

The Role of Learning for Asset Prices and Business Cycles

Fabian Winkler*

1 March 2017

Abstract

I examine the implications of learning-based asset pricing in a model in which firms face credit constraints that depend partly on their market value. Agents learn about stock prices, but have conditionally model-consistent expectations otherwise. The model jointly matches key asset price and business cycle statistics, while the combination of financial frictions and learning produces powerful feedback between asset prices and real activity, adding substantial amplification. The model reproduces many patterns of forecast error predictability in survey data that are inconsistent with rational expectations. A reaction of the monetary policy rule to asset price growth increases welfare under learning.

JEL: D83, E32, E44, G12

Keywords: Learning, Expectations, Financial Frictions, Asset Pricing, Survey Forecasts


1 Introduction

Financial frictions are a central mechanism by which asset prices interact with macroeconomic dynamics. But to understand this interaction, we need models that are able to capture the dynamics of asset prices in the data. Moreover, the interaction is likely two-sided, with feedback from asset prices to the real economy, but also from the real economy to asset prices. Asset price dynamics observed in the data should therefore arise endogenously in our models.

At the same time, there is evidence that measures of expectations do not conform to the rational expectations hypothesis, even in financial markets. The rational expectations hypothesis implies,

*Board of Governors of the Federal Reserve System, 20th St and Constitution Ave NW, Washington DC 20551, fabian.winkler@frb.gov. I thank my former advisors Wouter den Haan and Albert Marcet, my discussant at the NBER Summer Institute Martin Schneider, as well as Klaus Adam, Andrea Ajello, Johannes Boehm, Rachel Ngai, Markus Riegler, and participants at the NBER Summer Institute 2016, Bank of Spain, Econometric Society World Congress in Montreal, Econometric Society Winter Meeting in Madrid, EEA Annual Meeting in Toulouse, ECB-CFS-Bundesbank lunchtime seminar, and Singapore Management University, for 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.

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 2000Q3–2012Q4. Correlation coefficient for return forecasts $\rho = .54$, for realized returns $\rho = -.44$.

for example, that investors are fully aware of return predictability in the stock market, expecting low returns when prices are high and vice versa. Instead, in survey data they expect high returns when prices are high. 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. This contrasts sharply with the actual return predictability in the right panel of the figure, where the correlation is strongly negative. This discrepancy is easy to detect and therefore hard to reconcile with rational expectations-based asset pricing theories such as the long-run risk, habit, and disaster risk models, but it is consistent with asset pricing theories based on extrapolative expectations or learning. [Adam, Marcet and Nicolini \(2015\)](#) develop a learning-based asset pricing theory that, in an endowment economy, is able to reproduce the discrepancy between subjective and rational return expectations as well as many so-called asset price puzzles.

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 firms' access to credit depends on their market value. This dependency implies that stock market valuations affect the availability of credit and therefore investment. At the same time, agents do not have rational expectations but are instead learning about price growth in the stock market. Equilibrium prices are determined endogenously and depend on agents' beliefs.

Deviating from rational expectations is a tricky business. One needs to explicitly spell out the entire belief formation process, how agents form expectations about future income, interest rates

and so on. I develop a restriction on expectations which requires that agents' expectations remain consistent with equilibrium conditions other than stock market clearing. It implies that agents make the best possible forecasts compatible with their subjective beliefs about stock prices. These “conditionally model-consistent expectations” allow me to study the effects of asset price learning in isolation while retaining most of the familiar logic and parsimony of rational expectations. The method is general and can be of interest in other situations where limited departures from rational expectations are desirable.

The model with learning jointly matches key business cycle and asset price moments, including the volatility and predictability of returns, negative skewness and heavy tails. Learning also considerably magnifies the strength of the financial accelerator. The model therefore addresses the [Kocherlakota \(2000\)](#) critique that amplification of shocks through financial frictions is usually quite weak. This weakness is due in part to the low endogenous asset price volatility in standard models (as previously observed by [Quadrini, 2011](#)). Under learning, a positive feedback loop emerges between asset prices and the production side of the economy. When beliefs of learning investors are more optimistic, their demand for stocks increases. This raises firm valuations and relaxes credit constraints, in turn allowing firms to move closer to their profit optimum. Firms are able to pay higher dividends to their shareholders, raising stock prices further and propagating investor optimism. The interaction between learning and financial frictions is stronger when credit constraints are tighter.

I also compare the forecast error predictability implied by the model with that found in survey forecasts. Such comparisons can be used to discriminate among models of expectation formation ([Manski, 2004](#)). I compare forecast error predictability in the model and in the data on a range of variables including stock returns, real GDP and its main components, inflation and interest rates. Rational expectations imply the absence of any forecast error predictability and are rejected by the data. By contrast, the learning model replicates many patterns of forecast error predictability remarkably well. Although agents learn only about stock prices, their expectational errors spill over into their other forecasts as well. For example, when agents are too optimistic about future stock prices, they also become too optimistic about the tightness of credit constraints and therefore over-predict real economic activity. The price dividend-ratio therefore predicts not only forecast errors on stock returns but also on GDP, and this holds true in survey data, too. This and other patterns of predictability discussed in the paper provide strong evidence in favor of the model with learning.

Assumptions on expectation formation carry not only positive but also normative implications, which I explore at the end of the paper. Specifically, I find that a positive response to stock price growth in the interest rate rule—a form of “leaning against the wind”—is welfare-increasing in the model with learning because it helps to stabilize expectations in financial markets and to mitigate asset price fluctuations. In the rational expectations version of the model, such a reaction is not beneficial.

The remainder of the paper is structured as follows. Section [2](#) briefly discusses the related

literature. Section 3 presents the model and discusses the setup of expectation formation under learning. To gain some understanding of the interaction between asset price learning and financial frictions, Section 4 discusses a special case of the model that shuts off auxiliary frictions and allows for a closed-form solution. Section 5 discusses the quantitative fit of the full model on the asset pricing and business cycle side, while Section (6) compares survey data on expectations with model-implied forecasts. Section 7 provides a discussion of other aspects of the model, namely the effects of learning beyond the effect on asset prices; the possibility of allowing for mean reversion in asset price expectations; and the role of nominal rigidities for the quantitative fit of the model. Section 8 explores implications for monetary policy. Section 9 concludes.

2 Related literature

This paper builds on the learning-based asset pricing theory developed in Adam and Marcet (2011) and Adam et al. (2015, 2016) who show that, in an endowment economy, learning about asset prices can explain many so-called asset price puzzles. Barberis et al. (2015), using a similar model, show that the price fluctuations induced by learning even survive the introduction of some traders with rational expectations. This paper is the first to employ this theory in a business cycle context.

There are a number of papers in the adaptive learning literature that examine linkages between asset prices and the real economy, including Milani (2008, 2011), Caputo et al. (2010) and Gelain et al. (2013), to name just a few. In that literature, agents typically learn about all forward-looking variables simultaneously in an otherwise linear model. This paper introduces a different approach to learning, focusing on learning about one variable while keeping expectations close to rational otherwise, and thereby retaining most of the intuition and the parsimony of rational expectations models. It also does not require the underlying model to be linear.

The paper also relates to a number of studies that try to explain asset price fluctuations with non-rational beliefs about exogenous fundamentals. This idea goes back to Timmermann (1996) and has recently become more prominent with applications by Fuster et al. (2012), Hirshleifer et al. (2015), and Collin-Dufresne et al. (2016). The idea is that agents have an incomplete understanding of some exogenous process such as dividends or consumption in an endowment economy or productivity growth in a production economy and estimate a subjective model for that process. Because the subjective model is often simpler than the true data-generating process and hard to reject in small samples, Fuster et al. (2012) refer to such expectations as “natural expectations”. Pintus and Suda (2013) and Pancrazi and Pietrunti (2014) also use this type of expectations to study lending and borrowing in the housing market. One limitation of this approach is that there is no feedback from asset prices to beliefs, as beliefs depend on exogenous fundamentals only. As such, the amount of endogenous amplification is limited (Timmermann, 1996). In this paper, agents instead learn about the endogenous price that depends itself on beliefs. This leads to two-sided feedback between learning and asset prices (and also to and from the real economy), considerably magnifying volatility.

The paper also contributes to the literature on financial frictions as an amplification channel. The early literature on financial frictions emphasized amplification of standard productivity or monetary policy shocks (Kiyotaki and Moore, 1997; Bernanke and Gertler, 2001), but the quantitative importance of the “financial accelerator” mechanism was found to be small (Kocherlakota, 2000; Cordoba and Ripoll, 2004; Quadrini, 2011). The more recent literature on financial frictions has found sizable amplification effects in response to shocks to borrowing constraints (e.g. Jermann and Quadrini, 2012) or shocks that directly move collateral prices (e.g. Liu et al., 2013) and emphasizes their role as driving forces of the business cycle. Other proposed solutions to the Kocherlakota critique explore occasionally binding constraints (Mendoza, 2010) or other non-linearities (Brunnermeier and Sannikov, 2014) which lead to amplification in severe crisis states. Instead, this paper takes a different approach by going back to the question of financial frictions as an amplification mechanism for standard and frequent business cycle shocks that do not directly impact credit constraints or asset prices. Learning endogenously generates volatility in asset prices and interacts with financial frictions to form a feedback loop that amplifies even standard productivity shocks.

The dependency of the borrowing constraint on equity valuations is similar to the one developed in Miao et al. (2015). In their model, rational liquidity bubbles exist which allow for a sunspot shock that governs the size of the bubble and drives the bulk of the variation in equity prices.¹ In this paper, asset price volatility instead arises endogenously through learning, which interacts with financial frictions to form a feedback loop that amplifies standard business cycle shocks.

Finally, the paper relates to the literature testing the rational expectations hypothesis with survey data. It is well known that expectations measured in surveys fail to conform to the rational expectations hypothesis because forecast errors are statistically predictable (e.g. Bacchetta et al., 2009; Andrade and Le Bihan, 2013; Gennaioli et al., 2016). Coibion and Gorodnichenko (2015) document predictability by forecast revisions across a range of variables and interpret their results as evidence in favor of rational inattention models. The learning model in this paper also matches predictability by forecast revisions, as well as by the level and first difference of the price dividend ratio. It does so with learning about asset prices as the only departure from rational expectations. To my knowledge, a model-data comparison of forecast error predictability within a single, general equilibrium model and across the range of variables presented in this paper has not been carried out previously in the literature.

3 The model

3.1 Model setup

The economy is closed and operates in discrete time. It is populated by two types of households. *Lending households* consume final goods and supply labor. They are risk-averse, trade debt claims on intermediate goods producers and receive interest from them. *Firm owners* only consume final

¹In principle, such bubbly sunspot equilibria can arise in my model as well, but they do not exist in the range of parameter values considered (see the Appendix).

goods. They are risk-neutral and trade equity claims on firms and receive dividends from them. The intermediate goods producers, or simply *firms*, are at the heart of the model. They combine capital and labor into intermediate goods and are financially constrained.

In addition, the model has nominal price and wage rigidities and investment adjustment costs. These are introduced by adding other types of firms into the model (which are all owned by the lending households): Final good producers transform intermediate goods into differentiated final goods and set prices; labor agencies transform homogeneous household labor into differentiated labor services and set wages; and capital goods producers produce new capital goods from final consumption goods. 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. Since these additional elements of the model are standard (e.g. [Christiano et al., 2005](#)), their detailed description is relegated to the appendix. A simplified version of the model without nominal rigidities and adjustment costs is discussed in [Section 4](#).

3.1.1 Households

Households with time-separable preferences maximizes utility as follows:

$$\begin{aligned} \max_{(C_t, L_t, B_{jt}, B_t^g)_{t=0}^{\infty}} \quad & \mathbb{E}^{\mathcal{P}} \sum_{t=0}^{\infty} \beta^t \frac{C_t^{1-\theta}}{1-\theta} - \eta \frac{L_t^{1+\phi}}{1+\phi} \\ \text{s.t. } C_t = \quad & \tilde{w}_t L_t + B_t^g - (1 + i_{t-1}) \frac{p_{t-1}}{p_t} B_{t-1}^g + \int_0^1 (B_{jt} - R_{jt-1} B_{jt-1}) dj + \Pi_t \end{aligned}$$

Here, \tilde{w}_t is the real wage received by the household and L_t is the amount of labor supplied. Consumption C_t is a standard Dixit-Stiglitz aggregator of differentiated consumption varieties with elasticity of substitution σ_p . 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. 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. Households do not trade equity claims.² Π_t represents lump-sum profits received from price- and wage-setting firms and capital goods producers, as described in the appendix.

The expectation operator is evaluated under the probability measure \mathcal{P} . Agents use this measure when forming their expectations to solve their optimization problems. Under learning, the distribution of outcomes expected under \mathcal{P} does not necessarily coincide with the distribution induced

²The fact that households do not trade equity claims is not technically necessary, but helps the model fit business cycle properties. If households held equity in this model, fluctuations in beliefs under learning would introduce strong wealth effects: Agents that are optimistic about future stock returns have higher expected financial wealth and therefore work less and reduce their savings, which can lead to a counterfactually low rise or even a fall in investment and hours worked during expansions. Here, household wealth effects from changes in beliefs are present as well, but are more muted because they operate only through changes in expected labor income rather than expected financial wealth. A similar problem is known in the news shock literature ([Beaudry and Portier, 2007](#)).

in equilibrium.

The first-order conditions are standard. In what follows I define the stochastic discount factor of the household as $\Lambda_{t+1} = \beta (C_{t+1}/C_t)^{-\theta}$.

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 labor L_{jt} to produce output:

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

where A_t is aggregate productivity. Labor is a CES combination of differentiated labor services with elasticity of substitution σ_w , but the firm's problem can be treated as if the labor index was acquired in a competitive market at the real wage index w_t . Output is sold competitively to final good producers at price q_t . 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:

$$N_{jt} = q_t Y_{jt} - w_t L_{jt} + Q_t (1 - \delta) K_{jt-1} - R_{jt-1} B_{jt-1}. \quad (2)$$

I assume that firms exit with probability γ . This probability is exogenous and independent across time and firms. As in [Bernanke et al. \(1999\)](#), 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 a regular dividend (where earnings E_{jt} are given by $N_{jt} - Q_t K_{jt-1} + B_{jt-1}$). The number ζ therefore represents the dividend payout ratio for continuing firms.³ The 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:

$$Q_t K_{jt} = B_t^j + N_{jt} - \zeta E_{jt} \quad (3)$$

If a firm does exit, it pays out its entire net worth as a terminal dividend.

3.1.3 Firm owners

Firm owners differ from households in their capacity to own intermediate firms. They are risk-neutral and trade equity claims on firms indexed by $j \in [0, 1]$. As described above, if a firm exits, it pays out its net worth N_{jt} as a terminal dividend. Otherwise it pays a regular dividend of

³The 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 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.

ζE_{jt} . Denote the subset of firms alive at the end of period t by Γ_t . Then, a firm owner's utility maximization problem is given by:

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

$$\text{s.t. } C_t^f + \int_{j \in \Gamma_t} S_{jt} P_{jt} dj = \int_{j \in \Gamma_t} S_{jt-1} (P_{jt} + \zeta E_{jt}) dj \quad (4)$$

$$+ \int_{j \in \Gamma_{t-1} \setminus \Gamma_t} S_{jt-1} N_{jt} dj \quad (5)$$

$$S_{jt} \in [0, \bar{S}] \quad (6)$$

where $\bar{S} > 1$. Consumption C_t^f is the same Dixit-Stiglitz aggregator of consumption varieties as for households. Firm owners do not trade debt claims.

The first term on the right-hand side of the budget constraint deals with continuing firms while the second term deals with exiting firms. In addition, firm owners face upper and lower bounds on traded stock holdings. This renders demand for stocks finite under arbitrary beliefs. In equilibrium, the bounds are never binding. The first-order condition of the firm owner is:

$$\left. \begin{aligned} S_{jt} &= 0 & \text{if } P_{jt} > \\ S_{jt} &\in [0, \bar{S}] & \text{if } P_{jt} = \\ S_{jt} &= \bar{S} & \text{if } P_{jt} < \end{aligned} \right\} \beta \mathbb{E}_t^\mathcal{P} [N_{jt+1} \mathbf{1}_{\{j \notin \Gamma_{t+1}\}} + (\zeta E_{jt+1} + P_{jt+1}) \mathbf{1}_{\{j \in \Gamma_{t+1}\}}] \forall j \in \Gamma_t. \quad (7)$$

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 the market value of the firm. Effectively, it introduces a link between stock market valuations and investment.

