functions, Property 1 just follows from constant returns to scale. I then take Properties 2 and 3 as given and prove 4 to 6. I verify that 3 holds. The aggregation property 7 is then easily verified. I conclude by establishing the parameter restrictions for which 2 holds.

Value functions

An operating firm j enters period t with a predetermined stock of capital and debt. It is convenient to decompose its value function into two stages. The first stage is given by:

$$\Upsilon_1(K,B) = \max_{N,L,D} \gamma N + (1-\gamma)(D + \Upsilon_2(N-D))$$

$$\begin{aligned} \text{s.t. } N &= qY - wL + (1-\delta) QK - RB \\ Y &= K^{\alpha}(AL)^{1-\alpha} \\ D &= \zeta(N - QK + B) \end{aligned}$$

(I suppress the time and firm indices for the sake of notation.) After production, the firm exits with probability γ and pays out all net worth as dividends. The second stage of the value function consists in choosing debt and capital levels as well as a strategy in the default game:

$$\begin{aligned} \Upsilon_2\left(\tilde{N}\right) = \max_{K',B',\text{strategy in default game}} \quad & \tilde{\beta}\mathbb{E}\left[\Upsilon_1(K',B'),\ \text{no default}\right] \\ + \quad & \tilde{\beta}\mathbb{E}\left[\Upsilon_1(K',B^*),\text{debt renegotiated}\right] \\ + \quad & \tilde{\beta}\mathbb{E}\left[0,\ \text{lender seizes firm}\right] \end{aligned}$$

$$\text{s.t. } K' = N + B'$$

A firm that only enters in the current period starts directly with an exogenous net worth endowment and the value function Υ_2.

Characterizing the first stage

The first order conditions for the first stage with respect to L equalizes the wage with the marginal revenue: $w = q(1-\alpha)(K/L)^{\alpha} A^{1-\alpha}$. Since there is no firm heterogeneity apart from capital K and debt B, this already implies Property 1 that all firms choose the same capital-labor ratio. Hence the internal rate of return on capital is common across firms:

$$R^k = \alpha q \left((1-\alpha)\frac{qA}{w}\right)^{\frac{1-\alpha}{\alpha}} + (1-\delta) Q \qquad \text{(B.1)}$$

Taking Property 3 as given for now, Υ_2 is a linear function with slope strictly greater than one.

Then the following holds for the first-stage value function Υ_1:

$$\begin{aligned}\Upsilon_1(K,B) &= N+(1-\gamma)(D-N+\Upsilon_2(N-D))\\ &= N+(1-\gamma)(\Upsilon_2'-1)((1-\zeta)N+\zeta(QK-B))\\ &> N\\ &= R^kK-RB \end{aligned} \tag{B.2}$$

This property will be used repeatedly in the next step of the proof.

Characterizing the second stage

The second stage involves solving for the subgame-perfect equilibrium of the default game between borrower and lender. Pairings are anonymous, so repeated interactions are ruled out. Also, only the size B and the interest rate $\tilde{R}$ of the loan can be contracted (I omit primes for ease of notation and separate $\tilde{R}$ from the risk-free rate R). The game is played sequentially:

1. The firm (F) proposes a borrowing contract $\left(B,\tilde{R}\right)$.
2. The lender (L) can accept or reject the contract.
    - A rejection corresponds to setting the contract $\left(B,\tilde{R}\right)=(0,0)$.
      Payoff for L: 0. Payoff for F: $\tilde{\beta}\mathbb{E}\left[\Upsilon_1\left(\tilde{N},0\right)\right]$.
3. F acquires capital and can then choose to default or not.
    - If F does not default, it has to repay in the next period.
      Payoff for L: $\mathbb{E}Q_{t,t+1}\tilde{R}B-B$. Payoff for F: $\tilde{\beta}\mathbb{E}\left[\Upsilon_1\left(K,\frac{\tilde{R}}{R}B\right)\right]$.
4. If F defaults, the debt needs to be renegotiated. F makes an offer for a new debt level B^*.[21]
5. L can accept or reject the offer.
    - If L accepts, the new debt level replaces the old one.
      Payoff for L: $\mathbb{E}\Lambda\tilde{R}B^*-B$. Payoff for F: $\tilde{\beta}\mathbb{E}\left[\Upsilon_1\left(K,\frac{\tilde{R}}{R}B^*\right)\right]$.
6. If L rejects, then she seizes the firm. A fraction $1-\xi$ of the firm's capital is lost in the process. Nature decides randomly whether the firm can be "restructured."
    - If the firm cannot be restructured, or it can but the lender chooses not to do so, then the lender has to liquidate the firm.
      Payoff for L: $\mathbb{E}\Lambda\xi QK-B$. Payoff for F: 0.

[21] That the interest rate on the repayment is fixed is without loss of generality.

- If the firm can be restructured and the lender chooses to do so, she retains a debt claim of present value ξB and sells the residual equity claim in the firm to another investor. Payoff for L: $\xi B + \tilde{\beta}\mathbb{E}\left[\Upsilon_1\left(\xi K, \xi B\right)\right] - B$. Payoff for F: 0.

Backward induction leads to the (unique) subgame-perfect equilibrium of this game. Start with the possibility of restructuring. L prefers this to liquidation if

$$\xi B + \tilde{\beta}\mathbb{E}\left[\Upsilon_1\left(\xi K, \xi B\right)\right] \geq \mathbb{E}\Lambda\xi Q K. \tag{B.3}$$

This holds true at the steady state because $R^k > R$ (Property 2), $Q = 1$, $\tilde{\beta} = \Lambda$ and

$$\begin{aligned} \xi B + \tilde{\beta}\mathbb{E}\left[\Upsilon_1\left(\xi K, \xi B\right)\right] &> \xi B + \tilde{\beta}\mathbb{E}\left[R^k \xi K - R\xi B\right] \\ &= \tilde{\beta}\mathbb{E}\left[R^k \xi K\right] \\ &> \xi K \end{aligned} \tag{B.4}$$

Since the inequality is strict, the statement holds in a neighborhood around the steady-state as well. This establishes Property 5.

Next, L will accept an offer B^* if it gives her a better expected payoff (assuming that lenders can diversify among borrowers so that their discount factor is invariant to the outcome of the game). The probability of restructuring is given by x. The condition for accepting B^* is therefore that

$$\mathbb{E}\Lambda\tilde{R}B^* \geq x\left(\xi B + \tilde{\beta}\mathbb{E}\left[\Upsilon_1\left(\xi K, \xi B\right)\right]\right) + (1-x)\,\mathbb{E}\Lambda\xi Q K. \tag{B.5}$$

Now turn to the firm F. Among the set of offers B^* that are accepted by L, the firm will prefer the lowest one—i.e., that which satisfies (B.5) with equality. This follows from Υ_1 being a decreasing function of debt. This lowest offer will be made if it leads to a higher payoff than expropriation: $\tilde{\beta}\mathbb{E}\left[\Upsilon_1\left(K, \frac{\tilde{R}}{R}B^*\right)\right] \geq 0$. Otherwise, F offers zero and L seizes the firm.

Going one more step backwards, F has to decide whether to declare default or not. It is preferable to do so if the B^* that L will just accept is strictly smaller than B or if expropriation is better than repaying, $\tilde{\beta}\mathbb{E}\left[\Upsilon_1\left(K, \frac{\tilde{R}}{R}B\right)\right] \geq 0$.

What is then the set of contracts that L accepts in the first place? From the perspective of L, there are two types of contracts: those that will not be defaulted on and those that will. If F does not default ($B^* \geq B$), L will accept the contract simply if it pays at least the risk-free rate, $\tilde{R} \geq R$. If F does default ($B^* < B$), then L accepts if the expected discounted recovery value exceeds the size of the loan—i.e., $\mathbb{E}\Lambda\tilde{R}B^* \geq B$.