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 $\xi (P_{jt} + B_{jt})$ and this will always be higher than the recovery value after liquidation.

In the appendix, I show that the optimal debt contract in this limited commitment problem takes the form of a leverage constraint with a weighted average of liquidation and market value of the firm:⁴

$$B_{jt} \leq (1-x) \underbrace{\mathbb{E}_t^P \Lambda_{t+1} Q_{t+1} \xi K_{jt}}_{\text{liquidation value}} + x \underbrace{\xi (P_{jt} + B_{jt})}_{\text{market value}} \quad (8)$$

It is worth noting that even under rational expectations, the market value of the firm $P_{jt} + B_{jt}$ is different from the value of its capital stock, $Q_t K_{jt}$. Because financial frictions prevent firms from investing up to the efficient level, so that the marginal discounted revenue on capital exceeds the marginal cost of capital. Therefore, the market value of the firm is higher than the value of its capital stock.

The borrowing constraint acts as a link between firm investment and equity valuations. It is well known that in the data, stock prices comove with investment, too ([Barro, 1990](#)). This could be simply because news about investment opportunities affect stock prices and predict investment at the same time [Blanchard et al. \(1993\)](#). But the literature has documented evidence that firms’ investment depends on equity valuations beyond fundamentals (e.g. [Baker et al., 2003](#) for equity-dependent firms; and more recently [Hau and Lai, 2013](#) for firms whose shares were subjected to fire sales by distressed equity funds in the 2007–2009 financial crisis). While not the only one, the borrowing constraint developed here is one possible explanation for this dependency.⁵

3.1.5 Further model elements and shocks

To improve the quantitative fit of the model, I add standard nominal rigidities and adjustment costs (details are provided in the appendix). The prices for final goods and labor are subject to Calvo rigidities, with probabilities of non-adjustment κ and κ_w and elasticities of substitution σ and σ_w , respectively. The price for intermediate goods q_t equals the inverse of the gross markup of final goods producers. The monetary authority sets the nominal interest rate according to a

⁴The borrowing constraint is similar to that in [Miao and Wang \(2011\)](#) who develop a model where $x = 1$ but where only a fraction of firms are constrained, and show that this type of borrowing constraint can lead to sunspot equilibria. The appendix shows that such multiple equilibria do not arise for the parameter values considered in this paper.

⁵Other studies examine models in which firms’ credit constraints depend on the value of their real estate rather than their equity value ([Liu et al., 2013](#)). Here, too, the empirical evidence is not clear-cut. For example, [Chakraborty et al. \(2016\)](#) argue that price increases in real estate might induce lenders to substitute commercial lending with mortgage lending, thereby tightening credit constraints.

Taylor-type interest rate rule:

$$i_t = \rho_i i_{t-1} + (1 - \rho_i) (\beta^{-1} + \phi_\pi \pi_t + \varepsilon_{it}), \quad (9)$$

where ϕ_π is the reaction coefficient on consumer price inflation, ρ_i is the degree of interest rate smoothing, and ε_{it} is an interest rate shock. Capital needs to be purchased from capital goods producers, owned by the household, whose technology of transforming consumption to capital goods is subject to standard quadratic adjustment costs that move the price for capital goods:

$$Q_t = 1 + \psi \left(\frac{I_t}{I_{t-1}} - 1 \right) \quad (10)$$

Productivity evolves as an AR(1) process and there are exogenous shocks to productivity and the nominal interest rate.

$$\log A_t = (1 - \rho) \log \bar{A} + \rho \log A_{t-1} + \log \varepsilon_{At} \quad (11)$$

$$\varepsilon_{At} \sim iid \mathcal{N}(0, \sigma_A^2) \quad (12)$$

$$\varepsilon_{it} \sim iid \mathcal{N}(0, \sigma_i^2) \quad (13)$$

3.2 Rational expectations equilibrium

I first describe the equilibrium under rational expectations. A rational expectations 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} coincides with the actual distribution of equilibrium outcomes.

Under a mild restriction on the exit probability γ , there exists a rational expectations equilibrium characterized by the following properties in a neighborhood of the non-stochastic steady state.

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:

$$Y_t = \alpha K_{t-1}^\alpha \left(A_t \tilde{L}_t \right)^{1-\alpha} \quad (14)$$

$$R_t^k = q_t \alpha \frac{Y_t}{K_{t-1}} + Q_t (1 - \delta) K_{t-1} \quad (15)$$

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 a linear function of its post-dividend net worth $Q_t K_t - B_t$. This permits one to write an aggregate stock market index

as

$$P_t = \int_0^1 P_{jt} = \beta \mathbb{E}_t [D_{t+1} + P_{t+1}]. \quad (16)$$

where aggregate dividends are given by $D_t = \gamma N_t + \zeta E_t$.

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

$$R_{jt} = R_t = \frac{1}{\mathbb{E}_t \Lambda_{t+1}}. \quad (17)$$

The lender only accepts debt payments up to the limit given by (8), which is linear in the firm's net worth, and the firm always exhausts this limit, so that the borrowing constraint is always binding.

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:

$$N_t = R_t^k K_{t-1} - R_{t-1} B_{t-1} \quad (18)$$

$$Q_t K_t = (1 - \gamma) ((1 - \zeta) N_t + \zeta (B_{t-1} - Q_t K_{t-1})) + B_t \quad (19)$$

$$B_t = x \mathbb{E}_t \Lambda_{t+1} Q_{t+1} \xi K_t + (1 - x) \xi (P_t + B_t). \quad (20)$$

I prove these equilibrium properties in the appendix.

3.3 Learning equilibrium

I require that the equilibrium under learning satisfies internal rationality ([Adam and Marcet, 2011](#)). Specifically, given a subjective belief measure \mathcal{P} of the distribution of model outcomes, an internally rational equilibrium is a set of stochastic processes for prices and allocations, and a set of strategies in the limited commitment game 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} ; and the strategies in the limited commitment game are a subgame-perfect Nash equilibrium for all lender-borrower pairs. The only difference from a rational expectations equilibrium is that the measure \mathcal{P} under which agents evaluate expectations can differ from the actual distribution of model outcomes.⁶

How, then, should one choose the subjective beliefs \mathcal{P} ? In a forward-looking business cycle model, there are many ways in which expectations affect equilibrium. Households need to form expectations about future interest rates and wages to determine their savings, investors need to form expectations about stock prices, firms need to forecast future productivity and wages in order to decide their demand for capital and so forth. This leaves many degrees of freedom to be filled. My focus in this paper is to concentrate on the effects of stock price learning, while keeping the model dynamics as close as possible to a rational expectations equilibrium otherwise.

⁶It then becomes crucial to distinguish the subjective expectations $\mathbb{E}^{\mathcal{P}}[\cdot]$ taken under \mathcal{P} from the expectation $\mathbb{E}[\cdot]$ taken under the actual distribution of model outcomes.

To this end, I construct the belief \mathcal{P} in two steps: First, I specify the belief about stock prices that give rise to a learning problem; second, I impose what I call *conditionally model-consistent expectations* with respect to stock prices.

Under the subjective belief measure \mathcal{P} , agents are not endowed with the knowledge of the equilibrium pricing function for stocks, i.e. the mapping of state variables and shocks to prices P_{jt} . To keep things simple and retain the linear aggregation property, I assume that agents do believe that the value of an individual firm is proportional to its net worth, as under rational expectations:

$$P_{jt} = \frac{\mathbb{E}_t^{\mathcal{P}} N_{jt+1}}{\mathbb{E}_t^{\mathcal{P}} N_{t+1}} P_t. \quad (21)$$

However, they are uncertain about the evolution of the aggregate market value P_t . They employ a simple subjective model to forecast the market value, according to which P_t evolves as a random walk with a small unobservable drift:

$$\begin{aligned} \log P_t &= \log P_{t-1} + \mu_t + \eta_t, \quad \eta_t \sim \mathcal{N}(0, \sigma_\eta^2) \\ \mu_t &= \mu_{t-1} + \nu_t, \quad \nu_t \sim \mathcal{N}(0, \sigma_\nu^2) \end{aligned}$$

where the innovations η_t and ν_t are independent of each other and across time, and also independent of the other exogenous shocks in the model.⁷

Importantly, the innovations are unobservable, as is the mean growth rate μ_t . The only observable in the system above is the price P_t . Bayesian updating of this belief system amounts to a simple Kalman filtering problem. With an appropriate prior, the system above can be rewritten in terms of observables only:

$$\log P_t = \log P_{t-1} + \hat{\mu}_{t-1} + z_t \quad (22)$$

$$\hat{\mu}_t = \hat{\mu}_{t-1} + g z_t, \quad (23)$$

where $\hat{\mu}_t = \mathbb{E}_t^{\mathcal{P}} \mu_t$ is the mean belief about the trend in stock price growth, and z_t is the *subjective forecast error*. Under \mathcal{P} , z_t is normally distributed white noise with variance σ_z^2 , independent of the structural shocks ε_{A_t} and ε_{it} . The belief $\hat{\mu}_t$ is updated in the direction of the last forecast error: When agents see stock prices rising faster than they expected, they will also expect them to rise by more in the future. The parameter g is the *learning gain* which governs the speed of adjustment of price growth expectations.⁸

⁷This subjective model is the one selected by Adam et al. (2016) for its good asset pricing properties in an endowment economy. One particularly strong departure from rational expectations implied by the subjective model is that agents think that prices can diverge forever from economic fundamentals. In equilibrium however, prices and dividends (and other fundamentals) will turn out to be cointegrated, contrary to agents' expectations. Section 7.2 shows that it is possible to generalize this belief system to allow for mean reversion in prices, as long as beliefs about price growth remain very persistent in the short run.

⁸The parameters of the state-space system and the observable system map into each other through the relations $\sigma_\eta^2 = \sigma_p^2(1-g)$ and $\sigma_\nu^2 = \sigma_p^2 g^2$. Intuitively, a small learning gain g implies that agents believe there is little predictability in stock prices.

While agents think that P_t is a random walk with a drift, in equilibrium this will generally not be the case. The equilibrium stock price must be such that the stock market clears. This implies that the Euler equation (7) has to hold with equality in the aggregate:

$$\begin{aligned} P_t &= \beta \mathbb{E}_t^{\mathcal{P}} [P_{t+1} + D_{t+1}] \\ &= \beta \mathbb{E}_t^{\mathcal{P}} [P_t \exp(\hat{\mu}_t + z_{t+1})] + \beta \mathbb{E}_t^{\mathcal{P}} [D_{t+1}] \\ &= \frac{\beta \mathbb{E}_t^{\mathcal{P}} [D_{t+1}]}{1 - \beta \exp(\hat{\mu}_t + \frac{1}{2}\sigma_z^2)}. \end{aligned} \tag{24}$$

The value for the forecast error z_t that equates (22) and (24) is not white noise but depends on the belief about future dividends and price growth, even though agents remain ignorant of this fact.⁹ The expression also shows that, since the denominator is close to zero, small changes in the belief $\hat{\mu}_t$ can induce large swings in equilibrium asset prices.

The specification of the subjective probability measure \mathcal{P} is not complete yet. It does not yet specify how to calculate the expectation of future dividends $\mathbb{E}_t^{\mathcal{P}} [D_{t+1}]$ above, or indeed any expectation of other model variables. The subjective probability measure \mathcal{P} has to fully specify expectations about all variables in the model in order for an equilibrium to be well-defined. In principle, there is an enormous amount of degrees of freedom still left to specify \mathcal{P} : Once we move away from rational expectations, agents could entertain arbitrary beliefs about dividends, inflation, interest rates, wages and so on. Here is where I introduce the concept of conditionally model-consistent expectations in order to impose discipline on the expectation formation process.

Definition. Consider an internally rational equilibrium with beliefs \mathcal{P} . Agents' subjective expectations \mathcal{P} are *conditionally model-consistent* if:

1. the distribution under \mathcal{P} of the exogenous shocks, denoted u_t , coincides with their actual distribution, and
2. for any endogenous model variable y_t and $t \geq 0$:

$$\mathbb{E}_t^{\mathcal{P}} [y_{t+1} \mid u_{t+1}, P_{t+1}] = y_{t+1} \tag{25}$$

almost everywhere in equilibrium.

Conditional model consistency restricts the subjective belief \mathcal{P} to have the maximum degree of consistency with the model given agents' misspecified belief about stock prices. Agents endowed with such expectations may not know the equilibrium pricing function, but they make the smallest possible expectational errors consistent with their subjective view about the evolution of stock prices. Conditionally model-consistent expectations are close to but different from rational expectations, for which condition (25) would be strengthened to $\mathbb{E}_t^{\mathcal{P}} [y_{t+1} \mid u_{t+1}] = y_{t+1}$. In particular, agents' systematic forecast errors on stock prices spill over to their forecasts of other variables:

⁹The corresponding equilibrium values for the subjective forecasting error z_t and the belief $\hat{\mu}_t$ are $z_t = \Delta \log P_t - \hat{\mu}_{t-1}$ and $\hat{\mu}_t = (1 - g)\hat{\mu}_{t-1} + g\Delta \log P_{t-1}$.

When households are overly optimistic about the stock market, they expect borrowing constraints to be loose and are therefore overly optimistic about how much they will be able to consume in the future. These expectational spillovers are what enables the model to match the predictability of forecast errors in Section 6; they also lead to aggregate demand effects in the presence of nominal rigidities which are discussed in Section 7.3.

I construct expectations satisfying conditional model-consistency as follows. I solve for expectations that would be model-consistent in a model in which the stock price really evolved according to agents' subjective beliefs. I take the set of equilibrium conditions and remove the market clearing conditions for stocks and consumption goods,¹⁰ replacing it with the subjective belief system (22)–(23). I then solve the model as under rational expectations, which leads to a subjective policy function $y_t = h(y_{t-1}, u_t, z_t)$ where y_t is the collection of all endogenous model variables, u_t is the collection of exogenous shocks, and z_t is the subjective forecast error on stock prices. The policy function h together with the joint distribution of (u_t, z_t) defines the subjective probability measure \mathcal{P} . In order to solve for the equilibrium, I then impose the market clearing condition in the stock market. This leads to a solution for the equilibrium stock price and, through Equation (22), a solution for the equilibrium subjective forecast error $z_t = r(y_{t-1}, u_t)$. Importantly, under \mathcal{P} agents think that z_t is an unpredictable white noise process, while in equilibrium it is a function of the states and the exogenous shocks. Finally, the actual policy function describing the equilibrium is computed as $y_t = g(y_{t-1}, u_t) = h(y_{t-1}, u_t, r(y_{t-1}, u_t))$.

The appendix spells out this procedure in detail. I approximate the equilibrium both under learning and under rational expectations using a second order perturbation method.

4 Inspecting the mechanism

In this section, I examine a special case of the model which does away with nominal rigidities, adjustment costs, net worth dynamics and risk aversion. This special case reduces the number of state variables and allows me to solve the model in closed form.

One insight is that neither learning nor financial frictions alone generate sizable amplification of business cycle shocks or asset price volatility in a production economy. It is the interaction of these two features that leads to endogenous amplification.

4.1 Simplifying the model

First, I render the nominal rigidities redundant by setting $\kappa = \kappa_w = 0$ and $\sigma = \sigma_w = 0$. Investment adjustment costs are eliminated by setting $\psi = 0$.

Next, I simplify the financial structure of the firm. I set the exit rate of firms to $\gamma = 0$, so that firms are infinitely-lived. To ensure that they do not escape the borrowing constraint, I then require

¹⁰Dividends are paid in units of consumption goods, and so it is necessary to remove the market clearing condition for consumption goods in addition to that for stocks. Otherwise agents could infer the equilibrium amount of stocks traded from the equilibrium amount of dividends paid.

them to pay out their earnings entirely every period, by setting $\zeta = 1$. The equations (18)–(20) describing aggregate firm dynamics now simplify to:

$$N_t = R_t^k K_{t-1} - R_{t-1} B_{t-1} \quad (26)$$

$$K_t = B_t \quad (27)$$

$$B_t = \xi x K_t + \xi (1 - x) (P_t + B_t). \quad (28)$$

I further simplify preferences by setting the household coefficient of relative risk aversion to $\gamma = 0$ and the inverse elasticity of labor supply to $\phi \rightarrow \infty$, implying inelastic labor supply $L_t = \bar{L}$. I also set the autoregressive coefficient for the productivity process to $\rho = 1$, effectively making all productivity innovations permanent.