Finally, let us consider the contract offer. F can offer a contract on which it will not default. In this case, it is optimal to offer just the risk-free rate $\tilde{R} = R$. Also note that the payoff from this strategy is strictly positive since

$$\begin{aligned}\tilde{\beta}\mathbb{E}\left[\Upsilon_1\left(K,B\right)\right] &> \tilde{\beta}\mathbb{E}\left[R^k K - RB\right] \\ &= \tilde{\beta}\mathbb{E}\left[R^k \tilde{N} + \left(R^k - R\right) B\right] \\ &> 0. \end{aligned} \tag{B.6}$$

The payoff is also increasing in the size of the loan B. So conditional on not defaulting, it is optimal for F to take out the maximum loan size $B = B^*$, and this is preferable to default with expropriation. However, it might also be possible for F to offer a contract that only leads to a default with debt renegotiation. The optimal contract of this type is the solution to the following problem:

$$\max_{\tilde{R},B,B^*} \tilde{\beta}\mathbb{E}\left[\Upsilon_1\left(\tilde{N} + B, \frac{\tilde{R}}{R}B^*\right)\right]$$

$$\begin{aligned}\text{s.t. } \mathbb{E}\Lambda\tilde{R}B^* &\geq B \\ \mathbb{E}\Lambda\tilde{R}B^* &= x\left(\xi B + \tilde{\beta}\mathbb{E}\left[\Upsilon_1\left(\xi\left(\tilde{N} + B\right), \xi B\right)\right]\right) \\ &\quad + (1-x)\,\mathbb{E}\Lambda Q\xi\left(\tilde{N} + B\right)\end{aligned}$$

The first thing to note is that only the product $\tilde{R}B^*$ appears, so the choice of the interest rate $\tilde{R}$ is redundant. Further, $B = B^*$ and $\tilde{R} = R$ solve this problem, and this amounts to the same as not declaring default. This choice solves the maximization problem above if the following condition is satisfied at the steady state:

$$\frac{\xi}{R}\left(1 - x + xR + x\Upsilon_1'\left[\frac{R^k}{R} - 1\right]\right) < 1 \tag{B.7}$$

For the degree of stock price dependence x sufficiently small, this condition is satisfied. This establishes Properties 4 and 6.

Linearity of firm value

Since firms do not default and exhaust the borrowing limit B^*, the second-stage firm value can be written as follows:

$$\Upsilon_2\left(\tilde{N}\right) = \tilde{\beta}\mathbb{E}\left[\Upsilon_1\left(\tilde{N} + B, B\right)\right] \tag{B.8}$$

$$\text{where } B = x\left(\xi B + \tilde{\beta}\mathbb{E}\left[\Upsilon_1\left(\xi\left(\tilde{N} + B\right), \xi B\right)\right]\right) + (1-x)\,Q_t\xi\left(\tilde{N} + B\right) \tag{B.9}$$

We already know that if Υ_2 is a linear function, then Υ_1 is also linear. The converse also holds: The constraint above, together with linearity of Υ_1 imply that B is linear in $\tilde{N}$, and thus Υ_2 is linear, too.

To establish Property 3, it remains to show that the slope of Υ_2 is greater than one. This is easy to see in steady state:

$$\begin{aligned}
\Upsilon_2' &= \tilde{\beta}\frac{\Upsilon_1(K,B)}{\tilde{N}} \\
&= \tilde{\beta}\frac{\gamma\left(R^kK-RB\right)+(1-\gamma)\Upsilon_2\left(R^kK-RB\right)}{\tilde{N}} \\
&= \tilde{\beta}\left(\gamma+(1-\gamma)\Upsilon_2'\right)\left(R^k\frac{K}{\tilde{N}}-R\frac{B}{\tilde{N}}\right) \\
&= \left(\gamma+(1-\gamma)\Upsilon_2'\right)\underbrace{\frac{R^k+\left(R^k-R\right)\frac{B}{\tilde{N}}}{R}}_{=:c_0>1} \\
&= \frac{\gamma c_0}{1-(1-\gamma)c_0} \\
&> 1 \qquad \text{(B.10)}
\end{aligned}$$