4.2 Rational expectations equilibrium

The presence of financial frictions does not necessarily imply strong amplification, and this is true in particular in this model under rational expectations. Start first with the case $\xi = 1$, where the borrowing constraint (28) is never binding. In this case, the model is just a particular variant of the RBC model. Equilibrium in the labor market requires equalization of the marginal rate of substitution and the marginal rate of transformation between consumption and labor:

$$w_t = (1 - \alpha) A_t^{1-\alpha} \left(\frac{K_t}{\bar{L}} \right)^\alpha \quad (29)$$

which implies that the expected marginal return on investment can be written as:

$$\mathbb{E}_t [R_{t+1}^k] = R^k(K_t, A_t) = \alpha \left(A_t \bar{L} e^{\frac{\alpha \sigma^2 A}{2}} \right)^\alpha K_t^{-\alpha} + 1 - \delta$$

The firm equates this expected return with the interest rate, which is constant at $R = 1/\beta$. Capital is simply proportional to productivity: $K_t/A_t = K^*$ for some fixed value K^* .

Once we introduce financial frictions by setting $\xi < 1$, how much amplification do we get? The answer: none. For all values of ξ strictly below one, the borrowing constraint is always binding, but the capital stock and the stock price of the firm are still proportional to productivity. The equilibrium is characterized by the following two equations:

$$P_t = A_t \bar{P} = A_t \frac{\beta (R^k(\bar{K}, 1) - R) \bar{K}}{1 - \beta} \quad (30)$$

$$K_t = A_t \bar{K} = \frac{\xi x}{1 - \xi} P_t \quad (31)$$

The first equation pins down the stock market value of the firm, which depends on the capital stock through expected dividends in the numerator. These dividends depend on capital through the size of the firm and the rate of return on capital. The second equation determines the capital stock that

can be reached by exhausting the borrowing constraint that depends on the stock market value. In the unique equilibrium, the capital stock is proportional to productivity, just as was the case when $\xi = 1$.

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 (or x). Similarly, the behavior of asset prices is entirely independent of financial frictions. The stock price evolves simply as:

$$\log P_t = \log P_{t-1} + \varepsilon_t. \quad (32)$$

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.

The complete irrelevance of financial frictions for the model dynamics is particular to the assumptions in this section, which make prices in the collateral constraint move exactly in lockstep with the marginal product of capital. However, the problem that low endogenous asset price volatility leads to low amplification through financial frictions is illustrative of most models with financial frictions (Quadrini, 2011).

4.3 Learning equilibrium

I now describe the equilibrium under learning. The first-order conditions of the household imply that the interest rate still has to equal $R = 1/\beta$ regardless of expectations. Also, the static labor demand equation of the firm is unchanged, so that labor market equilibrium implies Equation (29) has to hold. The firm's investment decision depends on the expected return on investment. Here, I make use of conditionally model-consistent expectations: Expectations under the subjective measure \mathcal{P} can be calculated using the equilibrium conditions of the model except for the market clearing condition for stocks and consumption goods. In this setting, we then have $\mathbb{E}_t^{\mathcal{P}} [R_{t+1}^k] = R^k(K_t, A_t)$, as under rational expectations. Agents in the learning equilibrium are able to accurately infer the equilibrium wage, and therefore the optimal choice of labor and the return on capital, given today's productivity and capital stock.


By the same logic, they can also accurately predict next period's dividend.¹¹ They do not accurately predict next period's stock price, instead relying on the subjective law of motion (22)–(23). Applying the expression for the equilibrium stock price (24), here we have:

$$P_t = \frac{\beta (R^k(K_t, A_t) - R) K_t}{1 - \beta \exp(\hat{\mu}_t + \frac{1}{2}\sigma_z^2)}.$$

We can infer the equilibrium realization of the subjective forecast error as $z_t = \Delta P_{t-1} - \hat{\mu}_t$. So

¹¹They do not accurately predict dividends more than one period ahead, as those depend on (misperceived) future stock prices.

Figure 2: Stock price dynamics under learning.



while under \mathcal{P} , agents think that z_t is random white noise, in equilibrium it is a function of the state variables and fundamental shocks of the model. This is precisely where rational expectations break under learning.

The learning equilibrium is summarized by the following three equations:

$$P_t = \frac{\beta (R^k (K_t, A_t) - R) K_t}{1 - \beta \exp(\hat{\mu}_t + \frac{1}{2}\sigma_z^2)} \quad (33)$$

$$K_t = \frac{\xi x}{1 - \xi} P_t \quad (34)$$

$$\hat{\mu}_{t+1} = \hat{\mu}_t - \frac{\sigma_\nu^2}{2} + g(\Delta \log P_t - \hat{\mu}_t) \quad (35)$$

The first equation is the determination of the equilibrium stock price as a function of the capital stock and the belief of stock price growth. The second equation is the borrowing constraint, which is always binding in equilibrium for any $\xi < 1$ ¹² The third equation is the updating equation of the Kalman filtering problem. Here, the forecast error z_t that agents perceive as random noise has been substituted out by its equilibrium value.

Figure 2 depicts the dynamics of stock prices after a positive productivity innovation $\varepsilon_1 > 0$. The initial shock at $t = 1$ raises stock prices and the capital stock proportionally to productivity through Equations (33) and (34), just as under rational expectations. Learning investors observe the rise in P_1 and are unsure whether their positive forecast error z_1 it is due to a transitive shock ($\eta_1 > 0$) or a permanent increase in the growth rate of stock prices ($\nu_1 > 0$). They therefore revise their beliefs $\hat{\mu}_2$ upward in Equation (35). In the next period $t = 2$, the more optimistic beliefs increase the demand for stocks, and the market clearing price in Equation (33) has to be higher, in turn relaxing credit constraints and fueling investment. Beliefs continue to rise in subsequent

¹²To see this, assume that P_t were sufficiently high so that the borrowing constraint didn't bind. Then $R^k(K_t, A_t) = R$ and therefore $P_t = 0$, a contradiction.

periods 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 subjective 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. But in the absence of subsequent shocks, no change in $\hat{\mu}_t$ implies no change in the price P_t , so that realized asset price growth is zero at $t = 4$, at a time when agents expect strongly positive price growth. This triggers a downward revision in beliefs and an endogenous reversal in prices. Ultimately prices return to their steady-state levels. These learning dynamics produce return volatility and predictability.


Asset price learning affects economic activity because it influences current stock prices and therefore the tightness of the borrowing constraint. But the feedback in this model is two-sided, as real activity affects stock prices through dividend payments. The strength of this channel depends on general equilibrium effects. Equation (33) shows that K_t enters the expression for expected dividends twice. The multiplying factor K_t captures a partial equilibrium effect that is internalized by the firm. The internal rate of return on capital is higher than the cost of debt and the firm therefore wants to increase its capital stock until it exhausts the borrowing constraint, *increasing* expected dividends. At the same time though, higher levels of capital lower its marginal return $R^k(K_t, A_t)$ because of decreasing returns to scale at the aggregate level, increasing wages and *decreasing* expected dividends.

When financial frictions are severe enough (ξ is low), the difference between the return on capital R^k and the return on debt R is high, and the partial equilibrium effect dominates: An expansion of firm investment increases expected dividends. This case is depicted in Panel (a) of Figure 3. The figure plots the stock pricing equation (33) and the credit constraint (34). When the degree of financial frictions is high, the credit constraint line is steep. Consider the effect of a positive productivity shock at $t = 1$ as before, when the initial equilibrium is at P_1 and $\hat{\mu}_1$. The immediate effect will be a proportionate rise in stock prices and capital, together with a rise in beliefs from $\hat{\mu}_1$ to $\hat{\mu}_2$. This leads to higher stock prices at $t = 2$ and allows the firm to invest more and increase its expected profits—the partial equilibrium 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.

When financial frictions are not severe (ξ is high), the difference between R^k and R is low and the general equilibrium effect dominates: An expansion of firm investment decreases expected dividends, as in Panel (b). A relaxation of the borrowing constraint due to a rise in $\hat{\mu}_2$ still allows the firm to invest and produce more, but in general equilibrium wages rise by so much that dividends fall. This response of dividends dampens rather than amplifies the dynamics of investment and asset prices.

This dampening effect can be so strong as to eliminate the effects of learning altogether. The appendix shows that in the limit of vanishing financial frictions, stock prices become a pure random

Figure 3: Endogenous response of dividends.



walk again:

$$\log P_t - \log P_{t-1} \xrightarrow{\xi \rightarrow 1} \varepsilon_t. \quad (36)$$

As a consequence, the entire dynamics of the model become identical to those under rational expectations. As financial frictions disappear, the general equilibrium effects offset any dynamics from stock price learning. This shows that sizable amplification arises neither from learning nor from credit frictions alone, but only from their interaction.

5 Quantitative results

The general model is solved using a second-order approximation around the non-stochastic steady state (described in the appendix). First, I discuss the parameterization. I then review standard business cycle statistics. Learning and asset price volatility account for a third of the volatility of output, pointing to the strength of the endogenous amplification mechanism. Next, I look at asset pricing moments and find that the model with learning closely matches the volatility of stock prices (which is targeted by the estimation), but also the predictability of stock returns, skewness and kurtosis. Impulse response functions confirm the presence of a strong amplification mechanism. The main channel is the endogenous volatility of asset prices induced by learning. Finally, I show that stronger financial friction amplify the effects of learning, confirming the intuition from the simplified version of the model above.

5.1 Choice of parameters

I partition the set of parameters into two groups. The first set of parameters is calibrated, and the second set is estimated by simulated method of moments (SMM) on moments of US quarterly data (1962Q1–2012Q4).

5.1.1 Calibration

The capital share in production is set to $\alpha = 0.33$. The depreciation rate $\delta = 0.025$ corresponds to ten percent annual depreciation. The trend growth rate of productivity is set to 1.5 percent annually so that $g = 0.0037$, and the persistence of productivity shocks is set to $\rho = 0.95$.¹³

The household discount factor is set such that the steady-state real interest rate equals two percent per year, implying a discount factor $\beta = 0.9951$. 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 household labor supply is set to 3, implying $\phi = 0.33$, and risk aversion is set to $\theta = 1$ (log utility in consumption).

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

Three parameters describe the structure of financial constraints: x , the probability of restructuring after default; ξ , the tightness of the borrowing constraint; and γ , the rate of firm exit. 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.¹⁴ The remaining two parameters are chosen such that the non-stochastic steady state of the model jointly matches the US average investment share in output of 18 percent and an average ratio of debt to assets of one (the sample average in the Fed flow of funds). The corresponding parameter values are $\gamma = 0.0155$ and $\xi = 0.3094$.

5.1.2 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 pertaining to both business cycle and asset price statistics: The standard deviation of output; the standard deviations of consumption, investment hours worked, and dividends relative to output;

¹³The results are similar if one assumes permanent shocks to productivity, $\rho = 1$, which is more common in the asset pricing literature.

¹⁴2006 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>. Proprietary data on Chapter 11 outcomes are analyzed in various samples by Flynn and Crewson (2009), Warren and Westbrook (2009), Lawton (2012), and Altman (2014). The results in this paper are robust to choosing alternative values as long as x is not too low. For values lower than about $x = 0.05$, the link between stock market value and investment through the borrowing constraint is too weak to generate sizeable amplification through learning.

Table 1: Estimated parameters.

parameter	σ_a	σ_i	κ	κ_w	ψ	ζ	g
learning	.00726 (.00112)	.000101 (.00210)	.362 (.137)	.956 (.024)	17.22 (2.44)	.684 (.013)	.00483 (.00005)
RE	.01029 (.00242)	.001028 (.000302)	.704 (.065)	.779 (.230)	.451 (.765)	.251 (.407)	-

Parameters as estimated by simulated method of moments. Asymptotic standard errors in parentheses. Targeted data moments and estimated standard errors in Tables 2 and 3.

and the standard deviations of inflation, the nominal interest rate, and stock returns (see Tables 2 and 3 for the value of the data moments and estimated standard errors; all variables are in logs and all variables except stock returns are HP-filtered). The set of estimated parameters θ solves

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

where $m(\theta)$ 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.¹⁵ 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.¹⁶ Such oscillations are not observed in the data, but can emerge in the learning equilibrium when the learning gain is large. The restriction effectively constrains the degree of departure of subjective beliefs from rational expectations.

Table 1 summarizes the SMM estimates for both the learning and rational expectations version of the model. 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 second row contains the parameters estimated under rational expectations. The size of the shocks σ_a and σ_i is larger than under learning, despite the fact that the degree of investment adjustment costs is much smaller, implying a larger degree of endogenous amplification of the shocks under learning.

The Calvo price adjustment parameter κ implies retailers adjust their prices about every 2 months on average, while they do so about every 7 months in the rational expectations estimation. By contrast, the SMM procedure selects a high degree of nominal wage rigidities κ_w and of adjustment costs ψ under learning. These values are somewhat undesirable, but it is intuitive why they are needed to fit the data. A high degree of wage rigidity is needed in order to match the volatility of employment, which is about as high as the volatility of output. High wage rigidities help to match it by suppressing the wealth effect on labor supply that would otherwise dampen the rise of employment during asset price booms. High adjustment costs are needed because the high asset price volatility would otherwise translate into counterfactually large fluctuations in investment

¹⁵I choose $W = \text{diag}(\hat{\Sigma})^{-1}$ 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.

¹⁶ $\theta \notin \mathcal{A}$ if 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.

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

	moment	(1) data	(2) learning	(3) RE	(4) no fin. fric.	(5) RE re-estimated
output	$\sigma_{hp}(Y_t)$	1.43%	1.51%*	0.82	0.65	1.66%*
volatility		(0.14%)				
volatility rel. to output	$\sigma_{hp}(I_t) / \sigma_{hp}(Y_t)$	2.90 (.12)	2.99*	0.41	0.25	2.71*
	$\sigma_{hp}(C_t) / \sigma_{hp}(Y_t)$.60 (.035)	.58*	1.00	1.36	.60*
	$\sigma_{hp}(L_t) / \sigma_{hp}(Y_t)$	1.13 (.061)	1.11*	0.45	0.25	1.24*
	$\sigma_{hp}(D_t) / \sigma_{hp}(Y_t)$	3.00 (.489)	2.97*	2.21	-	1.41*
correlation with output	$\rho_{hp}(I_t, Y_t)$.95 (.0087)	.75	0.91	0.24	0.96
	$\rho_{hp}(C_t, Y_t)$.94 (.0087)	.81	0.97	0.99	0.80
	$\rho_{hp}(L_t, Y_t)$.85 (.035)	.93	0.83	0.30	0.84
	$\rho_{hp}(D_t, Y_t)$.56 (.080)	.60	0.51	-	0.45
inflation	$\sigma_{hp}(\pi_t)$.27% (.047%)	.31%*	0.28%	0.31%	.21%*
nominal rate	$\sigma_{hp}(i_t)$.37% (.046%)	.10%*	0.10%	0.12%	.09%*

Quarterly U.S. data 1962Q1–2012Q4. Standard errors in parentheses. π_t is quarterly CPI inflation. i_t is the federal funds rate. All following variables are in logarithms. 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. Dividends D_t are four-quarter moving averages of S&P 500 dividends. $\sigma_{hp}(\cdot)$ 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.

relative to output. In a sense, the amplification through asset price learning is so powerful that it needs to be dampened again to fit the data.

The fraction of earnings paid out as dividends ζ under learning is fitted to 68 percent, somewhat higher than the historical average for the S&P500 (about 50 percent). Finally, the learning gain g implies that agents believe the amount of predictability in stock price growth to be small: When stock prices today rise by 10 percent more than they expected, they update their belief about predictable future price growth by 0.05 percent.

5.2 Business cycle and asset price moments

I now review the key business cycle and asset pricing moments in the data and across model specifications. Table 2 starts with business cycle statistics. Moments in the data are shown in Column (1). Moments for the estimated learning model are shown in Column (2), while Columns

(3) and (4) contain the corresponding moments for the model under rational expectations and a comparison rational expectations model without financial frictions.¹⁷ The model parameters are held constant at the estimated values for the learning model in Columns (2) to through (4). 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 detrended output. By construction, it is matched well by the learning model in Column (2). When learning is shut off in Column (3), it drops almost one half. This shows the great degree of amplification that learning adds to the model. 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. But it is not as powerful as when it is combined with asset price learning. Of course, it is also possible to match output volatility with rational expectations, using larger shock sizes, as in Column (5).

The next set of rows report the standard deviation of consumption, investment, hours worked and dividends relative to output. Moving from Column (2) to (3), the removal of learning leads to a drop in the relative volatility of 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 volatility of dividends is well matched by the learning model, which is a prerequisite for stock prices in the model and in the data to be comparable.

Next, I report the contemporaneous correlations of consumption, investment, hours worked and dividends with output. The values in the model with learning are broadly in line with the data. The last two rows report the volatility of inflation and the nominal interest rate. Inflation volatility is also roughly in line with the data, but the nominal interest rate is less volatile across all model specifications.

Next, I present asset price statistics in Table 3.¹⁸ The statistics correspond to some well-known asset price puzzles. In an endowment economy, Adam et al. (2015) show that a learning model is able to match these statistics well, and the table makes it clear that this carries over into a general equilibrium model with production.¹⁹ Starting with excess volatility in Column (2), the model with learning reproduces the standard deviation of excess returns in the data and also comes close to the volatility of the P/D ratio.²⁰ By contrast, the model re-estimated under rational expectations

¹⁷In the economy without financial frictions, intermediate goods producers are owned directly by households and face no financial constraint. The model then reduces to a standard New-Keynesian model with adjustment costs and price and wage rigidities. The stock price in this economy is defined as the value of the capital stock, $Q_t K_t$.

¹⁸I omit a comparison to the frictionless economy benchmark as presented in Table 2 as the asset price statistics are very similar to those in Column (3).

¹⁹It is also worth noting that these statistics are produced using only a second-order approximation to the model equations, whereas Adam et al. (2015) and Adam et al. (2016) use global methods to solve their model.

²⁰The P/D ratio is not as high as in the data is because the P/D ratio correlates somewhat too strongly with dividends. The Campbell-Cochrane decomposition $r_t \approx \kappa_0 + \kappa_1 p d_t - p d_{t-1} + \Delta d_t$ implies that $\mathbb{V}[r_t] \approx (1 - 2\kappa_1 \rho(p d_t, p d_{t-1}) + \kappa_1^2) \mathbb{V}[p d_t] + \mathbb{V}[\Delta d_t] + 2\text{Cov}(\Delta d_t, \Delta p d_t)$. Since the learning model matches both the volatility of dividends and returns, and has an autocorrelation of $p d_t$ that is low relative to the data, the low volatility of $p d_t$ results from the covariance term of price and dividend growth.

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

		(1) data	(2) learning	(3) RE	(4) no fin. fric.	(5) RE re-estimated
excess	$\sigma(R_{t,t+1}^e)$	32.56%	33.99%*	0.61%	2.02%	0.97%*
volatility	$\sigma\left(\frac{P_t}{D_t}\right)$	(2.44%) 41.08% (6.11%)	32.62%	2.42%	2.48%	3.86%
return	$\rho\left(\frac{P_t}{D_t}, R_{t,t+4}^e\right)$	-.297 (.092)	-.486	-.14	.05	-.28
predictability	$\rho\left(\frac{P_t}{D_t}, R_{t,t+20}^e\right)$	-.585 (.132)	-.752	-.09	.43	-.17
	$\rho\left(\frac{P_t}{D_t}, \frac{P_{t+4}}{D_{t+4}}\right)$.904 (.056)	.54	.19	.32	.58
negative	skew($R_{t,t+1}^e$)	-.897 (.154)	-.188	-.031	.003	-.008
skewness	kurt($R_{t,t+1}^e$)	1.57 (.62)	3.33	0.07	0.00	-0.05
heavy tails						
risk-free rate	$\mathbb{E}(R_t^f)$	1.99% (.61%)	1.99%	1.99%	1.99%	1.99%
	$\sigma(R_t^f)$	2.34% (.29%)	0.74%	0.58%	.62%	0.86%
equity	$\mathbb{E}(R_{t,t+1}^e)$	4.06% (1.93%)	0.01%	0.00%	0.00%	-0.04%
premium						
price	$\rho_{hp}(P_t, Y_t)$.458 (.115)	.700	.979	.996	.748
correlation						
with output						

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. Excess returns R_t^e are annualized quarterly excess returns of the S&P 500 over 3-month Treasury yields. $\sigma(\cdot)$ is the standard deviation; $\rho(\cdot, \cdot)$ is the correlation coefficient; $\rho_{hp}(\cdot, \cdot)$ is the correlation coefficient of HP-filtered data (smoothing coefficient 1600); skew(\cdot) is skewness; kurt(\cdot) is excess kurtosis. Moments used in the SMM estimation are marked with an asterisk.

in Column (5) cannot produce a similar amount of volatility, despite the fact that return volatility is explicitly targeted by the SMM estimation.²¹

Stock returns also exhibit considerable predictability by the P/D ratio at business-cycle frequency. The same is true in the model with learning. Predictability is not targeted by the estimation, and in fact it is somewhat stronger than in the data, reflected in a persistence of the P/D ratio that is somewhat lower than in the data. Again, the rational expectations model is not able to produce sizable return predictability.

The learning model also produces a distribution of returns that is negatively skewed and heavy-

²¹The poor asset pricing performance of the rational expectations model is due to the very simple preferences and moderate risk aversion. Other models exist that can address asset pricing puzzles in a production economy at least as well, e.g. for habit and for long-run risk. However, these alternative ways of modelling asset prices would miss the disconnect of survey expectations from statistically expected stock returns discussed in the introduction and later in Section 6.


tailed as in the data, underlining the non-linearities in the asset price dynamics that arise from the learning mechanism. At the same time, the model delivers a low and smooth risk-free rate. Even though realized stock returns are very volatile, expected stock returns are not very volatile because discount factors are fairly stable under learning. However, the learning model is not able to produce a sizable average equity premium. The reason is that, even though returns are highly volatile, this volatility is not priced because it stems from the subjective updating of beliefs, and this problem has been encountered previously in the learning literature ([Barberis et al., 2015](#)).

Finally, the model has a rather large degree of comovement between stock prices and real activity. While stock prices are procyclical in the data, too, the correlation is only about 0.45, while in the learning model it is 0.70. One could think that this high correlation arises because of the link between stock prices and investment in the borrowing constraint, but in fact the correlation is even higher in the rational expectations economy without financial frictions in Column (3). Rather, the high correlation arises because the model is driven only by two business cycle shocks (productivity and monetary policy) that simultaneously move output and firm value. In the data, many movements in stock prices occur without immediate changes in fundamentals, and the model in this paper is not set up to capture those movements. Adding news shocks or discount factor shocks could bring the correlation in line with the data, but is beyond the scope of this paper.

5.3 Impulse response functions

Impulse response functions reveal 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 non-linearities present with 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 the frictionless comparison; this is the standard financial accelerator effect. But the strength of the financial accelerator is magnified considerably under learning. 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 7.1). 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 price-dividend ratio (not shown) also rises after the shock. The nominal interest rate falls less under learning as the monetary authority reacts to the inflationary pressures stemming from the relaxation in credit

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.

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.

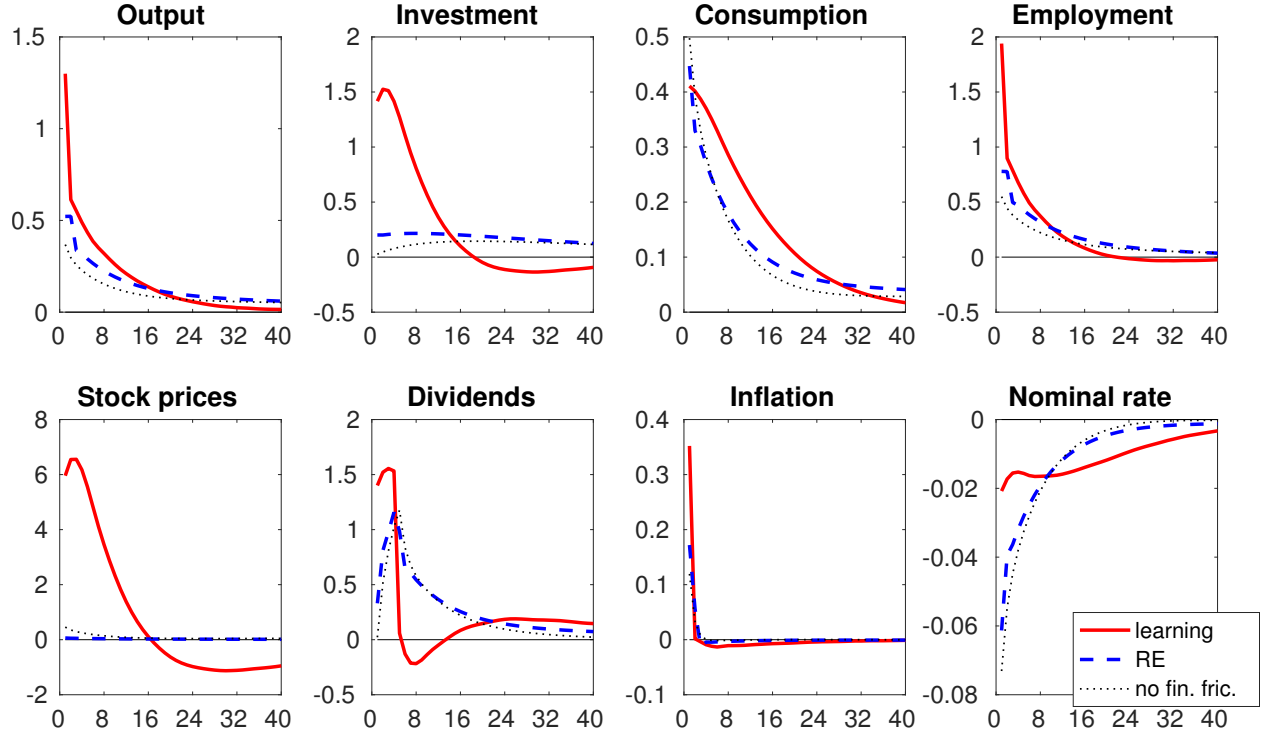
5.4 Interaction between learning and financial frictions

The analysis of the simplified model in Section 4 revealed that there can be an important interaction between learning and the financial friction. More precisely, in the simplified model the effects of learning are amplified by a tighter financial constraint, and conversely disappear in the limit when the constraint stops binding. Here, I show that a similar interaction is also present in the full model. Figure 6 plots the volatility of output and stock prices for different values of ξ , which parameterizes the overall tightness of the borrowing constraint in Equation (8).

As can be seen in the upper panel of the figure, the volatility of stock prices is decreasing in the tightness of the financial friction, just as in Section 6.²² A tighter constraint implies a higher wedge between the return on capital and the return on debt, which means that a temporary relaxation

²²Unlike in the simplified version of the model, the limit of the financial friction not binding cannot be reached here. For values of ξ higher than those shown in the figure, the borrowing constraint is still binding in steady state but the equilibrium becomes locally unstable.

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.


of the borrowing constraint due to changes in beliefs under learning increases the firm's expected dividend payouts and raises stock prices by more, strengthening the positive feedback loop between stock prices and real activity in the model.

The lower panel of the figure shows that a tighter credit constraint also increases output volatility under learning, unless ξ is very small. When ξ is near zero, almost no collateral is pledgeable and firms have to finance their capital stock almost entirely out of retained earnings. In this case, fluctuations in stock prices only have small effects on allocations, which brings the dynamics of output under learning and rational expectations closer together (indeed, they coincide perfectly at $\xi = 0$). But otherwise, the amount of amplification from learning is decreasing with ξ , just like in the simplified version of the model.

6 Survey data 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. If it were, then agents would quickly detect this predictability and improve their forecasts accordingly. Yet the

Figure 6: Interaction between the financial friction and learning.



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.


absence of predictability is easily rejected in the data, which is a longstanding challenge for rational expectations.

By contrast, agents in this model also make systematic, predictable forecast errors. This is true not only for stock prices but also for other endogenous model variables, despite the fact that, *conditional* on stock prices, agents' beliefs are model-consistent. A systematic mistake in predicting stock prices will 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.

Despite the parsimonious departure from rational expectations, the model fits several patterns of forecast predictability in the data. Since survey data can be used to put quantitative discipline on models of expectation formation (Manski, 2004), the evidence presented here favors the use of learning models such as this one over rational expectations models.

Figure 7 repeats the scatter plot of the introduction, 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 while realized returns will be low on average. This deviation of beliefs from rational expectations is easy to detect even in a small sample like the one generated in Figure 7, yet it is exactly this detectable pattern that is consistent with survey data. In the learning model, the discrepancy between expected returns and return expectations has a causal interpretation: High return expectations drive up stock prices. At the same time, realized future

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



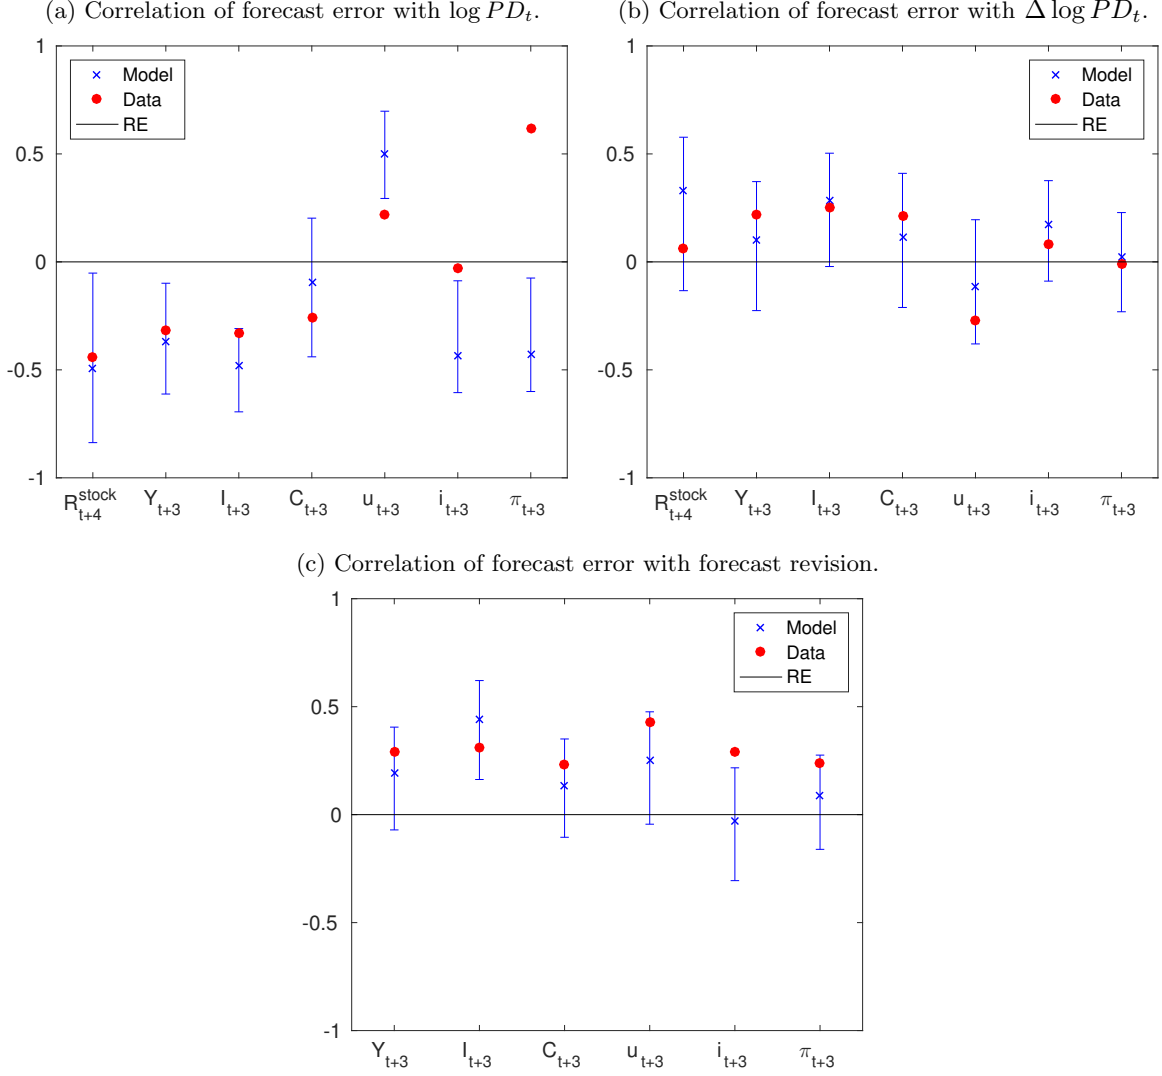
Expected and realized nominal returns along a simulated path of model with learning. Simulation length 200 periods. Theoretical correlation coefficient for subjective expected returns $\rho = .47$, for future realized returns $\rho = -.38$.