Finally, the aggregated law of motion for capital and net worth needs to be established (Property 7). Denoting again by $\Gamma_t \subset [0,1]$ the indices of firms that exit and are replaced in period t, we have

$$\begin{aligned}
K_t = \int_0^1 K_{jt}dj &= \int_{j\notin\Gamma_t}\left(N_{jt}-\zeta E_{jt}+B_{jt}\right)dj+\int_{j\in\Gamma_t}\left(\omega\left(N_t-\zeta E_t\right)+B_{jt}\right)dj \\
&= (1-\gamma+\gamma\omega)\left(N_t-\zeta E_t\right)+B_t \qquad \text{(B.11)} \\
N_t = \int_0^1 N_{jt}dj &= R_t^kK_{t-1}-R_{t-1}B_{t-1} \qquad \text{(B.12)} \\
B_t = \int_0^1 B_{jt}dj &= x\xi\left(B_t+P_t\right)+(1-x)\,\xi\mathbb{E}_t\Lambda_{t+1}Q_{t+1}K_t \qquad \text{(B.13)}
\end{aligned}$$

So far, then, all model properties are established except for $R^k > R$.

Return on capital

It can now be shown under which conditions the internal rate of return is indeed greater than the return on debt. From the steady-state versions of equations (B.11) and (B.12), it follows that

$$R^k = R+\left(G-R\left(1-\gamma+\gamma\omega\right)\right)\frac{\bar{N}}{\bar{K}}+Rc\left(1-\gamma+\gamma\omega\right)\frac{\bar{E}}{\bar{K}}. \qquad \text{(B.14)}$$

Sufficient conditions for $R^k > R$ are therefore that $\bar{N}/\bar{K}$ and $\bar{E}/\bar{K}$ are strictly positive and that the following holds:

$$\gamma > \frac{R-G}{G\left(1-\omega\right)}. \qquad \text{(B.15)}$$

C Approximation method for the learning equilibrium

The second-order perturbation method for the learning equilibrium follows Schmitt-Grohe and Uribe (2004) but has to be adapted to allow for relaxation of the rational expectations assumption. A rational expectations equilibrium can generally be described as a solution $(y_t)_{t\in\mathbb{N}}$ to

$$\mathbb{E}_t\left[f\left(y_{t+1}, y_t, x_t, u_t\right)\right] = 0, \tag{C.1}$$

where $\mathbb{E}_t$ is the expectations operator with respect to the probability measure and filtration induced by exogenous stochastic disturbances u_t. These disturbances are of dimensionality n_u, independent and identically distributed, of zero mean and variance $\sigma^2\Sigma_u$. The solution y_t is of dimensionality n, as is the image of f. x_t denotes a vector of predetermined state variables of dimensionality $n_x < n$: $x_{i,t} = y_{\iota(i),t-1}$ for an injective $\iota : \{1..n\} \to \{1..n_x\}$, or simply $x_t = Cy_{t-1}$ for an appropriate matrix C. One is interested in finding a policy function that generates solutions of the form

$$y_t = g\left(x_t, u_t, \sigma\right). \tag{C.2}$$

Perturbation methods for approximating the policy function to higher orders are straightforward. They compute Taylor expansions of g, typically around a non-stochastic steady state of the model: a constant solution $\bar{y}$ for $\sigma = 0$ such that $f\left(\bar{y}, \bar{y}, \bar{x}, 0\right) = 0$ and hence $g\left(\bar{y}, 0, 0\right) = \bar{y}$.