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.

Figure 8 compares the predictability of forecast error in the Federal Reserve’s Survey of Professional Forecasters (SPF) as well as the CFO survey to that obtained from simulated model data. Each red dot 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, all correlation coefficients should be zero. The blue crosses show the corresponding correlation coefficient in the model (obtained using a long simulation of 50,000 periods), while the blue bands represent 95% confidence intervals of the model correlation coefficient in small samples of the same length as those in the data. The model-data comparison is also tabulated in Table 4. The table additionally shows the p-values from the corresponding regression of the forecast errors on the predictors. It is immediate that the null of no predictability is often rejected in the data.

Panel (a) of the figure and Columns (1) and (2) of the table contain the correlations of future forecast errors with the P/D ratio. When stock prices are high, people systematically over-predict future stock returns and economic activity. For stock returns the model reproduces this pattern, as was already shown in the scatter plot above. But the model also reproduces this pattern for output, investment, consumption and unemployment (proxied for by one minus employment in the model). Where the model fails to replicate the data is for inflation: When stock prices are high, forecasters under-predict future inflation, while in the model the opposite result obtains. This discrepancy arises essentially because stock market booms are inflationary in the model, which seems not to be

Figure 8: Forecast error predictability.



Red dots show 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, unemployment rate, CPI inflation and 3-month treasury bill (SPF). Regressors: Panel (a) is the S&P 500 P/D ratio and Panel (b) is its first difference. Panel (c) is the forecast revision as in [Coibion and Gorodnichenko \(2015\)](#). Data from Graham-Harvey covers 2000Q3–2012Q4. Data for the SPF covers 1981Q1–2012Q4. Blue crosses show corresponding correlation coefficients in the model, 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 $u_t = 1 - L_t$. Stock returns in the model $R_{t,t+4}^{stock}$ are quarterly nominal aggregate market returns. Blue lines show 95% confidence bands of the correlation coefficients in the model in small samples of the same size as the data (123 quarters in the SPF and 49 quarters in the Graham-Harvey survey) from 5,000 simulations with a burn-in period of 1,000 periods.

true in the data (Christiano et al., 2010).

Panel (b) of the figure and Columns (3) and (4) of the table repeat the exercise for the first difference of the P/D ratio, and here the model replicates the patterns in the data across all variables remarkably well. Price-dividend ratio growth positively predicts forecast errors, suggesting that agents' expectations adjust slowly: They under-predict an expansion in its beginning but then overshoot and over-predict it when it is about to end. In the model, this pattern also emerges because expectations about asset prices adjust slowly.

Finally, Panel (b) of the figure and Columns (5) and (6) report 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 $\hat{E}_t[x_{t+3} - x_t] - \hat{E}_{t-1}[x_{t+3} - x_t]$. Forecast errors are positively predicted by this revision measure across all variables. Coibion and Gorodnichenko take this as evidence for sticky information models in which information sets are gradually updated over time, but concede that "deviations from FIRE [full information rational expectations] may exist above and beyond those captured by simple models of information rigidities" (p. 2655). It turns out that a learning model can capture the predictability of forecast errors by forecast revisions as well, and furthermore also be consistent with additional patterns of predictability in the data.

7 Discussion

7.1 Does learning matter?

In this model, large swings in stock prices lead to large swings in real activity through their effect on credit constraints. But maybe all that matters for amplification is asset price volatility, and learning is but one way of getting there. Here, I show that learning does have create amplification *beyond* its effect on asset price volatility. Extrapolative expectations about asset prices also cause procyclical movements in aggregate demand that do not obtain under rational expectations, leading to additional amplification in the presence of nominal rigidities.

To this end, I replace the stock market value P_t in the borrowing constraint (20) with an exogenous process V_t that has the same law of motion as the stock price under learning. More precisely, I fit a linear 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 identical dynamics to the model under learning.²³

Figure 9 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 practically indistinguishable. But after a positive productivity shock, output, investment, and consumption

²³For this exercise I only compute a first-order approximation to the model equations, so that an ARMA process has a fair chance of fitting the learning dynamics.

Table 4: Forecast error predictability, tabulated.

	(1)	(2)	(3)	(4)	(5)	(6)
	log PD _t		Δ log PD _t		forecast revision	
forecast variable	data	model	data	model	data	model
<i>R</i> _{<i>t,t+4</i>} ^{stock}	-.44 [.000]	-.49	.06 [.566]	.33	-	-.47
<i>Y</i> _{<i>t,t+3</i>}	-.32 [.021]	-.37	.22 [.068]	.10	.29 [.004]	.20
<i>I</i> _{<i>t,t+3</i>}	-.33 [.060]	-.48	.25 [.028]	.29	.31 [.001]	.44
<i>C</i> _{<i>t,t+3</i>}	-.26 [.062]	-.09	.21 [0.057]	.12	.23 [.103]	.13
<i>u</i> _{<i>t,t+3</i>}	.22 [.135]	.50	-.27 [.043]	-.12	.43 [.000]	.25
<i>π</i> _{<i>t,t+3</i>}	.62 [.000]	-.43	-.01 [.910]	.02	.24 [.008]	.08
<i>i</i> _{<i>t+3</i>}	-.03 [.874]	-.40	.08 [.399]	.27	.29 [.002]	-.05


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, unemployment rate, CPI inflation and 3-month treasury bill (SPF). Newey-West adjusted p-values from least squares regressions in brackets. 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 $u_t = 1 - L_t$. Stock returns in the model $R_{t,t+4}^{stock}$ are quarterly nominal aggregate market returns.

still increase by 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.

7.2 Do non-stationary beliefs matter?

This section considers a generalized belief system which allows agents' subjective belief about stock prices to be mean-reverting. I show that expectations about the expected long-run dynamics of stock prices do not affect the equilibrium dynamics, as long as expectations about price growth remain persistent in the short run.

Figure 9: Effects of learning beyond asset prices.



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 fitted to those under learning by a linear ARMA(10,5) process. All impulse responses in the figure are produced using a first-order approximation to the model equations.

Consider extending the belief system (22)–(23) as follows:

$$\log P_t = \log P_{t-1} + \hat{\mu}_{t-1} + z_t \quad (37)$$

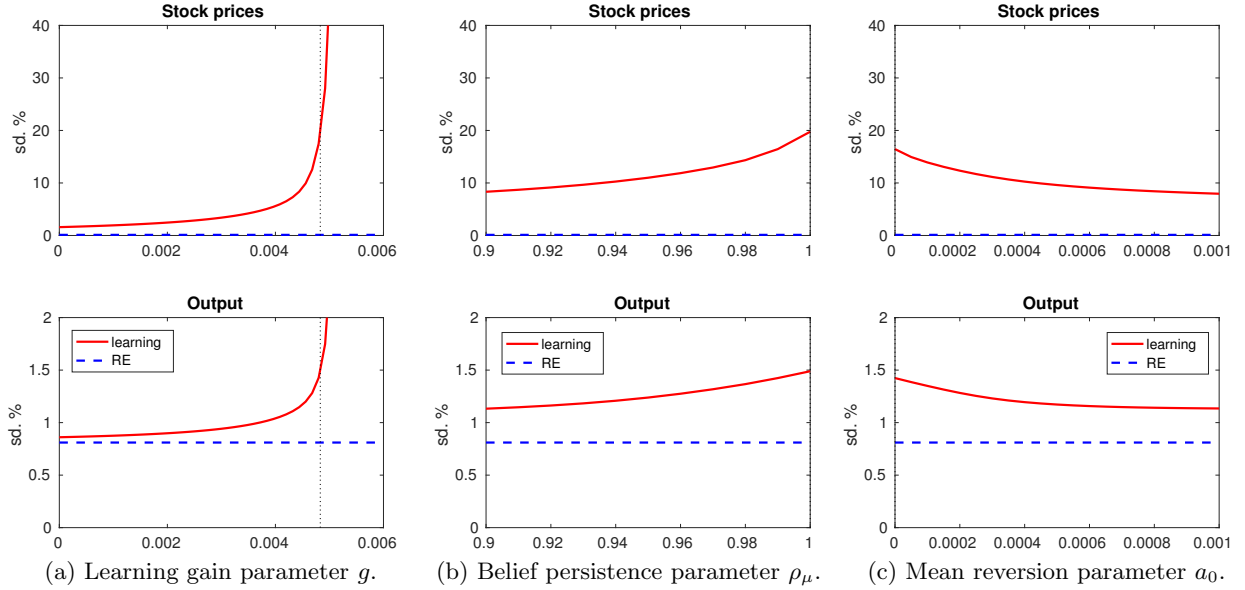
$$\hat{\mu}_t = \rho_\mu \hat{\mu}_{t-1} - a_0 \log \frac{P_{t-1}}{\bar{P}} + g z_t. \quad (38)$$

This specification nests the simpler belief system when $\rho_\mu = 1$ and $a_0 = 0$. Now if $\rho_\mu < 1$ and $a_0 = 0$, agents believe stock price growth to be mean-reverting, although they still believe the level of stock prices to follow a random walk. If $\rho_\mu < 1$ and $a_0 > 0$, then agents' also believe that the level of stock prices is mean reverting to some long-run value \bar{P} . Figure 10 plots the volatilities of output and stock prices as functions of each of the three parameters of this extended belief system.

The left panel of the figure shows the sensitivity to the learning gain when $\rho_\mu = 1$ and $a_0 = 0$. This parameter governs the speed with which expectations adjust to recent observations of price growth. Asset price volatility as well as the amount of amplification obtained under learning are very sensitive to this parameter. In the estimated learning model, it is set to $g = 0.00483$ in order to match the return volatility observed in the data.

In the middle panel, the learning gain is fixed at its estimated value, $a_0 = 0$ and ρ_μ is being varied. The strength of the effects of learning are increasing in the persistence of the perceived

Figure 10: Sensitivity to the specification of the belief system.



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

predictable component of prices. However, the value $\rho_\mu = 1$ is not special. Even if agents believe that price growth eventually returns to a stable mean in the long run, the effects of learning remain sizable: At $\rho_\mu = 0.9$ (an expected yearly mean reversion of price growth by about 35 percent), the volatility of stock prices under learning is still about eight times larger than under rational expectations.

In the right panel, I set $\rho_\mu = 0.99$ and vary the value for a_0 , the strength of perceived mean reversion in levels. The effects of learning decrease but remain sizable as a_0 increases. Again, the value $a_0 = 0$ is not special. Whether agents believe in mean reversion in the long run or not is not important as long as short-run expected price growth remains persistent.

In sum, the effects of learning do not depend on subjective long run expectations about prices and can be modified to incorporate mean reversion in the growth rate or even the level of prices. What is important though is that agents overestimate the persistence of price growth relative to the data. The fact that measured stock return expectations comove positively with past returns supports such a misperception.²⁴

7.3 Do nominal rigidities matter?

The model includes nominal rigidities. These allow for an investigation of the role of monetary policy in the next section, but they are also crucial for the quantitative fit of the model. It is difficult to obtain comovement of consumption, investment and employment in response to changes

²⁴ Adam et al., 2016 provide further evidence of the persistence of expected price growth by comparing stock price survey forecasts across multiple horizons.

in investment demand that are unrelated to TFP (Beaudry and Portier, 2007). Nominal rigidities help with this problem, as shown by Kobayashi and Nutahara (2010) in the context of news shocks and Ajello (2016) in the context of credit shocks, and they also help here.

An increase in stock prices relaxes credit constraints and allows firms to invest more and pay out higher dividends. But general equilibrium effects can dampen or even overturn these effects, as has already been discussed in the simplified model of Section 4. Higher demand for investment implies higher labor demand and therefore an increase in the real wage. It also tends to crowd out consumption today, implying higher real interest rates through the consumption Euler equation. Both channels dampen the rise in firm profits and investment.

Wage rigidities will counteract the dampening effects of real wage responses to changes in the tightness of credit constraints, allowing for a larger response of employment and dividends. Price rigidities, together with a relatively loose monetary policy rule, render the price of intermediates q_t pro-cyclical, implying larger dividend responses. They also dampen real interest rate movements in response to changes in investment demand, mirrored by less crowding out of consumption. Together, nominal rigidities help the comovement of investment, consumption and employment, and they also amplify the response of dividends to changes in borrowing constraints, reinforcing the learning dynamics in the stock market.

To illustrate this point, I re-compute impulse responses of the model with learning, but without nominal rigidities (setting $\kappa = \kappa_w = 0$). I also reduce the size of investment adjustment costs to $\psi = 0.1$. 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. Lowering adjustment costs also gives the real version a better chance at delivering strong impulse responses. Even then, the nominal version delivers greater amplification in output and consumption. Figure 11 plots impulse responses to a positive productivity shock for both the nominal and real version 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 more and dividends rise less. This considerably dampens the learning dynamics and also mutes the response of stock prices.

8 Implications for monetary policy


If belief distortions from asset price learning have important consequences for the real economy, then the question arises whether policy should intervene to mitigate them. Here, I specifically ask whether it would be beneficial for the central bank in the model to “lean against the wind”, i.e. to systematically raise interest rates in response to asset price increases.

Consider extending the interest rate rule (9) as follows:

$$i_t = \rho_i i_{t-1} + (1 - \rho_i) (1/\beta + \phi_\pi \pi_t + \phi_Y \Delta \log Y_t + \phi_P \Delta \log P_t) \quad (39)$$

In addition to raising interest rates when inflation is above its target level (taken to be zero), the monetary authority can raise interest rates $\phi_{\Delta Y}$ percentage points when real GDP growth increases

Figure 11: Counterfactual impulse responses without 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 5.1.2 to fit the data (“real” comparison). The size of the shock shown is the same as in the nominal model.

one percentage point and $\phi_{\Delta P}$ percentage points when stock price growth increases one percentage point. Conditionally model-consistent expectations imply that under learning, agents have full knowledge of the policy rule as well as the monetary transmission mechanism, but continue to believe in the misspecified law of motion for stock prices.

Table 5 shows numerically computed policy rules that maximize expected conditional welfare at the non-stochastic steady state. Welfare calculations in the model are complicated by the fact that there are two types of households in the model. I compute²⁵ a weighted average of household and firm owner utility, expressed in units of steady state consumption and weighed by their share in aggregate consumption:

$$\chi = \frac{\bar{C}_{welfare} + \bar{C}_{welfare}^f}{\bar{C}_{SS} + \bar{C}_{SS}^f} - 1 \quad (40)$$

²⁵To compute conditional welfare, I run 4,000 model simulations of length 1,000 periods, each starting at the non-stochastic steady state, and compute series for C_t , C_t^f and L_t using the exact formulas given in the appendix. I then evaluate the discounted sum of utility in each simulation and compute conditional welfare as the average across simulations.

Table 5: Optimized interest rate 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	1.36	2.84	1.50	4.82	3.59
$\phi_{\Delta Y}$		3.48	2.63		1.10	0.85
$\phi_{\Delta P}$			1.82			-1.00
$\sigma_{hp}(Y)$	1.51%	.30%	.51%	1.66%	.93%	.93%
$\sigma_{hp}(P)$	12.14%	3.12%	1.17%	.10%	.10%	.10%
$\sigma_{hp}(\pi)$.31%	.43%	.32%	.21%	.14%	.14%
$\sigma_{hp}(i)$.10%	.09%	.04%	.09%	.05%	.05%
welfare gain relative to baseline	-	1.244%	1.381%	-	.314%	.314%
welfare gain χ relative to steady state	0.084%	1.187%	1.289%	-0.092%	0.222%	0.222%

Standard deviations of output, stock prices, inflation, and interest rates (unfiltered) under learning in percent. The welfare gain is computed relative to the respective baselines in Columns (1) and (4) and expressed in percent of steady-state consumption (see Footnote 25). The interest rate smoothing coefficient is kept at $\rho_i = 0.85$ and interest rate shocks are absent in all rules considered.

where

$$u(\bar{C}_{welfare}, \bar{L}_{SS}) = (1 - \beta) \mathbb{E} \left[\sum_{t=0}^{\infty} \beta^t u(C_t, L_t) \right]$$

$$\bar{C}_{welfare}^f = (1 - \beta) \mathbb{E} \left[\sum_{t=0}^{\infty} \beta^t C_t^f \right].$$

In both the learning and the rational expectations versions of the model, steady-state consumption of the risk-averse households is about 9 times larger than that of risk-neutral firm owners, implying a welfare weight of households of about 90 percent. The results in this section are qualitatively unchanged if the conditional welfare measure is replaced by unconditional expected utility or a simple quadratic loss function of output and inflation volatility. Also, note that the expectation is evaluated under the actual equilibrium distribution of model outcomes, rather than under the subjective probability measure \mathcal{P} . Under learning, this welfare measure is therefore paternalistic, in that maximizing it does not necessarily maximize agents' subjectively expected welfare.

Column (1) shows the baseline calibration under learning. The welfare measure is slightly higher than in steady-state, owing to the fact that the steady state itself is inefficient and no mean correction has been applied to the exogenous driving processes. Column (2) displays optimized

coefficients ϕ_π and ϕ_Y without a reaction to stock prices. The resulting rule implies a large welfare gain of 1.19% of steady state consumption. Column (3) additionally allows for a reaction to stock prices. The optimal coefficient on stock prices is computed as $\phi_P = 1.82$. Allowing for a reaction to stock prices increasing the volatility of output but further reduces the volatility of asset prices and inflation. Overall, the welfare gain relative to the baseline is larger than without the reaction to stock prices, at 1.29%. It is worth noting that despite a strong reaction of interest rates to stock prices, the volatility of the nominal interest rate in Column (3) is lower than without the reaction. In the model with learning, leaning against the wind reduces endogenous asset price volatility so that equilibrium rates do not end up being excessively volatile. Intuitively, raising rates when asset prices are rising (and vice-versa) acts to stabilize expectations in financial markets. In the model, this stabilization works mainly through changes in firms' borrowing costs and dividend payouts, offsetting the self-amplifying learning dynamics in the stock market.

By contrast, under rational expectations this benefit of stabilizing financial market expectations is absent. Columns (4) to (6) repeat the discussed calculations for the rational expectations version of the model, re-estimated to fit the moments in the data. The optimal coefficient ϕ_P found by the numerical optimization is in fact negative, but to be precise, the optimal combination of coefficients in Column (6) is not uniquely determined: Higher values of ϕ_P together with lower values for ϕ_Y lead to identical allocations and welfare gains relative to the (rational expectations) baseline. In other words, the optimal allocations and welfare gains cannot be improved upon by reacting to stock prices when expectations are rational. This result is similar in spirit to [Faia and Monacelli \(2007\)](#) or [Cúrdia and Woodford \(2016\)](#) who find the benefits of reacting to financial conditions to be small.

These results on optimal policy rules come with several caveats. First, the results presented here are meant to illustrate that optimal policy prescriptions in a model depend on the underlying asset pricing theory, rather than to provide a comprehensive evaluation of the merits of leaning against the wind. Second, because risk aversion is relatively low, the welfare cost of business cycles is small to start with. This low risk aversion is also reflected in the small equity premium that the model produces. Higher risk aversion could potentially affect how agents react to changes in policy. Finally, it was implicitly assumed that the subjective model that agents use to forecast stock prices under learning is invariant to policy, which can be legitimately questioned even in the absence of fully rational expectations.

9 Conclusion

In this paper, I have analyzed the implications of learning-based asset pricing in a business cycle model with financial frictions. When firms' borrowing constraints depend on their market value, 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. At the same time, it makes the financial accelerator mechanism more powerful, amplifying standard business cycle shocks. The

model jointly matches business cycle and asset pricing moments. One interesting avenue to explore in further research is to incorporate firm heterogeneity into the model, taking into account the fact that some firms are less credit-constrained than others;

The model also replicates the predictability of forecast errors in survey data across a range of variables, as agents' forecast errors on stock prices spill over into their forecasts of economic activity. The fact that many patterns of predictability match up with the data is strong evidence in favor of the expectation formation mechanism laid out in this paper.

An innovation in developing the model was to introduce a belief system that combines learning about stock prices with a high degree of rationality and internal consistency. Beliefs are restricted in a way such that forecast errors conditional on future prices and fundamentals are zero. This differs from most of the existing adaptive learning literature where every forward-looking equation is parameterized separately. The method is general and can be applied in other models of the business cycle.

An examination of the sensitivity of the amplification mechanism to the monetary policy rule revealed that a reaction of interest rates to stock price growth is beneficial under learning. This is because such a reaction effectively stabilizes expectations in financial markets. The same is not true in a rational expectations framework, illustrating that the choice of an asset price theory can have important normative implications. Further research could examine in detail the policy implications of learning-based asset pricing and establish whether a policy of "leaning against the wind" is desirable in more general settings.

References

- Adam, K. and A. Marcet**, "Internal rationality, imperfect market knowledge and asset prices," *Journal of Economic Theory*, 2011, 146 (3), 1224–1252.
- Adam, Klaus, Albert Marcet, and Juan Pablo Nicolini**, "Stock Market Volatility and Learning," *Journal of Finance*, 2015.
- , **Johannes Beutel, and Albert Marcet**, "Stock Price Booms and Expected Capital Gains," Working paper 2016.
- Ajello, Andrea**, "Financial intermediation, investment dynamics and business cycle fluctuations," *American Economic Review*, 2016, 106 (8), 2256–2303.
- Altman, Edward I**, "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*, 2014, 22 (1), 75–267.
- Andrade, Philippe and Hervé Le Bihan**, "Inattentive professional forecasters," *Journal of Monetary Economics*, 2013, 60 (8), 967 – 982.

- Bacchetta, Philippe, Elmar Mertens, and Eric van Wincoop**, “Predictability in financial markets: What do survey expectations tell us?,” *Journal of International Money and Finance*, April 2009, 28 (3), 406–426.
- Baker, Malcolm, Jeremy C. Stein, and Jeffrey Wurgler**, “When Does the Market Matter? Stock Prices and the Investment of Equity-Dependent Firms,” *Quarterly Journal of Economics*, 2003, 118 (3), 969–1005.
- Barberis, Nicholas, Robin M Greenwood, Lawrence J Jin, and Andrei Shleifer**, “X-Capm: An Extrapolative Capital Asset Pricing Model,” *Journal of Financial Economics*, January 2015, 115 (1), 1–24.
- Barro, RJ**, “The stock market and investment,” *Review of Financial Studies*, 1990, 3 (1), 115–131.
- Beaudry, Paul and Franck Portier**, “When can changes in expectations cause business cycle fluctuations in neo-classical settings?,” *Journal of Economic Theory*, 2007, 135 (1), 458 – 477.
- Bernanke, Ben S. and Mark Gertler**, “Should Central Banks Respond to Movements in Asset Prices?,” *American Economic Review*, 2001, 91 (2), 253–257.
- , —, and **Simon Gilchrist**, “The financial accelerator in a quantitative business cycle framework,” in J. B. Taylor and M. Woodford, eds., *Handbook of Macroeconomics*, Vol. 1, Elsevier, 1999, chapter 21, pp. 1341–1393.
- Blanchard, Olivier, Changyong Rhee, and Lawrence Summers**, “The stock market, profit, and investment,” *Quarterly Journal of Economics*, 1993, 108 (1), 115–136.
- Brunnermeier, Markus K. and Yuliy Sannikov**, “A Macroeconomic Model with a Financial Sector,” *American Economic Review*, February 2014, 104 (2), 379–421.
- Caputo, Rodrigo, Juan Pablo Medina, and Claudio Soto**, “The Financial Accelerator Under Learning and The Role of Monetary Policy,” Working Papers of the Central Bank of Chile 590, Central Bank of Chile September 2010.
- Chakraborty, Indraneel, Itay Goldstein, and Andrew MacKinlay**, “Housing Price Booms and Crowding-Out Effects in Bank Lending,” Working paper 2016.
- Christiano, Lawrence, Cosmin Ilut, Roberto Motto, and Massimo Rostagno**, “Monetary policy and stock market booms,” *Proceedings, Jackson Hole Economic Policy Symposium*, 2010, pp. 85–145.
- Christiano, Lawrence J., Martin Eichenbaum, and Charles L. Evans**, “Nominal Rigidities and the Dynamic Effects of a Shock to Monetary Policy,” *Journal of Political Economy*, February 2005, 113 (1), 1–45.

- Coibion, Olivier and Yuriy Gorodnichenko**, “Information Rigidity and the Expectations Formation Process: A Simple Framework and New Facts,” *American Economic Review*, 2015, 105 (8), 2644–2678.
- Collin-Dufresne, Pierre, Michael Johannes, and Lars A. Lochstoer**, “Parameter Learning in General Equilibrium: The Asset Pricing Implications,” *American Economic Review*, March 2016, 106 (3), 664–98.
- Cordoba, Juan-Carlos and Marla Ripoll**, “Credit Cycles Redux,” *International Economic Review*, 2004, 45 (4), 1011–1046.
- Cúrdia, Vasco and Michael Woodford**, “Credit Frictions and Optimal Monetary Policy,” *Journal of Monetary Economics*, 2016, 84, 30 – 65.
- Faia, Ester and Tommaso Monacelli**, “Optimal interest rate rules, asset prices, and credit frictions,” *Journal of Economic Dynamics and Control*, 2007, 31 (10), 3228–3254.
- Flynn, Ed and Phil Crewson**, “Chapter 11 Filing Trends in History and Today,” *American Bankruptcy Institute Journal*, 2009, 14, 65 ff.
- Fuster, Andreas, Benjamin Hebert, and David Laibson**, “Natural Expectations, Macroeconomic Dynamics, and Asset Pricing,” in Daron Acemoglu and Michael Woodford, eds., *NBER Macroeconomics Annual 2011*, Vol. 26 of *NBER Chapters*, National Bureau of Economic Research, October 2012, pp. 1–48.
- Gelain, Paolo, Kevin J Lansing, and Caterina Mendicino**, “House Prices, Credit Growth, and Excess Volatility: Implications for Monetary and Macroprudential Policy,” *International Journal of Central Banking*, 2013, 9 (2), 219–276.
- Gennaioli, Nicola, Yueran Ma, and Andrei Shleifer**, “Expectations and Investment,” Working paper, Bank for International Settlements June 2016.
- Greenwood, Robin and Andrei Shleifer**, “Expectations of Returns and Expected Returns,” *Review of Financial Studies*, 2014, 27 (3), 714–746.
- Hau, Harald and Sandy Lai**, “Real effects of stock underpricing,” *Journal of Financial Economics*, 2013, 108 (2), 392 – 408.
- Hirshleifer, David, Jun Li, and Jianfeng Yu**, “Asset pricing in production economies with extrapolative expectations,” Working paper 2015.
- Jermann, Urban and Vincenzo Quadrini**, “Macroeconomic Effects of Financial Shocks,” *American Economic Review*, February 2012, 102 (1), 238–71.
- Kiyotaki, Nobuhiro and John Moore**, “Credit Cycles,” *Journal of Political Economy*, April 1997, 105 (2), 211–48.

- Kobayashi, Keiichiro and Kengo Nutahara**, “Nominal Rigidities, News-Driven Business Cycles, and Monetary Policy,” *B.E. Journal of Macroeconomics*, 2010, 10 (1), 1–26.
- Kocherlakota, Narayana R.**, “Creating business cycles through credit constraints,” *Quarterly Review*, 2000, (Sum), 2–10.
- Lawton, Anne**, “Chapter 11 Triage: Diagnosing a Debtor’s Prospects for Success,” *Arizona Law Review*, 2012, 54 (4), 985–1028.
- Liu, Zheng, Pengfei Wang, and Tao Zha**, “Land-Price Dynamics and Macroeconomic Fluctuations,” *Econometrica*, 2013, 81 (3), 1147–1184.
- Manski, Charles F.**, “Measuring Expectations,” *Econometrica*, 2004, 72 (5), 1329–1376.
- Mendoza, Enrique G.**, “Sudden Stops, Financial Crises, and Leverage,” *American Economic Review*, December 2010, 100 (5), 1941–66.
- Miao, Jianjun and Pengfei Wang**, “Bubbles and Credit Constraints,” Boston University - Department of Economics - Working Papers Series WP2011-031, Boston University - Department of Economics January 2011.
- , —, and **Zhiwei Xu**, “A Bayesian dynamic stochastic general equilibrium model of stock market bubbles and business cycles,” *Quantitative Economics*, 2015, 6 (3), 599–635.
- Milani, Fabio**, “Learning about the Interdependence between the Macroeconomy and the Stock Market,” Technical Report May 2008.
- , “Expectation Shocks and Learning as Drivers of the Business Cycle,” *Economic Journal*, May 2011, 121 (552), 379–401.
- Pancrazi, Roberto and Mario Pietrunti**, “Natural Expectations and Home Equity Extraction,” Working paper 2014.
- Pintus, Patrick A. and Jacek Suda**, “Learning Financial Shocks and the Great Recession,” AMSE Working Papers 1333, Aix-Marseille School of Economics, Marseille, France June 2013.
- Quadrini, Vincenzo**, “Financial frictions in macroeconomic fluctuations,” *FRB Richmond Economic Quarterly*, 2011, 97 (3), 209–254.
- Sargent, Thomas J.**, “Rational Expectations,” in David R. Henderson, ed., *The Concise Encyclopedia of Economics*, 2nd ed., Library of Economics and Liberty, 2008.
- Timmermann, Allan**, “Excess Volatility and Predictability of Stock Prices in Autoregressive Dividend Models with Learning,” *Review of Economic Studies*, 1996, 63 (4), 523–557.
- Warren, Elizabeth and Jay Lawrence Westbrook**, “The Success of Chapter 11: A Challenge to the Critics,” *Michigan Law Review*, February 2009, 107 (4), 604–641.