In a learning equilibrium, (C.1) does not fully characterize the equilibrium because the probability measure used by agents to form expectations does not coincide with the actual probability measure of the model. The stock price in the model of this paper is determined by the usual market-clearing condition, but agents think it is determined by random unpredictable shocks that are not necessarily related to the rest of the economy. A model described by (C.1) cannot contain a shock that is perceived as exogenous but at the same time is determined endogenously.

The model in this paper belongs to a class that can be written as follows:

$$\mathbb{E}_t^{\mathcal{P}}\left[f\left(y_{t+1}, y_t, x_t, u_t, z_t\right)\right] = 0 \tag{C.3}$$

$$\mathbb{E}_t^{\mathcal{P}}\left[\phi\left(y_{t+1}, y_t, x_t, u_t, z_t\right)\right] = 0 \tag{C.4}$$

Here, the probability measure $\mathcal{P}$ denotes beliefs for which the disturbances z_t (of dimensionality n_z) are perceived as exogenous, independent, and identically distributed with zero mean and variance $\sigma^2\Sigma_z$. They are also perceived as independent of u_t, although this can be relaxed. These disturbances have the interpretation of forecast errors. The iid assumption then amounts to imposing that agents holding the belief $\mathcal{P}$ think their forecasts cannot be improved upon. As before, f is of dimensionality n. The system (C.3) is assumed to have a unique solution for each initial condition x_t and path of disturbances u_t and z_t which can be described by a *subjective policy function:*

$$y_t = h\left(x_t, u_t, z_t, \sigma\right) \tag{C.5}$$

Contrary to agents' beliefs, z_t is not an exogenous disturbance but is determined endogenously by the second set of equilibrium conditions (C.4). The function ϕ is of dimension n_z. This set of conditions is not known to agents. The actual probability measure $\mathcal{P}_0$, induced by (C.3)–(C.4) and the disturbances u_t, is thus different from $\mathcal{P}$. Under $\mathcal{P}_0$, z_t is a function of the state and the fundamental disturbances:

$$z_t = r\left(x_t, u_t, \sigma\right) \tag{C.6}$$

This leads to the *objective policy function:*

$$g\left(x_t, u_t, \sigma\right) = h\left(x_t, u_t, r\left(x_t, u_t, \sigma\right), \sigma\right) \tag{C.7}$$

All functional forms are assumed to be such that the functions h and r are uniquely determined.

In the case of the model of this paper, lagged belief updating requires two pseudo-disturbances, $n_z = 2$. Agents cannot update their beliefs about future stock prices at the same time as they observe current prices, yet by observing the price they can infer the current forecast error. This implies the following subjective belief equations that are part of (C.3):

$$\log P_t = \log P_{t-1} + \hat{\mu}_t - \frac{\sigma_\eta^2 + \sigma_\nu^2}{2} + z_{1t} \tag{C.8}$$

$$\hat{\mu}_t = \hat{\mu}_{t-1} - \frac{\sigma_\nu^2}{2} + g z_{2t} \tag{C.9}$$

The conditions (C.4) that pin down the values for the forecast errors z_t in equilibrium are then described as follows:

$$P_t - \tilde{\beta}\mathbb{E}_t^{\mathcal{P}}\left[D_{t+1} + P_{t+1}\right] = 0 \tag{C.10}$$

$$z_{2t} - z_{1t-1} = 0 \tag{C.11}$$

Going back to the general case, the goal is to derive an accurate second-order approximation of the objective policy function g around the non-stochastic steady state:

$$\begin{aligned} g\left(x_t, u_t, \sigma\right) &\approx g\left(\bar{x}, 0, 0\right) \\ &+ g_x\left(x_t - \bar{x}\right) + g_u u_t + g_\sigma \sigma \\ &+ \frac{1}{2} g_{xx}\left[\left(x_t - \bar{x}\right) \otimes \left(x_t - \bar{x}\right)\right] + \frac{1}{2} g_{xu}\left[\left(x_t - \bar{x}\right) \otimes u_t\right] + \frac{1}{2} g_{uu}\left[u_t \otimes u_t\right] \\ &+ \frac{1}{2} g_{\sigma\sigma}\sigma^2 \end{aligned} \tag{C.12}$$