Online Appendix

A Details on the model

A.1 Setup of adjustment costs and nominal rigidities

In the full model, final good producers (indexed by $i \in [0, 1]$) 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, and sets a nominal price p_{it} for its production. A standard price adjustment friction à la Calvo means that a retailer cannot adjust her price with probability κ . Hence, the retailer solves the following optimization:

$$\begin{aligned} \max_{p_{it}} \sum_{s=0}^{\infty} \left(\prod_{\tau=1}^s \kappa \Lambda_{t+\tau} \right) ((1 + \tau) p_{it} - q_{t+s} p_{t+s}) Y_{it+s} - T_{t+s} \\ \text{s.t. } Y_{it+s} = \left(\frac{p_{it}}{p_{t+s}} \right)^{-\sigma} \tilde{Y}_{t+s}, \end{aligned}$$

where \tilde{Y}_t is aggregate demand for the composite final good and $p_t = \left(\int_0^1 p_{it}^{1-\sigma} \right)^{1/(1-\sigma)}$ is the aggregate price level. I assume that the government sets subsidies such that $\tau = 1/(\sigma - 1)$ so that the steady-state markup over marginal cost is zero, and levies a lump-sum tax T_{t+s} on retailers to finance the subsidy. Since all retailers that can re-optimize at t are identical, they all choose the same price $p_{it} = p_t^*$. The derivation of the non-linear aggregate law of motion for the retail sector is standard and the final equations are:

$$\begin{aligned} \frac{p_t^*}{p_t} &= \frac{1}{1 + \tau} \frac{\sigma}{\sigma - 1} \frac{\Gamma_{1t}}{\Gamma_{2t}} \\ \Gamma_{1t} &= q_t + \kappa \mathbb{E}_t^{\mathcal{P}} \Lambda_{t+1} \frac{\tilde{Y}_{t+1}}{\tilde{Y}_t} \pi_{t+1}^{\sigma} \\ \Gamma_{2t} &= 1 + \kappa \mathbb{E}_t^{\mathcal{P}} \Lambda_{t+1} \frac{\tilde{Y}_{t+1}}{\tilde{Y}_t} \pi_{t+1}^{\sigma-1}. \end{aligned}$$

Inflation $\pi_t = p_t/p_{t-1}$ and the reset price are linked through the price aggregation equation which can be written as

$$1 = (1 - \kappa) \left(\frac{p_t^*}{p_t} \right)^{1-\sigma} + \kappa \pi_t^{\sigma-1}$$

and the Tak-Yun distortion term is

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

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.

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:

$$\begin{aligned} \max_{w_{ht}} \mathbb{E}_t^P \sum_{s=0}^{\infty} \left(\prod_{\tau=1}^s \kappa_w \Lambda_{t+\tau} \right) ((1 + \tau_w) w_{ht} - \tilde{w}_{t+s} p_{t+s}) L_{ht+s} - T_{wt+s} \\ \text{s.t. } L_{ht} = \left(\frac{w_{ht}}{\tilde{w}_t} \right)^{-\sigma_w} \tilde{L}_t. \end{aligned}$$

Since all labor agencies that can re-optimize at t are identical, they all choose the same price $w_{ht} = w_t^*$. Again, I assume that the government sets wage subsidies $\tau = 1/(\sigma_w - 1)$ such that the steady-state markup over marginal cost is zero, and levies a lump-sum tax on labor agency profits to finance the subsidy. The first-order conditions are analogous to those for retailer, and the aggregate nominal wage level that firms face is where $w_t = \left(\int_0^1 w_{ht}^{1-\sigma_w} \right)^{1/(1-\sigma_w)}$. Wage inflation is $\pi_{wt} = w_t/w_{t-1}$ and the Tak-Yun wage distortion Δ_{wt} is defined analogously to that for final good producers.

Capital good producers operate competitively in input and output markets, producing capital goods using final consumption goods. For the latter, they have a CES aggregator just like households. There is no distinction between new and used capital and depreciation takes place within intermediate firms. The maximization program of capital producers 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. This setup is simpler than the one in [Bernanke et al. \(1999\)](#) where the price of used and new capital goods differ. The first-order condition defines the price for capital goods in the main text.

All of the profits made by the firms described above accrue to households. Similarly, all subsidies by the government are financed by lump-sum taxes on households. Taken together, the term Π_t defined in the main text is:

$$\Pi_t = \tilde{Y}_t - q_t Y_t + \tilde{w}_t L_t - w_t \tilde{L}_t + (Q_t - 1) I_t - \frac{\psi}{2} \left(\frac{I_t}{I_{t-1}} - 1 \right)^2$$

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

on the right-hand side.

$$\begin{aligned}
Y_t = \int_0^1 Y_{jt} dj &= \int_0^1 Y_{it} di \\
\tilde{Y}_t = \frac{Y_t}{\Delta_t} &= C_t + I_t + -\frac{\psi}{2} \left(\frac{I_t}{I_{t-1}} - 1 \right)^2 + C_t^f \\
L_t &= \int_0^1 L_{ht} dh \\
\tilde{L}_t = \frac{L_t}{\Delta_{wt}} &= \int_0^1 L_{jt} dj \\
K_t = \int_0^1 K_{jt} dj &= (1 - \delta) K_{t-1} + I_t \\
1 &= S_{jt}, j \in [0, 1] \\
0 &= B_t^g
\end{aligned}$$

A.2 Properties of the rational expectations equilibrium

I consider a rational expectations equilibrium with the following properties that hold in a neighborhood of the non-stochastic steady-state.

1. The expected net discounted return on capital is strictly positive for both firm owners and households: $\beta \mathbb{E}_t R_{t+1}^k > 1$ and $\mathbb{E}_t \Lambda_{t+1} R_{t+1}^k > 1$.
2. At any time t , the stock market valuation P_{jt} of a firm j is linear in its net worth, with a slope that is strictly greater than one.
3. All firms choose the same capital-labor ratio K_{jt}/L_{jt} .
4. All firms can be aggregated. Aggregate debt, capital and net worth are sufficient to describe the intermediate goods sector.
5. 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.
6. If the firm defaults and the lender seizes the firm, it always prefers restructuring to liquidation.
7. The firm always exhausts the borrowing limit.

Here, I derive restrictions on the parameters for existence of such an equilibrium. I first take the first two properties as given and show under which conditions the remaining ones hold, and then derive conditions for the first two properties be verified.

Value functions

An operating firm j enters the period 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:

$$\begin{aligned}\Upsilon_1(K_{jt-1}, B_{jt-1}, s_t) &= \max_{N_{jt}, L_{jt}, D_{jt}, Y_{jt}} \gamma N_{jt} + (1 - \gamma)(D_{jt} + \Upsilon_2(N_{jt} - D_{jt}, s_t)) \\ \text{s.t. } N_{jt} &= q_t Y_{jt} - w_t L_{jt} + (1 - \delta) Q_t K_{jt-1} - R_{t-1} B_{jt} \\ Y_{jt} &= K_{jt-1}^\alpha (A_t L_{jt})^{1-\alpha} \\ D_{jt} &= \zeta (N_{jt} - Q_t K_{jt-1} + B_{jt-1})\end{aligned}$$

The aggregate state of the economy is denoted by s_t . In what follows, I will 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:

$$\begin{aligned}\Upsilon_2(\tilde{N}, s) &= \max_{K, B, \text{strategy in default game}} \beta \mathbb{E}[\Upsilon_1(K, B, s'), \text{no default}] \\ &\quad + \beta \mathbb{E}[\Upsilon_1(K, B^*, s'), \text{debt renegotiated}] \\ &\quad + \beta \mathbb{E}[0, \text{lender seizes firm}] \\ \text{s.t. } QK &= \tilde{N} + B\end{aligned}$$

Note that, since net worth \tilde{N} is non-negative around the steady state, the firm's debt B cannot exceed its capital stock K .

In the first stage, the first order condition with respect to L equalizes the wage with the marginal revenue. Since there is no firm heterogeneity apart from capital K and debt B and the production function has constant returns to scale, this already implies Property 3 that all firms choose the same capital-labor ratio. Hence the internal rate of return on capital is common across firms:

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

Taking Property 2 as given for now, Υ_2 is a linear function

$$\Upsilon_2(\tilde{N}, s) = v_s \tilde{N}$$

with slope $v_s > 1$. Then Υ_1 is homogeneous of degree one, and at the steady state (and therefore

in a neighborhood):

$$\begin{aligned}
\mathcal{V}_1(K, B, s) &= N + (1 - \gamma)(D - N + \mathcal{V}_2(N - D, s)) \\
&= N + (1 - \gamma)(v_s - 1)((1 - \zeta)N + \zeta(QK - B)) \\
&> N = R^k K - RB.
\end{aligned}$$

Limited commitment problem

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 (in equilibrium $\tilde{R} = R$ but this is to be established first). The game is played sequentially:

1. The firm (F) proposes a borrowing contract (B, \tilde{R}) .
2. The lender (L) can accept or reject the contract.
 - A rejection corresponds to setting the contract $(B, \tilde{R}) = (0, 0)$.
Payoff for L: 0. Payoff for F: $\beta \mathbb{E} [\mathcal{V}_1(\tilde{N}, 0, s')]$.
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: $\mathbb{E} \Lambda \tilde{R} B - B$. Payoff for F: $\beta \mathbb{E} [\mathcal{V}_1(K, \frac{\tilde{R}}{R} B, s')]$.
4. If F defaults, the debt needs to be renegotiated. F makes an offer for a new debt level B^* .²⁶
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: $\beta \mathbb{E} [\mathcal{V}_1(K, \frac{\tilde{R}}{R} B^*, s')]$.
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: $\mathbb{E} [\Lambda Q'] \xi K - B$. Payoff for F: 0.
 - 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: $\xi B + \beta \mathbb{E} [\mathcal{V}_1(\xi K, \xi B, s')] - B$. Payoff for F: 0.

²⁶That the interest rate on the repayment is fixed is without loss of generality.

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

$$\xi B + \beta \mathbb{E} [\gamma_1 (\xi K, \xi B, s')] \geq \mathbb{E} \Lambda \xi Q' K.$$

This holds true at the steady state, as we have $\beta R^k > 1$ (Property 1), $\Lambda = \beta = 1/R$ and $Q' = 1$:

$$\begin{aligned} \beta \gamma_1 (\xi K, \xi B, s) &> \beta (R^k \xi K - R \xi B) \\ &> \xi (K - \tilde{\beta} R B) \\ &> \xi (\beta K - B). \end{aligned}$$

Since the inequality is strict, it holds around the steady-state as well. This establishes Property 6.

Next, L will accept an offer B^* if it gives her a better expected payoff (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

$$\mathbb{E} [\Lambda] \tilde{R} B^* \geq x (\xi B + \tilde{\beta} \mathbb{E} [\gamma_1 (\xi K, \xi B, s')]) + (1 - x) \mathbb{E} [\Lambda Q'] \xi K.$$

Now turn to the firm F. Among the set of offers B^* that are accepted by L, the firm will prefer the lowest one which satisfies the above restriction with equality. This follows from γ_1 being a decreasing function of debt. This lowest offer will be made if it leads to a higher payoff than expropriation: $\beta \mathbb{E} [\gamma_1 (K, \frac{\tilde{R}}{R} B, s')] \geq 0$. 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 B^* can be set smaller than B or if expropriation is better than repaying, $\beta \mathbb{E} [\gamma_1 (K, \frac{\tilde{R}}{R} B, s')] \geq 0$.

What is 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 (B, \tilde{R}) 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 for any non-negative B that does not trigger default, since at the steady-state $R^k > 1/\beta > R$ and therefore

$$\begin{aligned} \beta \mathbb{E} [\gamma_1 (K, B, s')] &> \beta \mathbb{E} [R^k K - R B] \\ &= \beta \mathbb{E} [R^k \tilde{N} + (R^k - R) B] \\ &> 0. \end{aligned}$$

F therefore prefers this contract to one that leads to default with expropriation. The payoff is increasing in the size of the loan B , since

$$\begin{aligned} & \frac{\partial}{\partial B} \mathbb{E} \left[\Upsilon_1 \left(\frac{\tilde{N} + B}{Q}, B, s' \right) \right] \\ &= \beta \mathbb{E} \left[\frac{R^k}{Q} - R + (1 - \gamma) (v_{s'} - 1) \left((1 - \zeta) \left(\frac{R^k}{Q} - R \right) + \zeta \left(\frac{Q'}{Q} - 1 \right) \right) \right] \\ &> 0. \end{aligned}$$

Therefore, of all values for B that do not lead to default, F will want to choose the largest one, defined as:

$$\bar{B} = \max \left\{ B \left| x \left(\xi B + \beta \mathbb{E} \left[\Upsilon_1 \left(\xi \left(\frac{\tilde{N} + B}{Q}, s' \right), \xi B, s' \right) \right] \right) + (1 - x) \mathbb{E} [\Lambda Q'] \xi \frac{\tilde{N} + B}{Q} - B \geq 0 \right. \right\}.$$

In order for the borrowing constraint to be binding, it must be finite. Since the set above contains $B = 0$, this amounts to the condition that

$$x \left(1 + \beta \frac{\partial}{\partial B} \mathbb{E} \left[\Upsilon_1 \left(\frac{\tilde{N} + B}{Q}, B, s' \right) \right] \right) + (1 - x) \mathbb{E} [\Lambda Q'] < \frac{1}{\xi} \quad (41)$$

which is satisfied for ξ small enough. Because Υ_1 is homogenous of degree one, the borrowing limit is linear in \tilde{N} and can be written as $\bar{B} = v_{B,s} \tilde{N}$.

F could also offer a contract (B, \tilde{R}) that only leads to a default with debt renegotiation. The optimal contract of this type is the solution to the following problem:

$$\begin{aligned} & \max_{\tilde{R}, B, B^*} \beta \mathbb{E} \left[\Upsilon_1 \left(\tilde{N} + B, \frac{\tilde{R} B^*}{R}, s' \right) \right] \\ \text{s.t. } & \frac{\tilde{R} B^*}{R} \geq B \\ & \frac{\tilde{R} B^*}{R} = x \left(\xi B + \beta \mathbb{E} \left[\Upsilon_1 \left(\xi \frac{\tilde{N} + B}{Q}, \xi B, s' \right) \right] \right) \\ & \quad + (1 - x) \mathbb{E} [\Lambda Q'] \xi \frac{\tilde{N} + B}{Q} \end{aligned}$$

It is clear that the value of this problem is solved by setting $\tilde{R} = R$ and $B = B^* = \bar{B}$, which amounts to not defaulting. This establishes Properties 5 and 7.

Linearity of firm value

Since firms do not default and exhaust the borrowing limit \bar{B} , the second-stage firm value is

$$\begin{aligned}\mathcal{V}_2(\tilde{N}) &= \beta \mathbb{E} \left[\mathcal{V}_1 \left(\frac{\tilde{N} + \bar{B}}{Q}, \bar{B}, s' \right) \right] \\ &= \beta \mathbb{E} \left[\mathcal{V}_1 \left(\frac{1 + v_{B,s}}{Q} \tilde{N}, v_{B,s} \tilde{N}, s' \right) \right] \\ &= \beta \mathbb{E} \left[\mathcal{V}_1 \left(\frac{1 + v_{B,s}}{Q}, v_{B,s}, s' \right) \right] \tilde{N}.\end{aligned}$$

We have therefore verified the linearity of \mathcal{V}_2 . To establish Property 2, it remains to show that the slope of \mathcal{V}_2 is greater than one. At the steady state:

$$\begin{aligned}v_s &= \beta \mathcal{V}_1(1 + v_{B,s}, v_{B,s}, s) \\ &= \beta \left(R^k + \underbrace{v_{B,s}(R^k - R)}_{>0} \right) \underbrace{\left(1 + (1 - \gamma) \overbrace{(v_s - 1)(1 - \zeta)}^{>-1} \right)}_{>0} + (1 - \gamma)(v_s - 1)\zeta \\ &> \beta R^k (1 + (1 - \gamma)(v_s - 1)(1 - \zeta)) + (1 - \gamma)(v_s - 1)\zeta \\ &> 1 + (1 - \gamma)(v_s - 1) \\ &> 1.\end{aligned}$$

Finally, the aggregated law of motion for capital and net worth needs to be established (Property 4). Denoting again by $\Gamma_t \subset [0, 1]$ the indices of firms that are alive at the end of period t , we have

$$\begin{aligned}Q_t K_t = Q_t \int_0^1 K_{jt} dj &= \int_{j \in \Gamma_t} (N_{jt} - \zeta E_{jt} + B_{jt}) dj \\ &= (1 - \gamma)(N_t - \zeta E_t) + B_t \\ N_t = \int_0^1 N_{jt} dj &= R_t^k K_{t-1} - R_{t-1} B_{t-1} \\ B_t = \int_0^1 B_{jt} dj &= x \xi (B_t + P_t) + (1 - x) \xi \mathbb{E}_t \Lambda_{t+1} Q_{t+1} K_t.\end{aligned}$$

Return on capital

We can now establish a condition under which $\beta R^k > 1$ holds (Property 1). From the aggregate equations above, and the definition of earnings $E = N - QK + B$, it follows that in steady state:

$$R^k = \frac{RB}{K} + \frac{1 - \zeta(1 - \gamma)}{(1 - \zeta)(1 - \gamma)} \left(1 - \frac{B}{K} \right). \quad (42)$$

Rearranging the above expression, one obtains that $\beta R^k > 1$ holds at the steady state if and only if:

$$\gamma > 1 - \frac{\beta}{1 - \zeta(1 - \beta)}. \quad (43)$$

A.3 Conditions to rule out multiple equilibria

Collateral constraints often give rise to multiple equilibria due to their feedback effects: Low asset prices reduce borrowing constraints and activity, which in depress asset prices and so on. This multiplicity appears even in the very early literature (Kiyotaki and Moore, 1997). More recently, Miao and Wang (2011) have shown that when firm borrowing constraints depend on equity value, multiple steady states are possible. In this section, I give conditions under which this type of multiplicity does not arise. These conditions are satisfied for the parameter values at which the model is simulated.

Miao and Wang look for an equilibrium in which firm value $\Upsilon_2(\tilde{N}, s)$ is not linear but affine in net worth \tilde{N} . Even a firm with zero net worth has positive value. This can be an equilibrium: The positive equity value enables the firm to borrow, acquire capital and pay dividends from the returns; those expected dividends can justify the positive equity value.

Suppose that $\Upsilon_2(\tilde{N}, s) = v_s \tilde{N} + \vartheta_s$ with $\vartheta_s \geq 0$ and $v_s > 1$, and that $\beta \mathbb{E}_t R_{t+1}^k > 1$. Then the proof for the existence of an equilibrium satisfying properties 3–7 above still goes through under the same conditions (41) and (43). The equation determining the coefficient ϑ_s is:

$$v_s \tilde{N} + \vartheta_s = \beta \mathbb{E} \left[\kappa_{N,s'} \tilde{N} + \kappa_{B,s'} \bar{B} + (1 - \gamma) \vartheta_{s'} \right]$$

where

$$\begin{aligned} \kappa_{N,s'} &= (1 + (1 - \gamma)(v_{s'} - 1)(1 - \zeta)) \frac{R^k}{Q} - (1 - \gamma)(v_{s'} - 1)\zeta \frac{Q'}{Q} \\ \kappa_{B,s'} &= (1 + (1 - \gamma)(v_{s'} - 1)(1 - \zeta)) \left(\frac{R^k}{Q} - R \right) - (1 - \gamma)(v_{s'} - 1)\zeta \left(\frac{Q'}{Q} - 1 \right). \end{aligned}$$

The borrowing limit \bar{B} depends itself on equity value and therefore on the coefficients v_s and ϑ_s :

$$\bar{B} = \frac{\mathbb{E} \left[\xi x \beta (1 - \gamma) \vartheta_{s'} + \xi \tilde{N} \left((1 - x) \Lambda \frac{Q'}{Q} + x \beta \kappa_{N,s'} \right) \right]}{1 - \xi \mathbb{E} \left[x + (1 - x) \Lambda \frac{Q'}{Q} + x \beta \kappa_{B,s'} \right]}. \quad (44)$$

Comparing coefficients, the equations determining v_s and ϑ_s are:

$$v_s = \beta \mathbb{E} [\kappa_{N,s'}] + \beta \frac{\xi \mathbb{E} \left[(1-x) \Lambda \frac{Q'}{Q} + x \beta \kappa_{N,s'} \right]}{1 - \xi \mathbb{E} \left[x + (1-x) \Lambda \frac{Q'}{Q} + x \beta \kappa_{B,s'} \right]} \mathbb{E} [\kappa_{B,s'}]$$

$$\vartheta_s = \beta (1-\gamma) \left(1 + \frac{\xi x \mathbb{E} [\kappa_{B,s'}]}{1 - \xi \mathbb{E} \left[x + (1-x) \Lambda \frac{Q'}{Q} + x \beta \kappa_{B,s'} \right]} \right) \mathbb{E} [\vartheta_{s'}].$$

Clearly, $\vartheta_s \equiv 0$ is a solution to the second equation and corresponds to the equilibrium considered in this paper. It is the unique solution if the term multiplying $\mathbb{E} [\vartheta_{s'}]$ is always strictly smaller than one. Around a steady state in which ϑ_s is zero, a sufficient condition to guarantee uniqueness is that

$$\beta (1-\gamma) \left(1 + \frac{\xi x \beta \kappa_{B,s}}{1 - \xi x + (1-x) \beta + x \beta \kappa_{B,s}} \right) < 1. \quad (45)$$

This always holds for x small enough. It remains to establish conditions under which a steady state with $\vartheta_s > 0$ can also be ruled out. Such a steady state would necessarily have the term multiplying $\mathbb{E} [\vartheta_{s'}]$ equal to one at the steady state. From this, it follows that necessarily

$$\begin{aligned} \kappa_{B,s} &= \frac{1 - \beta (1-\gamma)}{\xi \beta x} (1 - \xi x - \xi (1-x) \beta) \\ v_s &= \frac{1 - \beta (1-\gamma)}{\xi \beta x} \frac{1 - \xi x}{1 - \gamma} \\ R^k &= R + \frac{\kappa_{B,s}}{1 + (1-\gamma) (1-\zeta) (v_s - 1)} \end{aligned} \quad (46)$$

holds at the steady state. Now, note that the values of $\kappa_{B,s}$ and $\kappa_{N,s}$ at the steady-state are do not depend on ϑ_s , and that therefore the equilibrium borrowing limit \bar{B} in (44) is increasing in ϑ_s for any level of \tilde{N} . In particular then, equilibrium leverage B/K is also an increasing in ϑ_s . Since the equilibrium return on capital is a decreasing function of leverage through Equation (42), the steady state with $\vartheta_s > 0$ has a lower R^k than in the steady state with $\vartheta_s = 0$. A sufficient condition to guarantee that $\vartheta_s = 0$ is the unique steady state is therefore that the corresponding steady-state value of R^k is higher than the one computed in (46).

A.4 Proof of the limit in (36)

Combine Equation (33) and (34) to obtain:

$$\frac{1-\xi}{\xi} \left(R - \exp \left(\hat{\mu}_t + \frac{1}{2} \sigma_\mu^2 \right) \right) = \alpha \left(\frac{A_t}{K_t} \right)^{1-\alpha} \mathbb{E}_t [\varepsilon_{t+1}^{1-\alpha}] + 1 - \delta - R.$$

It follows that:

$$\begin{aligned}
\Delta \log P_t &= \Delta \log K_t \\
&= \Delta \log A_t - \frac{1}{1-\alpha} \Delta \log \left(\frac{1-\xi}{\xi} \left(R - \exp \left(\hat{\mu}_t + \frac{1}{2} \sigma_\mu^2 \right) \right) - 1 + \delta + R \right) \\
&\xrightarrow{\xi \rightarrow 1} \Delta \log A_t = \varepsilon_t.
\end{aligned}$$

B Solving for the learning equilibrium

This section describes how to construct and solve for the learning equilibrium with conditionally model-consistent expectations.

B.1 General formulation

To set the notation, I start with the solution of the standard rational expectations equilibrium. Denote the n endogenous model variables by y_t and the n_u exogenous shocks by u_t . The exogenous shocks are independent across time with joint distribution F_σ , mean zero and variance $\sigma^2 \Sigma_u$. The solution of a (recursive) rational expectations equilibrium satisfies the equilibrium conditions:

$$\mathbb{E}_t [f_{-P}(y_{t+1}, y_t, y_{t-1}, u_t)] = 0 \quad (47)$$

$$\mathbb{E}_t [f_P(y_{t+1}, y_t)] = 0 \quad (48)$$

where f_P denotes the stock market clearing condition (16) and f_{-P} collects the remaining $n-1$ equilibrium conditions.²⁷ A recursive solution takes the form:

$$y_t = g_{RE}(y_{t-1}, u_t, \sigma).$$

By the definition of rational expectations, the expectations in (47)–(48) are taken under the probability measure induced by g_{RE} and F_σ , so that the policy function g_{RE} itself can be found by solving:

$$\int f_{-P} \left(\begin{array}{c} g_{RE}(g_{RE}(y_{t-1}, u_t, \sigma), u_{t+1}, \sigma), \\ g_{RE}(y_{t-1}, u_t, \sigma), y_{t-1}, u_t \end{array} \right) dF_\sigma(u_{t+1}) = 0 \quad (49)$$

$$\int f_P \left(\begin{array}{c} g_{RE}(g_{RE}(y_{t-1}, u_t, \sigma), u_{t+1}, \sigma), \\ g_{RE}(y_{t-1}, u_t, \sigma) \end{array} \right) dF_\sigma(u_{t+1}) = 0. \quad (50)$$

In the learning equilibrium, the probability measure \mathcal{P} used by agents to form expectations

²⁷There are usually $n+1$ equilibrium conditions in total, but one of the market clearing conditions is redundant due to Walras' law. While under rational expectations, it is immaterial for the computation of the equilibrium which market clearing condition is left out, this choice can matter when constructing the learning equilibrium with conditionally model-consistent expectations. Here I choose to omit the market clearing condition for final consumption goods.

does not coincide with the actual probability measure describing the equilibrium outcomes. In particular, agents are not endowed with the knowledge that the stock price is determined by the market clearing condition f_P , and instead they form expectations about future prices using a subjective law of motion. This law of motion can be summarized in a function:

$$\phi(\tilde{y}_t, \tilde{y}_{t-1}, z_t) = \begin{pmatrix} \Delta \log P_t - \hat{\mu}_{t-1} - z_t \\ \Delta \hat{\mu}_t - g z_t \end{pmatrix} = 0$$

where $\tilde{y}_t = (y_t, \hat{\mu}_t)$ incorporates the belief state introduced by the learning process, and z_t is the *subjective forecast error*. In the mind of agents, this forecast error is an exogenous iid shock with distribution G_σ , mean zero and variance $\sigma^2 \Sigma_z$. I assume that agents believe that z and u are mutually independent as well.

I impose discipline on the expectation formation process by requiring that agents have conditionally model-consistent expectations, as defined in the main text. I find such expectations by computing a *subjective policy function*

$$\tilde{y}_t = h(\tilde{y}_{t-1}, u_t, z_t, \sigma)$$

which satisfies:

$$\int f_{-P} \begin{pmatrix} Ch(h(\tilde{y}_{t-1}, u_t, z_t, \sigma), u_{t+1}, z_{t+1}, \sigma), \\ Ch(\tilde{y}_{t-1}, u_t, z_t, \sigma), C\tilde{y}_{t-1}, u_t \end{pmatrix} dF_\sigma(u_{t+1}) dG_\sigma(z_{t+1}) = 0 \quad (51)$$

$$\phi(h(\tilde{y}_{t-1}, u_t, z_t, \sigma), \tilde{y}_{t-1}, z_t) = 0. \quad (52)$$

Here, the matrix C just selects the original model variables y_t from the augmented vector \tilde{y}_t ($y_t = C\tilde{y}_t$). Solving for h effectively amounts to solving a different rational expectations model in which the market clearing condition for the stock market is replaced by the subjective law of motion for stock prices. Once computed, the policy function h together with F_σ and G_σ defines a complete internally consistent probability measure \mathcal{P} on all endogenous model variables. Under \mathcal{P} , agents believe that the stock price follows the subjective law of motion ϕ , and \mathcal{P} also satisfies the equilibrium conditions f_{-P} :

$$\mathbb{E}_t^\mathcal{P} [f_{-P}(y_{t+1}, y_t, y_{t-1}, u_t)] = 0.$$

This subjective belief is very close to rational expectations and preserves as much as possible of its forward-looking, model-consistent logic while allowing for subjective expectations about stock prices.

Now, the subjective policy function h depends on the subjective forecast error z_t , which under \mathcal{P} is believed to be a white noise process. In equilibrium however, z_t is instead determined endogenously by the equilibrium stock price that clears the stock market. That is, the equilibrium value

of the subjective forecast error is itself a function of the states and the shocks:

$$z_t = r(\tilde{y}_{t-1}, u_t, \sigma) \quad (53)$$

The function r can be computed by imposing equilibrium in the stock market, represented by the equation $\mathbb{E}_t^{\mathcal{P}} [f_P(\tilde{y}_{t+1}, \tilde{y}_t, \tilde{y}_{t-1}, z_t)] = 0$. Substituting the functional forms:

$$\int \psi \left(\begin{array}{c} h(h(\tilde{y}_{t-1}, u_t, r(\tilde{y}_{t-1}, u_t, \sigma), \sigma), u_{t+1}, z_{t+1}, \sigma), \\ h(\tilde{y}_{t-1}, u_t, r(\tilde{y}_{t-1}, u_t, \sigma), \tilde{y}_{t-1}, z_t, \sigma) \end{array} \right) dF_{\sigma}(u_{t+1}) dG_{\sigma}(z_{t+1}) = 0. \quad (54)$$

Note that while the current value of the forecast error z_t was substituted out, the future value z_{t+1} was not substituted out, as this value is still taken under the subjective expectation \mathcal{P} which treats it as an exogenous random disturbance.

The final equilibrium of the model is described by the *objective policy function*:

$$\tilde{y}_t = g(\tilde{y}_{t-1}, u_t, \sigma) = h(\tilde{y}_{t-1}, u_t, r(\tilde{y}_{t-1}, u_t, \sigma), \sigma).$$

By construction, this policy function satisfies all equilibrium conditions of the model. This function g together with F_{σ} defines the equilibrium probability distribution of the model variables. It differs from the subjective distribution \mathcal{P} only in that under \mathcal{P} , z_t is an unpredictable exogenous shock, whereas in equilibrium z_t is a function of the state variables and the structural shocks u_t .

It is straightforward to see that the expectations thus constructed satisfy conditional model-consistency:

$$\begin{aligned} \mathbb{E}_t^{\mathcal{P}} [\tilde{y}_{t+1} \mid u_{t+1}, P_{t+1}] &= h(\tilde{y}_t, u_{t+1}, z_{t+1}, \sigma) \\ &= h(\tilde{y}_t, u_{t+1}, r(\tilde{y}_t, u_{t+1}, \sigma), \sigma) \\ &= g(\tilde{y}_t, u_{t+1}, \sigma) \\ &= \tilde{y}_{t+1}. \end{aligned}$$

B.2 Approximation with perturbation methods

I now describe how to compute an approximation of the objective policy function g around the non-stochastic steady state \bar{y} . The procedure has two steps and does not require iteration. The first step consists in deriving a perturbation approximation of the subjective policy function h . This can be done using standard methods, as the system of equations (51)–(52) can be solved as if it were a standard rational expectations model. The second step consists in finding the derivatives of the function r . Applying the implicit function theorem to Equation (54), one can compute the

first-order derivatives as:

$$\begin{aligned} r_y &= -A^{-1} \left(\left(\frac{\partial \psi}{\partial \tilde{y}_{t+1}} h_y + \frac{\partial \psi}{\partial \tilde{y}_t} \right) h_y + \frac{\partial \psi}{\partial \tilde{y}_{t-1}} \right) \\ r_u &= -A^{-1} \left(\frac{\partial \psi}{\partial \tilde{y}_{t+1}} h_y + \frac{\partial \psi}{\partial \tilde{y}_t} \right) h_u \\ r_\sigma &= -A^{-1} \left(\frac{\partial \psi}{\partial \tilde{y}_{t+1}} h_y + \frac{\partial \psi}{\partial \tilde{y}_t} \right) h_\sigma \end{aligned}$$

where the matrix A is given by $A = \left(\frac{\partial \psi}{\partial \tilde{y}_{t+1}} h_y + \frac{\partial \psi}{\partial \tilde{y}_t} \right) h_z + \frac{\partial \psi}{\partial z_t}$. This matrix needs to be invertible for the learning equilibrium to exist. The first-order derivatives of the actual policy function g can be obtained by applying the chain rule:

$$\begin{aligned} g(\tilde{y}_{t-1}, u_t, \sigma) &\approx g(\bar{y}, 0, 0) + g_y(\tilde{y}_{t-1} - \bar{y}) + g_u u_t + g_\sigma \sigma \\ g_y &= h_y + h_z r_y \\ g_u &= h_u + h_z r_u \\ g_\sigma &= h_\sigma + h_z r_\sigma \end{aligned}$$

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.

Second- and higher-order perturbation approximations of g can be computed analogously. As in first order, only invertibility of the matrix A is required for a unique local solution under learning.