The first step in deriving the approximation is calculating this approximation for the subjective policy function h. This can be done using standard methods as implemented in Dynare. The second step is finding the derivatives of the function r in (C.6). Substituting it into the equilibrium conditions (C.4) gives

$$
\begin{aligned}
0=\Phi(x,u,\sigma) &= \mathbb{E}_t^{\mathcal{P}}\left[\phi\left(y',y,x,u,z\right)\right] \\
&= \mathbb{E}_t^{\mathcal{P}}\left[\phi\left(\begin{array}{c} h\left(Ch\left(x,u,z,\sigma\right),u',z',\sigma\right), \\ h\left(x,u,z,\sigma\right),x,u,z \end{array}\right)\right] \\
&= \mathbb{E}_t^{\mathcal{P}}\left[\phi\left(\begin{array}{c} h\left(Ch\left(x,u,r\left(x,u,\sigma\right),\sigma\right),u',z',\sigma\right), \\ h\left(x,u,r\left(x,u,\sigma\right),\sigma\right),x,u,r\left(x,u,\sigma\right) \end{array}\right)\right].
\end{aligned}
\tag{C.13}
$$

Here, I drop time subscripts and denote by prime variables at $t+1$. Note that the term z' must not be substituted out when the expectation is taken under $\mathcal{P}$. Doing so would imply that agents know the true relationship between z_{t+1} and the model variables instead of taking it as an exogenous disturbance. Total differentiation at the non-stochastic steady state leads to the following first-order derivatives:

$$
0=\frac{d\Phi}{dx}\left(\bar{x},0,0\right) = \left(\phi_{y'}h_xC+\phi_y\right)\left(h_x+h_zr_x\right)+\phi_x+\phi_zr_x \tag{C.14}
$$

$$
0=\frac{d\Phi}{du}\left(\bar{x},0,0\right) = \left(\phi_{y'}h_xC+\phi_y\right)\left(h_u+h_zr_u\right)+\phi_u+\phi_zr_u \tag{C.15}
$$

$$
0=\frac{d\Phi}{d\sigma}\left(\bar{x},0,0\right) = \left(\phi_{y'}h_xC+\phi_y\right)\left(h_\sigma+h_zr_\sigma\right)+\phi_zr_\sigma \tag{C.16}
$$

Since the existence of a unique solution for r is assumed, the first derivatives can be solved for. I also assume that the equilibrium conditions imply that $\bar{z}=0$ at the steady state. This means that in the absence of shocks, agents make no forecast errors under learning.

Define the matrix $A=\left(\phi_{y'}h_xC+\phi_y\right)h_z+\phi_z$. Then the first-order derivatives of r are given by

$$
r_x = -A^{-1}\left(\left(\phi_{y'}h_xC+\phi_y\right)h_x+\phi_x\right) \tag{C.17}
$$

$$
r_u = -A^{-1}\left(\left(\phi_{y'}h_xC+\phi_y\right)h_u+\phi_u\right) \tag{C.18}
$$

$$
r_\sigma = 0. \tag{C.19}
$$

Up to first order, the existence and uniqueness of the function r is equivalent to invertibility of the matrix A. The first-order derivatives of the actual policy function g can be obtained by applying the chain rule. The certainty-equivalence property holds for the subjective policy function h, hence $h_\sigma=0$. This implies that $r_\sigma=0$ and $g_\sigma=0$ as well, so certainty equivalence also holds under learning.

The second-order calculations are similar, if more tedious. The second-order derivative of Φ with respect to x is

$$0=\frac{d^2\Phi}{dx^2}\left(\bar{x},0,0\right) \;=\; \left(\phi_{y'}g_xC+\phi_y\right)\left(h_{xx}+2h_{xz}\left[I_{n_x}\otimes r_x\right]+h_{zz}\left[r_x\otimes r_x\right]\right) + \phi_{y'}h_{xx}\left[Cg_x\otimes Cg_x\right]+B_{xx}+Ar_{xx}. \tag{C.20}$$

This equation isn_x^2-dimensional and linear in r_{xx} and thus can be solved easily. As in first order, only invertibility of the matrix A is required for a unique local solution under learning. I have collected all cross-derivatives of ϕ inside the matrix B_{xx} (of size $n_z\times n_x^2$), which contains only first-order derivatives of the policy functions:

$$\begin{aligned} B_{xx} \;=\;& \phi_{y'y'}\left[h_xCg_x\otimes h_xCg_x\right]+\phi_{yy}\left[g_x\otimes g_x\right]+\phi_{xx}+\phi_{zz}\left[r_x\otimes r_x\right] \\ +\;& 2\phi_{y'y}\left[h_xCg_x\otimes g_x\right]+2\phi_{y'x}\left[h_xCg_x\otimes I_{n_x}\right]+2\phi_{y'z}\left[h_xCg_x\otimes r_x\right] \\ +\;& 2\phi_{yx}\left[g_x\otimes I_{n_x}\right]+2\phi_{yz}\left[g_x\otimes r_x\right]+2\phi_{xz}\left[I_{n_x}\otimes r_x\right] \end{aligned} \tag{C.21}$$

The formulae to solve for r_{xu} and r_{uu} are analogous. It remains to look at the derivatives involving σ. This simplifies considerably because the first derivatives of the policy functions g and h with respect to σ are zero. The cross-derivative of Φ with respect to x and σ thus reads:

$$0=\frac{d^2\Phi}{dxd\sigma}\left(\bar{x},0,0\right) \;=\; \left(\phi_{y'}g_xC+\phi_y\right)\left(h_{x\sigma}+h_zr_{x\sigma}\right)+\phi_zr_{x\sigma} \tag{C.22}$$

But because $h_{x\sigma}=0$, as under rational expectations, $r_{x\sigma}=0$ holds as well. The same applies to $r_{u\sigma}=0$. Finally, the second derivative with respect to σ involves the variance of the disturbances:

$$\begin{aligned} 0=\frac{d^2\Phi}{d\sigma^2}\left(\bar{x},0,0\right) \;=\;& \phi_{y'y'}\left(h_{\sigma\sigma}+h_{uu}\mathrm{vec}\left(\Sigma_u\right)++h_{xx}\mathrm{vec}\left(\Sigma_z\right)\right) \\ +\;& \phi_{y'y'}\left(\mathrm{vec}\left(h_u'\Sigma_uh_u\right)+\mathrm{vec}\left(h_z'\Sigma_zh_z\right)\right) \\ +\;& \left(\phi_{y'}g_xC+\phi_y\right)\left(h_{\sigma\sigma}\right)+Ar_{\sigma\sigma} \end{aligned} \tag{C.23}$$

Again, this can be solved for $r_{\sigma\sigma}$ when A is invertible. Note that the perceived variance Σ_z appears in the calculation because it matters for expectations for the future (unless $\phi_{y'y'}=0$). This variance is not necessarily equal to the objective variance of z.

The second-order derivatives of the actual policy function g are calculated easily once those of r are known:

$$g_{xx}=h_{xx}+2h_{xz}\left[I_{n_x}\otimes r_x\right]+g_{zz}\left[r_x\otimes r_x\right]+g_zr_{xx} \tag{C.24}$$

and analogously for g_{xu}, g_{uu} and $g_{\sigma\sigma}$. The cross-derivatives $g_{u\sigma}$ and $g_{x\sigma}$ are zero.